

BRIAN A. JACKSON, VIKRAM KILAMBI, DAVID R. FRELINGER, THOMAS LIGHT,
AISHA NAJERA CHESLER, PAUL EMSLIE, ANTHONY LAWRENCE

Using Modeling and Simulation to Advance Effects-Based Security Forces Planning

Developing Prototype Approaches to Estimate Risk Reduction Across Security Missions



For more information on this publication, visit www.rand.org/t/RR1754-1.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.

© 2024 RAND Corporation

RAND® is a registered trademark.

Library of Congress Cataloging-in-Publication Data is available for this publication.

ISBN: 978-1-9774-1293-5

Cover: Adapted from Eugene Oliver/U.S. Air Force.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

About This Report

U.S. Air Force Security Forces (SF) are responsible for detecting, deterring, denying, and defeating current and future security threats in both traditional and expeditionary basing environments. At one end of the threat spectrum, the SF must address everyday threats that are very similar to those that a civilian law enforcement organization responds to on a day-to-day basis. At the other end of the spectrum are threats that are distinctly military in nature, focused on protecting bases and the assets and personnel located on those bases from potentially highly capable attackers. The SF must also be ready to deploy and support expeditionary operations in wartime while maintaining protection at garrison locations, which could be under greater risk of threat and attack.

Recognizing the varied risks that must be managed to protect personnel and assets on air bases requires an evolution in the processes and doctrine that support security operations. This report explores how modeling and simulation can be used as part of an effects-based security force and how the base defense planning process might be used to allocate resource and balance risks in a transparent and systematic way. We conducted this research in the context of SF operations at Air Force Global Strike Command (AFGSC) bases, which face unique challenges associated with the nuclear mission. We developed prototype models to explore several example scenarios with respect to various SF risks facing the AFGSC, with the scenarios informed by both long-term planning challenges facing the SF and key questions and concerns raised in our discussions with SF subject-matter experts during the study.

The research reported here was commissioned by the Director of Operations and Communications, Headquarters Air Force Global Strike Command (AFGSC/A3/6) and conducted within the Workforce, Development, and Health Program of RAND Project AIR FORCE as part of a fiscal year 2022 project “Asset Security Posture.”

RAND Project AIR FORCE

RAND Project AIR FORCE (PAF), a division of RAND, is the Department of the Air Force’s (DAF’s) federally funded research and development center for studies and analyses, supporting both the U.S. Air Force and the United States Space Force. PAF provides the DAF with independent analyses of policy alternatives affecting the development, employment, combat readiness, and support of current and future air, space, and cyber forces. Research is conducted in four programs: Strategy and Doctrine; Force Modernization and Employment; Resource Management; and Workforce, Development, and Health. The research reported here was prepared under contract FA7014-22-D-0001.

Additional information about PAF is available on our website:

www.rand.org/paf/

This report documents work originally shared with the DAF on September 27, 2022. The draft report, dated September 2022, was reviewed by formal peer reviewers and DAF subject-matter experts.

Acknowledgments

This research project was sponsored by Brigadier General Michael A. Miller, director of operations and communications, Headquarters Air Force Global Strike Command (AFGSC/A3/6). Major Cory Cook and Curtis Wells were the project's technical liaisons and provided excellent research support.

This research would not have been possible without the support of staff from various organizations who assisted us during our research, including Capt Michael Thrasher (5 Security Forces Squadron), TSgt Daren Marshall (791 Missile Security Forces Squadron), SMSgt Jaime Rivera (377 Weapons System Security Squadron), Capt Dylan Foster (377 Security Forces Squadron), Capt Mark Rapala (7 Security Forces Squadron), and Bridgette Simpkins (AF/A4SR). We also acknowledge all the other SF subject-matter experts we spoke to during the project in both virtual and in-person interviews. The insights they provided were critical to our framing of the issues and to the design of the prototype analytic tools we developed in the study.

We also want to acknowledge various colleagues at RAND who influenced and contributed to this research. Scott Savitz provided the CATAPULTA model and his knowledge and expertise using it, which allowed us to explore a type of SF risk that would not have otherwise been possible. Heather Williams shared unpublished work on incidents of radicalization among U.S. military personnel, which helped us define and illustrate that particular component of risk the SF must address. Edward Parker consulted with us on a facet of our modeling. We thank our program director, Nelson Lim, and associate program director, Miriam Matthews, for their support and guidance on this research. Lisa Harrington and Albert Robbert provided excellent input and guidance as this project progressed. Scott Savitz and Shannon Caudill provided valuable input during the peer review process for the report.

Summary

Issue

U.S. Air Force (USAF) Security Forces (SF) are responsible for detecting, deterring, denying, and defeating security threats in both garrison and expeditionary environments. The SF must address everyday threats that are very similar to those that a civilian law enforcement organization responds to on a day-to-day basis. Simultaneously, the SF must address threats that are distinctly military in nature: protecting bases from potentially highly capable attackers. Air Force Global Strike Command (AFGSC) has the added requirement of protecting critical nuclear weapons assets, as well as the aircraft and personnel involved in the management and deployment of those assets.

Although existing security planning supports detailed bottom-up, asset-based security planning (e.g., the current *Enterprise Protection Risk Management* process and supporting tools), such processes do not fully explore the risk trade-offs that are associated with different security strategies, nor do they identify opportunities for SF strategies to manage multiple risks simultaneously.

Approach

This report summarizes research that explored how top-down risk analysis models could help inform decisions regarding SF staffing, systems, and strategies. Given the project's scale, the scope was limited. The goal was not to build a fully fledged risk tool that could be immediately applied to USAF SF personnel and other planning but to take a substantial step in building the foundation for such a tool. During this research, we modeled and explored five different scenarios that involved risks to USAF (1) personnel, systems, and facilities and (2) nuclear assets and sites. Those risks were

- incursions within a base perimeter
- missile field operations and security threats
- active shooter threats
- outside-the-perimeter indirect fire or small unmanned aircraft systems threats
- routine law enforcement incidents.

The process of building the models was informed by existing analysis of USAF SF concerns, interviews with SF subject-matter experts (SMEs) at AFGSC bases, and analysis of USAF data provided to us.

To minimize the need for significant amounts of base-specific sensitive information during this prototyping effort, we applied the models to a notional AFGSC base. The models are relatively simple and transparent, representing SF personnel and their capabilities to perform their roles to protect against each of the five threat scenarios. The models were challenged with security demands or incidents of different characteristics, from which the goal was to explore how the risk of negative

outcomes changes with different levels of SF manning, training, and equipping or with different assumptions about base infrastructure or technical protections.

Conclusions

The prototype tools that were developed in this project demonstrate approaches to estimate the effects of the SF on those different risks. Some takeaways from this research include the following:

- **Modeling can show SF outcomes as different types of risk reduction, which can better inform effects-based planning.** The fact that the SF manages a wide variety of risks makes assessing the full consequences of manning, technology, and other changes in practice more difficult. Models, such as those developed in this research, can show how choices affect day-to-day risks (e.g., responding to routine law enforcement incidents), and potentially much higher-consequence incidents, to allow more-informed choices and to make consequences clearer.
- **Models can demonstrate the tangible effects of key SF concerns.** In our work, SF SMEs raised concerns about the potential for stress to reduce SF effectiveness in protecting USAF assets and personnel. Modeling can make the potential effects of such abstract issues more tangible and better inform decisions.
- **Models can show trade-offs among different risk types.** In the missile field, the SF protect silos from incursion and are responsible for securing maintenance operations on missile systems. If personnel numbers are constrained, trade-offs might have to be made between managing adversary risk and readiness risk from deferred maintenance operations. Tools, such as these prototype models, can make those trade-offs explicit, better defining how security decisions could be forced to trade off one type of risk for another.
- **Risk modeling can help explore how the SF in one role can reduce multiple risks simultaneously.** SF personnel in different posts can play a role in managing multiple threats. For example, neutralizing a base incursion attempt on nuclear armed platforms would draw on SF personnel involved in gate security and mobile units conducting daily, routine law enforcement operations, in addition to providing dedicated SF to those assets. A system- or effects-based approach to security planning, such as modeled in this report, provides an approach for estimating the effectiveness of those combined forces, potentially allowing for strategies that better address multiple types of risks simultaneously.

The USAF integrated defense approach has moved from a compliance-based to an effects-based framing for security planning and execution. This research demonstrates how top-down, outcome-based risk analysis tools could contribute to that effort, better informing planning and risk tolerance decisions made by commanders at Air Force bases at home and abroad. Future efforts could build on this foundation to use more-detailed combat models, modeling the SF response to different risks simultaneously, and to use other innovations to link these types of tools to the current Enterprise Protection Risk Management process and related planning efforts. That the basic prototypes developed here could successfully examine multiple key SF planning challenges suggests that continued effort to improve tools for effects-based planning is warranted and could make a tangible contribution to manpower, planning, and resource allocation decisionmaking.

Contents

About This Report.....	iii
Summary	v
Figures and Tables	ix
CHAPTER 1.....	1
Introduction	1
Threat Landscape for U.S. Air Force Base Security	3
Anticipating Future Threats.....	8
Effects-Based Security Planning for Security Forces Missions	8
Current Challenges for Air Force Global Strike Command Security Risk Management	10
Study Approach.....	11
Organization of This Report	13
CHAPTER 2.....	14
Prototyping Tools to Better Support Effects-Based Security Planning.....	14
Developing Risk Analysis Tools Using a Hypothetical Air Force Global Strike Command Base.....	16
Modeling Security Forces Risk Management Across Different Threats to the Notional Air Force Base	19
Chapter Summary	39
CHAPTER 3.....	41
Exploring Example Scenarios and Challenges Using the Prototype Tools	41
Incursion Response Risk	41
Missile Field Security Risk.....	48
Active Shooter Risk	54
Indirect Fire and Small Unmanned Aerial System Attack Risk	58
Law Enforcement Response Risk.....	66
Chapter Summary	74
CHAPTER 4.....	75
Conclusions	75
Outcome-Focused Analysis Can Support Effects-Based Security Forces Planning.....	75
Considering Future Steps and Applications	79
APPENDIX A.....	81
POLICE Law Enforcement Incident Response Risk Model.....	81
APPENDIX B	94
DEFENSE Base Incursion Response Risk Model.....	94
APPENDIX C.....	108
SILO PROTECT Missile Field Security Forces Risk Model	108
APPENDIX D.....	114
CATAPULTA Indirect Fire Model	114

Abbreviations	116
References.....	117

Figures and Tables

Figures

Figure 2.1. Different U.S. Air Force Assets Protected by Security Forces and the Risk Scenarios Included in the Model Prototyping Effort.....	16
Figure 2.2. Map of Notional Air Force Base.....	17
Figure 2.3. Example Red Incursion into Notional Air Force Base via Roadways from ECHO Gate to the Aircraft Parking Apron	23
Figure 2.4. Notional Air Force Base Intercontinental Ballistic Missile Field.....	26
Figure 2.5. Base Security Zone of Notional Air Force Base Relevant to Indirect Fire and Small Unmanned Aerial System Threats	31
Figure 2.6. Incident Response at Notional Air Force Base.....	34
Figure 2.7. Modeling Cycle and Times Simulated.....	36
Figure 2.8. Different U.S. Air Force Assets Protected by the Security Forces, Risk Scenarios Included in the Prototyping Effort, and Types of Security Forces Risk Reduction Explored.....	40
Figure 3.1. Modeled Direct Assault Scenario Schematic	43
Figure 3.2. Probability of at Least 1 Red Opposing Force Member Reaching Their Target by Red Force Size	44
Figure 3.3. Average Number of Red Opposing Force Members Reaching Their Target by Red Force Size	44
Figure 3.4. Modeled Clandestine Infiltration Scenario Schematic (Low Detection)	46
Figure 3.5. Chance of Red Neutralization Before Reaching Weapons Storage Area	47
Figure 3.6. Average Surviving Red Personnel Reaching Weapons Storage Area	47
Figure 3.7. Effect of Security Forces Squadron Size on Time Before Different Incident Types Can Be Fully Staffed (Including Dispatch and Travel Time)	50
Figure 3.8. Major Maintenance Events Delayed or Postponed Because of Manpower Constraints	51
Figure 3.9. Variation in Expected Travel Time from Dispatch to Incident Locations Within Two Adjacent Ten Silo/One Missile Alert Facility Portion of the Missile Field by Road Surface	52
Figure 3.10. Change in Security Forces Time to Scene for Different Percentages of Road Surface in the Missile Field	53
Figure 3.11. Modeled Active Shooter Scenario Schematic.....	56
Figure 3.12. Effect of Modeled Security Forces Proficiency on the Average Duration of Active Shooter Scenario	57
Figure 3.13. Expected Number of Engagements Between Attacker and Defenders with Increasing Modeled Security Forces Proficiency	58
Figure 3.14. Representative Model Runs of Mortar Impact Points in Simulations for High and Low Accuracy, Dense and Sparse Target Sets	61
Figure 3.15. Average Aircraft Receiving Direct Hits and Severe and Moderate Damage by Mortar Accuracy and Target Set.....	62
Figure 3.16. Illustration of a Smaller, Critical Cross Section of a Larger Simulated Aircraft.....	64
Figure 3.17. Average Expected Direct Hits by Direct Attack Small Unmanned Aerial Systems at 1-Meter Circular Error Probability	65

Figure 3.18. Average Expected Direct Hits by Direct Attack Small Unmanned Aerial Systems at 3-Meter Circular Error Probability	65
Figure 3.19. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels	68
Figure 3.20. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels, Ignoring Post-Incident Administrative Requirements.....	69
Figure 3.21. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels, Doubling Incident Frequencies from Recent Base Police Blotter Data	70
Figure 3.22. Effect of Delegated Response, Percentage of Incidents with Dispatch, plus Travel Time Greater Than Ten Minutes	72
Figure 3.23. Percentage of Incidents with Total of Dispatch and Travel Time Greater Than Ten Minutes	73
Figure 3.24. Percentage of Incidents Without the Desired Number of Response Units Available at Dispatch.....	73
Figure 4.1. Different U.S. Air Force Assets Protected by the Security Forces, Risk Scenarios Included in the Prototyping Effort, and Types of Security Forces Risk Reduction Explored.....	77
Figure A.1. POLICE Law Enforcement Incident Response Risk Model Architecture.....	82
Figure B.1. DEFENSE Base Incursion Response Risk Model Architecture	95
Figure B.2. Example Red Incursion into Notional Air Force Base via Roadways from ECHO Gate to the Aircraft Parking Apron	98
Figure B.3. Map of Notional Air Force Base	105
Figure C.1. SILO PROTECT Missile Field Security Forces Risk Model Architecture.....	109
Figure C.2. Variation in Expected Travel Time from Dispatch to Incident Locations Within Two Adjacent Ten Silo and One Missile Alert Facility Portion of the Missile Field by Road Surface...	112

Tables

Table 2.1. Simulated Missile Field Security Forces Activities.....	27
Table 2.2. Simulated Routine Law Enforcement Incident Types, Priorities, and Required Manpower.....	35
Table A.1. Summary of Received Incident Data.....	85
Table A.2. Summary of Interarrival Distributions	87
Table A.3. Summary of Travel Time Distributions.....	89
Table A.4. Summary of Time on the Scene Distributions	91
Table C.1. SILO PROTECT Key Model Inputs	111

Introduction

As observed by Giulio Douhet at the beginning of the airpower era, it is a fundamental reality that it is much easier to neutralize aircraft and their combat capability when the aircraft are on the ground rather than in the air.¹ Experience from numerous past conflicts has demonstrated that significant damage can be inflicted on airpower assets through penetrating attacks on air bases, striking parked aircraft, and through standoff attacks from the area around bases to disrupt operations, damage aircraft, and injure or kill personnel on the base.² Changes in technology and the development of increasingly advanced weapons have extended the range from which such attacks can be staged. Air Force Global Strike Command (AFGSC) has the added requirement to protect critical locations and nuclear weapons systems, as well as the aircraft and personnel involved in the management and deployment of those assets. In addition to direct threats to bases, assets, and personnel, maintaining base security also involves a variety of peacetime tasks that are captured by routine law enforcement or policing operations, which can vary from responding to criminal behavior on base, to maintaining order, and to enforcing law and safety requirements for residents and visitors to base locations.

U.S. Air Force (USAF) Security Forces (SF) are the central component of the USAF's effort to perform these roles. The SF mission is to "protect, defend, and fight to enable Air Force (AF), Joint and Coalition missions."³ Both military and civilian defenders perform a set of mission-essential tasks within and in the immediate vicinity of Air Force bases (AFBs) and operating locations that both fulfill day-to-day needs and hedge against rare, but potentially catastrophic, events.⁴ According to Air Force Instruction 36-2646, "Security Forces Training and Standardization Evaluation Programs," these tasks encompass the following:

- **Provide Installation and Asset Protection:** Defenders plan for and employ the capabilities of Integrated Defense to mitigate potential risks and defeat adversarial threats to the AF [Air Force] operations within the Base Boundary and Base Security Zone [the area around a base where standoff attacks can be staged against targets inside the base perimeter]. . . .
- **Provide Security and Protection for Nuclear Assets:** Defenders provide the highest degree of security possible for nuclear munitions in all circumstances (e.g., weapons storage areas, nuclear convoys, and uploaded

¹ Giulio Douhet, *The Command of the Air*, Office of Air Force History, 1983 (originally published in 1921), as cited in Alan J. Vick, *Snakes in the Eagle's Nest: A History of Ground Attacks on Air Bases*, RAND Corporation, MR-553-AF, 1995.

² Vick, 1995.

³ Air Force Instruction 36-2646, "Security Forces Training and Standardization Evaluation Programs," U.S. Air Force, December 4, 2020, p. 6.

⁴ Air Force Instruction 31-118 specifies the complementary roles of civilian and SF defenders in base law enforcement and security (Air Force Instruction 31-118, "Security Forces Standards and Procedures," U.S. Air Force, August 18, 2020).

aircraft) in accordance with . . . applicable Department of Defense, AF, and National Security Presidential Directives. . . .

- **Conduct Law and Order Operations:** Defenders directly contribute to installation's integrated defense via law and order operations, which encompasses crime prevention, criminal investigations, traffic enforcement, and corrections. In planning, the specific authorities for law and order operations may depend upon jurisdictional status of the installation. . . .
- **Provide Training and Maintenance of Small Arms and Light Weapons:** Defenders provide weapons qualifications training, forecast for sufficient ammunition in support of training, inspect, and service small arms and light weapons for AF personnel. . . .
- **Provide Military Working Dog Support:** Defenders equip, train, and manage military working dog teams to integrate into defense operations supporting Department of Defense military working dog taskings and the integrated defense plan.⁵

Beyond allocating staff and other resources to performing mission essential tasks, the SF must maintain sufficient personnel to support, train, and manage the functions required to preserve and advance their capabilities. These functions include ongoing training, training for new tasks or operational concepts, and supporting personnel medical and health needs during home station duty and deployed operations.

In recognition of the importance of airpower in modern conflict and the vulnerability of that power at base locations, the United States and other countries have explored a variety of doctrinal and organizational approaches for base protection, particularly at expeditionary bases during conflict situations. These approaches have included the viewpoint that *all* personnel present on a base (beyond designated SF) as responsible for contributing to base protection; varied perspectives on which service or forces—or, for expeditionary bases abroad, which *country's forces*—are primarily responsible for ensuring that standoff or indirect fire (IDF) attacks are not staged from the area around the base; and the incorporation of different roles for technology or static protections to reduce risk to aircraft and other key assets.

In U.S. doctrine, responsibility for base protection was once formalized with complementary roles for the USAF and U.S. Army forces, which had the Army carrying out security-related operations outside expeditionary installation boundaries. That formal arrangement ended in 1995 and was broadly replaced with joint doctrine for base defense that “made the joint forces commander responsible for installation defense regardless of service affiliation.”⁶ However, that doctrine did not prevent the return of ambiguity around roles and responsibility for the protection of air bases during conflict, particularly from threats originating outside the base boundary, and concerns about the

⁵ Air Force Instruction 36-2646, 2020, p. 6.

⁶ Milner ascribed this development to realizations made during Operation Desert Storm: “Operation Desert Storm quickly showed this agreement was not workable for the same reasons it did not work in Vietnam, namely that an offensive Army does not have the resources nor the inclination to dedicate resources to a defensive fight in the rear area” (Joseph A. Milner, *Integrated Defense: Lessons Learned from Joint Base Balad*, Research Paper 2012-3, Air Force Research Institute, 2013, p. 8). Milner, 2013, p. 8. See also the discussion in Robert D. Sagraves, “Air Support for Base Defense: Lessons for the Noncontiguous Battlefield,” in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014, pp. 169–170; Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.

USAF relying on others risked that there would not be sufficient forces allocated to perform those missions.⁷

During Operation Iraqi Freedom, the implementation of the new integrated defense (ID) concept for base security meant that there were situations during which the USAF SF managed both protective efforts inside base boundaries and in the surrounding base security zone (BSZ). During that conflict, BSZ security was carried out in cooperation with host country forces, with a focus on apprehension and justice system response to threats rather than kinetic action.⁸ At bases within the United States, the authorities of military forces generally end at the base boundary, which means that cooperation between base SF and local law enforcement can be critical for managing threats that originate from the BSZ. Depending on the characteristics of the land and the communities that surround the base, SF and local law enforcement cooperation can also be needed when crime or other events occurring in the surrounding area cross the base perimeter and create concerns for the base SF.

The combination of wartime and peacetime demands in expeditionary and home locations means that the SF must be prepared to address a wide variety of incidents and threats. The SF manage these risks through personnel, training, and technology. Appropriately preparing the SF to protect key assets at home and abroad, in circumstances in which the ability of technology or infrastructure to manage risk will vary and the associated personnel demands can fluctuate from modest to extensive, requires approaches for assessing how SF risk reduction can be assured in a variety of threat and operating environments.

Threat Landscape for U.S. Air Force Base Security

At one end of the threat spectrum, the USAF SF must address everyday threats that are very similar to those a civilian law enforcement organization responds to on a day-to-day basis. At the other end of the spectrum are threats that are distinctly military in nature, focused on protecting the base and the assets and personnel on that base from potentially highly capable attackers. During wartime, such threats can be relatively frequent and even routine occurrences; during peacetime, those threats could rarely, if ever, occur, but the SF must be vigilant in case they do. For analysis, we have broken the types of threats and incidents the SF must manage into different classes that are related to both the nature of SF activities to respond to them and the types of risk to personnel and military capability that those threats entail.

Incursion Threats Targeting Key Assets or Personnel

The SF mission set includes the day-to-day requirement to protect extremely high value assets that are critical to U.S. combat capability and to be prepared to repulse opposing forces (OPFORs) seeking to damage or seize those assets. The seriousness of attacks on air bases and the challenge of

⁷ See discussion in Milner, 2013.

⁸ Shannon W. Caudill, Anthony M. Packard, and Raymund M. Tembreull, "Defending the Joint Force: Lessons Learned from Joint Base Balad," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014, p. 204; Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.

halting them is driven by the size and capabilities of the OPFOR carrying out those attacks. In doctrine, threats have been categorized into three levels:

- Level 1: “Agents, saboteurs, sympathizers, terrorists, civil disturbances”
- Level 2: “Small tactical units, unconventional warfare forces, guerrillas, may include significant stand-off weapons threats”
- Level 3: “Large tactical force operations, including airborne, heliborne, amphibious, infiltration, and major air and space operations”⁹

Base security and defense forces are expected to protect against level 1 and 2 threats, while level 3 threats are “beyond the capability of base and base cluster security forces, and can only be effectively countered by a tactical combat force or other significant force.”¹⁰

Different attackers could seek to achieve a variety of goals by striking or staging an incursion onto an air base, shaping the variety of threats the SF must be prepared to address.¹¹ Attackers might be seeking to kill or injure personnel, destroy aircraft, or damage infrastructure, including buildings, communications, fueling, and other facilities that base operations require. Entry onto bases could be a prelude to a subsequent standoff or a conventional military attack that has attackers causing damage to protective measures, such as aircraft shelters and other defensive and security measures, or collecting intelligence to find and fix targets on base. The theft of resources or key assets can be another instrumental aim of penetrating attacks, with the nuclear assets protected by the AFGSC SF among the highest value assets that could be targeted by such operations. But the goals for attacking a base could be more modest through the use of penetrating or standoff operations to produce the transient disruption of base operations (e.g., halting the launch of combat sorties from a base temporarily), to undermine morale, or to gain a symbolic victory independent of the attack’s instrumental effects.

Because such actions could occur during periods before active conflict begins, even home station bases within U.S. territory must be prepared to detect and defeat attempted incursions onto bases. In some contexts, such as the Korean peninsula where—if conflict occurs—bases will be in close proximity to conflict zones, such actions could be the prelude to or a part of major military action.¹² During past conflicts, air bases faced significant numbers of penetrative threats—small teams seeking to enter bases to attack personnel and aircraft—some of which were successful in causing damage to aircraft and on-base personnel casualties. A notable example was the 2012 attack on Camp Bastion in Afghanistan, where a successful penetration attack resulted in significant damage to aircraft and personnel casualties. In a previous RAND report, Alan Vick analyzed past ground attacks against air bases and observed that approximately one-quarter of attacks sought to penetrate the base perimeter.¹³

⁹ Air Force Doctrine Publication 3-10, “Force Protection,” U.S. Air Force, November 19, 2019, p. 20.

¹⁰ Air Force Doctrine Publication 3-10, 2019, p. 22.

¹¹ A *base incursion* or *penetration* is defined as attackers crossing the perimeter of a base (either clandestinely or overtly via force) to act against sites or personnel within the perimeter.

¹² Joseph T. Buontempo and Joseph E. Ringer, “Airbase Defense Falls Between the Cracks,” *Joint Force Quarterly*, No. 97, 2nd Quarter, 2020.

¹³ Vick, 1995. See also David A. Shlapak and Alan J. Vick, “*Check Six Begins on the Ground*”: *Responding to the Evolving Ground Threat to U.S. Air Force Bases*, RAND Corporation, MR-606-AF, 1995.

In peacetime, these types of incidents are rare, although even then, there is the risk of larger-scale terrorist incidents with high human and military consequences originating from foreign and domestic threat actors.

- For many years, the U.S. Department of Defense (DoD) and the military services have focused on the prevention of terrorist attacks at base facilities in light of major incidents, such as the attack on the Khobar Towers housing facility in Saudi Arabia. In 2006, a group planning an attack on Fort Dix reportedly considered AFBs and other service air stations among their potential targets.¹⁴
- Particular concern has been raised about the involvement of active-duty, reserve, and veteran service members in domestic terrorist groups, particularly since the January 6, 2021, attack on the U.S. Capitol.¹⁵ In recent years, several active-duty members of the USAF have been identified as having ties to domestic and other extremist organizations or as posing mass shooting threats.¹⁶ Given the antigovernment ideology of some of these groups, military bases, personnel, and assets could be considered legitimate targets. Other domestic groups across the ideological spectrum have also targeted military and government facilities.¹⁷

Broader insider threats are also types of incursion attacks, by which legitimate access is used to facilitate infiltration and avoidance of security measures. There have also been extremely costly cases of sabotage perpetrated domestically that had no ideological component at all: A fire set during work on the USS *Miami* by a disgruntled civilian (reportedly in an effort to get off work early) caused sufficient damage to result in the scrapping of the vessel when repair estimates escalated to hundreds of millions of dollars.¹⁸

For our analysis on SF issues for the AFGSC, we looked at three categories of incursion threats on the basis of location and, therefore, the different types of security measures that potentially contribute to responding to the threat and managing the posed risk.¹⁹

¹⁴ Shannon W. Caudill, Andrea D. Tullos, and Erik K. Rundquist, *Defending Air Bases in an Age of Insurgency: Integrated Base Defense Principles for Commanders*, Vol. III, Air University Press, 2021, p. 5.

¹⁵ Caudill, Tullos, and Rundquist, 2021, p. 37; Shannon W. Caudill and Bryan A. Keeling, "Law Enforcement and Base Defense: Improving Interoperability to Benefit the War Fighter," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014, p. 309; Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.

¹⁶ These cases were identified during the course of a broader RAND research effort to collect data on publicly known cases of radicalized individuals across the U.S. military services carried out by Heather Williams, Caitlin McCulloch, Erik E. Mueller, and Shannon Prier (Heather Williams, personal communication with the authors, August 2022). See, for example, Stephen Losey, "Demoted Master Sergeant with White Nationalist Ties Is Now Out of the Air Force," *Air Force Times*, September 9, 2020b; Stephen Losey, "Security Forces Airman Accused of Trying to Burn Cop Car During Protests While Wearing Air Force Gas Mask," *Air Force Times*, August 20, 2020a; and Diana Stancy Correll, "Kirtland Airman Charged for Unlawful Importation of Firearm After Investigators Find Trove of Weapons at His On-Base Residence," *Air Force Times*, March 20, 2020.

¹⁷ See discussion in Seth G. Jones, "Violent Domestic Extremist Groups and the Recruitment of Veterans," statement before the House Committee on Veterans' Affairs, October 13, 2021.

¹⁸ David Sharp, "Damaged Sub Miami Leaves Shipyard, to Be Scrapped," *Navy Times*, June 13, 2015.

¹⁹ Because it is qualitatively different from threats posed to static sites, this effort did not look at modeling attacks on nuclear assets when they were being transported from one secure location to another (i.e., convoy operations).

- **Incursion targeting nuclear assets, critical locations, or airframes within the base perimeter:** Study interviewees emphasized that the most important mission of the AFGSC SF is the protection of the country's nuclear assets. A base incursion targeting those assets would have catastrophic consequences if it successfully seized or damaged those assets. Base incursions could also target non-nuclear assets, such as aircraft in a base parking area that might or might not have dedicated protective forces.²⁰ Although the greatest risk would be associated with incursions that are intended to cause damage (e.g., damage parked aircraft), there have been examples of actions at base perimeters or incursions into defense facilities for demonstrative purposes (e.g., protesting U.S. military policies or actions).²¹ Because these incidents involve attackers crossing the base perimeter, multiple layers of defense come into play and the SF posture of the entire base become involved in reducing risk (which is of particular relevance for nuclear assets, which have their own dedicated SF protection).
- **Missile field security threats:** A requirement for the AFGSC SF that is specific to their nuclear protection mission is the management of threats to missile silos: sites that are protected by hardened infrastructure and detection capability, but where the SF response occurs from longer distances. In addition to responding to potential threat detections (e.g., an intrusion alarm at a missile site), the SF perform protective operations during periods when silos must be opened for maintenance operations. These operations involve the deployment of personnel to a site that would normally not have SF present. Because the exposure of systems during maintenance obviates the risk reduction of the physical protection provided by hardened silos, the requirements for such deployed security operations can be a significant draw on SF personnel resources.
- **Active shooter threats that target personnel:** Active shooter and other types of mass casualty attacks are a distinctly different type of incursion threat. For these incidents, a significant concern is the potential for insiders—members of the military services, civilian employees, or others with a legitimate reason to be present on a base—to perpetrate attacks. Active shooting attacks have been staged at military facilities by insiders, including the 2009 Fort Hood and 2013 Washington, D.C., Navy Yard attacks.²² In both the military and civilian context, active shooter incidents have resulted in outcomes ranging from tens to even hundreds of individuals injured or killed, emphasizing the potentially large risks posed by such attacks to air bases. Unlike most incursion attacks, active shootings often have no specific geographic location that is targeted on the base; instead, the focus is on attacking personnel in exposed locations, with the shooter moving around the base in search of targets.

²⁰ Although threats to nuclear assets are a unique facet of AFGSC SF missions, this type of threat—posed to critical locations or sophisticated airframes—is relevant across the USAF SF.

²¹ For a historical discussion of these types of actions, see Elizabeth Heneghan Ondaatje, *Trends in Anti-Nuclear Protests in the United States, 1984–1987*, RAND Corporation, R-3696, 1989. More recent examples are described in “Clowns Break into a Nuclear Missile Launch Facility,” KLTV 7, July 1, 2006; Nancy McLaughlin, “Nuns Proudly Recount Protest Against Missiles,” *Greensboro News & Record*, November 11, 2008; and John LaForge, “First US Citizen Convicted for Protests at Nuclear Weapons Base in Germany,” *Counterpunch*, May 26, 2020.

²² DoD, *Protecting the Force: Lessons from Fort Hood*, January 2010; Under Secretary of Defense for Intelligence, *Internal Review of the Washington Navy Yard Shooting: A Report to the Secretary of Defense*, Department of Defense, November 20, 2013.

Outside-the-Perimeter Indirect Fire or Small Unmanned Aerial System Threats

In past conflicts, standoff IDF attacks using mortars, rockets, and other weapons have accounted for the vast majority of attacks on U.S. air bases. In his analysis of attacks during and before the Vietnam War, Vick observed that approximately three-quarters of attacks were staged with standoff weapons, such as mortars or rockets.²³ In recent conflicts, IDF attacks, including the use of man-portable antiaircraft systems (MANPADS), dominated even more significantly.²⁴ Concerns during conflict about effective defense outside the wire of air bases have been raised repeatedly in the literature as creating a seam from which standoff attacks pose a significant threat to bases, aircraft, personnel, and, as a result, military capability.²⁵ During the period that this study was conducted, the use of unmanned aerial systems (UAS) for surveillance and attack by Russia and Ukraine has emphasized the threat that these systems pose, including attacks on targets outside the area of immediate conflict.

Although the ability of IDF strikes to disrupt operations can be near certain (since halting flight operations and forcing personnel to take cover are primary responses to incoming fires), the risk that IDF strikes pose with respect to damage and casualties is more uncertain. The infrastructure at a base can be a strong determinant: If most personnel and systems are well sheltered in hardened structures, the probability of damage or casualties even from a strike that hits one of those structures might be quite low. The skill level of the OPFOR matters, as do such variables as the number of rounds they are able to launch in their firing window. The technologies involved are also important, with likely hit probabilities and damage from a successful strike varying considerably between simple mortars, advanced guided mortar systems, other IDFs, and small unmanned aerial systems (sUAS).

Routine Law Enforcement Incidents

During peacetime and at home stations, base law enforcement operations are a constant and many tasks closely resemble those of civilian police departments. On the basis of data on the annual activities of the SF at two AFGSC bases, incident types include engaging in traffic enforcement, responding to vehicle collisions, serving arrest warrants, and responding to criminal and violent incidents (e.g., theft, domestic violence, other crime). Other duties are very different from the civilian law enforcement context. For example, the SF perform customs duties for aircraft that land at bases from international points of origin. They also respond to events that might be initial indicators of larger-scale incidents, including resolving security alarms in facilities and investigating abandoned vehicles and breaches of the base perimeter. Although a perimeter breach could be the start of a base incursion, the breaches often are not intentionally threatening (e.g., vehicles proceeding through a checkpoint before being cleared or individuals under the influence of alcohol or other drugs climbing perimeter fencing). Similar to civilian policing, a subset of routine law enforcement incidents can include significant

²³ Vick, 1995.

²⁴ Caudill, Packard, and Tembreull, 2014, p. 204.

²⁵ Vick, 1995; Caudill, 2014; Buontempo and Ringer, 2020.

individual safety risk, both to members of the base community and to the responding SF officers. However, such situations are less frequent than less risky, everyday incidents.

Anticipating Future Threats

Although historical experience provides a window into the variety of threats and risks the USAF SF must manage, threats are not static and will shift in the future. Small-scale forces are gaining tools and capabilities that could increase their potency relative to defensive forces. As USAF platforms and systems evolve, new vulnerabilities could be created to particular types of attacks.

More-recent analyses have explored a broader variety of threats to air bases (including such threats as cruise missiles, which are well beyond what base SF are tasked with managing).²⁶ That work highlighted both sUAS and rockets, as well as mortars and non-line-of-sight missile systems. Other work has highlighted the potential use of more readily available weapons (e.g., large-caliber sniper rifles that are legally obtainable within the United States) as a threat to high-cost fighter or bomber platforms in exposed locations.²⁷ As a result, standoff threats now include much more than mortar teams staging attacks from the base security zone and expand the area outside base perimeters, where threats to USAF systems and personnel could originate.²⁸

Other non-traditional threats to air bases and their operations could become more prominent, further challenging SF capabilities. Lasing incidents have occurred in the civilian and military context, in which the illumination of planes with commercial lasers has created serious safety concerns. Other types of disruption, including communications jamming, disruption of positioning systems, or interfering with other electromagnetic signals, is also becoming of increasing importance, potentially expanding the ranges from which disruptive threats can be staged and changing the capabilities SF might need to rapidly respond and protect bases.

Effects-Based Security Planning for Security Forces Missions

The types of threats that the SF must be prepared to manage create a wide variety of risks that their operations seek to deter, defend against, and respond to. Some types of incidents occur frequently (e.g., routine law enforcement incidents) but the consequences associated with any individual incident are comparatively low. Threats to missile sites, critical air platforms, or large-scale attacks on personnel are a much lower probability, but the consequences of such attacks are exponentially higher. This reality requires focus at both ends of this spectrum because the total risk created by frequent but less consequential events can still add up to have considerable effects on the SF over time. The risk challenges faced by the SF also shift considerably between peacetime and wartime conditions and in expeditionary versus home locations. During wartime, past experience has shown

²⁶ Alan J. Vick, Sean M. Zeigler, Julia Brackup, and John Speed Meyers, *Air Base Defense: Rethinking Army and Air Force Roles and Functions*, RAND Corporation, RR-4368-AF, 2020.

²⁷ Buontempo and Ringer, 2020.

²⁸ See the discussion in Caudill, Tullos, and Rundquist, 2021, Chapter 10.

major increases in risk of direct attacks on air bases and—as dramatically shown in the historical data discussed previously—IDF attacks aimed at high-value targets and personnel.

Recognizing the varied risks that must be managed to protect personnel and assets on air bases, the USAF has facilitated a considerable evolution in security operations doctrine, including a focus on defensive operations in deployed or contested environments. As summarized by Briar in his examination of air base defense doctrine, the move toward a new framing of ID “fundamentally shifted security operations from a Cold War compliance-based model to a capabilities-based construct” and refocused on what the SF *did* rather than on compliance with standards for the numbers of SF and where they were supposed to be posted.²⁹ When implemented, the ID construct focused inside and outside the wire of expeditionary bases and its first major application was demonstrated at Joint Base Balad in Iraq in 2008.³⁰

Because the SF are tasked with protecting base personnel and assets from damaging incidents or attacks, security planning is ultimately a risk management task.³¹ As a result, SF *capabilities*, which are at the heart of capabilities-based planning, are intended to produce SF-specific *effects* or *outcomes*: the potential to deter actors from staging potentially damaging incidents and, if those incidents do occur, to respond effectively to halt the incident or contain the resulting damage. The ability of the SF to do this counters enemies’ ability to target Blue critical systems and personnel; greater SF capability will produce greater effects or outcomes, neutralizing higher levels of adversary (and other) risk to the USAF.

As a result, a critical transition occurred with the implementation of ID to use *risk assessment* as the basis for SF planning.³² Tools have been developed, including the Enterprise Protection Risk Management program, to support the systematic inventorying and risk evaluation of assets and the development of risk reduction approaches.³³ By reflecting best practice approaches for risk analyses, these tools prompt users to perform a bottom-up assessment of risk by looking at each asset and making estimates of the threats that asset faces, vulnerabilities to those threats given existing security plans, and the consequences (to military capability, time needed to recover from an attack, and so on)

²⁹ David P. Briar, “Sharpening the Eagle’s Talons: Assessing Advances in Air Base Defense Doctrine,” in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014, p. 265; Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.

³⁰ Caudill, Packard, and Tembreull, 2014.

³¹ This transition to a risk assessment approach is also consistent with other changes in planning. As summarized by Camm and colleagues (2009), “In many ways, structured risk assessment simply provides a set of accounts for keeping track of the application of common sense” but it also disciplines decisionmakers to be cognizant of uncertainties and explicit about the fact that choices involve risk. Explicit approaches to considering and talking about security risks allow command decisionmakers to “express their intuitive beliefs [which will always be an element of military judgment] more precisely.” And that more precise language could also help military leaders better articulate the level of risk involved with particular courses of action or resourcing decisions to stakeholders outside the Air Force (Frank Camm, Lauren Caston, Alexander C. Hou, Forrest E. Morgan, and Alan J. Vick, *Managing Risk in USAF Force Planning*, RAND Corporation, MG-827-AF, 2009, pp. xxiii, 110).

³² See a description in Kenneth Knox, “Automated Risk Analysis Tool to Support Transformation of US Air Force Security Forces (Briefing Slides),” June 2009.

³³ See a discussion of the historical ForcePro tool in Knox (2009) and descriptions of the current Enterprise Protection Risk Management tool, which supports the USAF ID Risk Management Process.

of a successful attack on the asset.³⁴ Such tools aggregate these risk assessments for individual assets, locations, or potential targets across a facility as a whole to provide a composite picture of facility risk.

Risk management processes—particularly when done under resource constraints—often require leaders to accept some risks to some assets in order to better reduce risk to more critical assets. The USAF ID risk management planning process is no exception, requiring commanders to make decisions about levels of tolerable risk when resources or personnel are limited. Planning and making these risk tolerance decisions is complicated by the need to consider and balance the full scope of risks discussed in previously discussed, from day-to-day incidents, to attacks of potentially high military significance.

The sensitivity of the assets and the platforms on which the AFGSC mission capability depends makes this task all the more difficult, since a successful attack on any U.S. nuclear asset could result in catastrophic consequences. As a result, requirements for the protection of nuclear assets are defined under separate security standards and, unlike planning for other types of risks, there is no flexibility in the process to meet the personnel and capability requirements involved.³⁵ Although assuring compliance with high levels of protection for such assets is understandable, the additional constraint that it creates within AFGSC SF planning can create hard trade-offs in non-nuclear SF missions.

Current Challenges for Air Force Global Strike Command Security Risk Management

Current and future threats to AFBs and assets pose a variety of challenges for security force planning, including the potential requirement for significant manpower to execute ID effectively at expeditionary bases.³⁶ The evolution of future basing constructs, including the Agile Combat Employment concept, could further stress the SF because of the training, personnel, and equipment challenges that are associated with protecting assets on bases that vary widely in number and location.

³⁴ The planning assumption that there can be significant differences in risk across locations is also reflected in DoD-level security regulations. For example, the Under Secretary of Defense for Intelligence and Security's *Physical Security Program*, updated October 19, 2020, includes that security of weapons systems can differ from location to location because of differences in mission; location and vulnerability; operational readiness; and value, classification and replacement costs (Under Secretary of Defense for Intelligence and Security, *Physical Security Program*, DoD 5200.08-R, U.S. Department of Defense, April 9, 2007, incorporating change 2, October 9, 2020, p. 19).

³⁵ See, for example, DoD Manual S-5210.41, "Nuclear Weapon Security Manual," U.S. Department of Defense, May 4, 2022, and associated supplements.

³⁶ For example, see discussion in Robert H. Holmes, Bradley D. Spacy, John M. Busch, and Gregory J. Reese, "The Air Force's New Ground War: Ensuring Projection of Air and Space Power Through Expeditionary Security Operations," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014, pp. 256–257; and Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014. Additionally,

Air Force manpower policies have provisions for deployment credit, but they were not implemented because of the resources required to fund the additional manpower and because they provided only static resources to offset dynamically shifting requirements. As a result, decades of continuous deployment demand have resulted in varying levels of workforce stress, degraded garrison service levels, or some combination of the two (Albert A. Robbert, Lisa M. Harrington, Louis T. Mariano, Susan A. Resetar, David Schulker, John S. Crown, Paul Emslie, Sean Mann, and Gary Massey, *Air Force Manpower Determinants: Options for More-Responsive Processes*, RAND Corporation, RR-4420-AF, 2020, p. xiii).

For SF, "degraded garrison service levels" potentially translates into acceptance of greater security risk.

Similarly, emerging threats, such as sUAS, are a challenge across USAF elements, and all base locations are facing constraints in their ability to prevent and respond to shifting IDF and sUAS surveillance threats. Across the USAF SF enterprise, increasing demands will require the force to flex in response, but doing so could require the acceptance of some risks, such as greater risk at domestic bases when forces are deployed abroad or increased risk to lower criticality assets to cover higher priority assets.

Along with these future demands that could increase SF manpower requirements to protect against a wider variety of threats, at least some USAF SF *currently* face manpower challenges. In this report, which focuses on the AFGSC, bases that contributed to our research reported issues filling key positions, although the nature of these manning shortfalls differed across locations.³⁷ Additionally, there were concerns expressed by SMEs at one location that the manpower studies underlying the unit manning document are outdated and, therefore, not reflective of current missions, activities, and SF requirements at the base. During our discussions, base staff also noted that some routine activities that the SF support, such as the transport of nuclear weapons, require significant manpower surges that are not fully accounted for in the manning requirements calculations. This can disrupt training activities and pull manpower away from other SF activities, and the demands of meeting mission requirements can cause significant stress on the force (e.g., requiring the use of long shift structures) that can degrade personnel proficiency and skills that are critical to responding when incidents occur.

The AFGSC's specific requirements for safeguarding nuclear assets, while also managing other risks to base assets and personnel, create a unique set of demands that are making it more difficult for the SF to accomplish the full scope of its missions. Because such requirements have no flexibility, the AFGSC's options for risk management are inherently constrained, complicating options for flexing to manage risk within this constrained manpower.

Study Approach

Although approaching security planning from an effects-based perspective is critical, the nature of risk and risk analysis means that this approach also takes on challenges that are common to nearly all efforts that are designed to address uncertain future events. Given the presence of resource constraints, there will always be the need to accept some risk in some areas. Subsequent events (i.e., whether the particular risk that was accepted manifested or not) could make a reasonable decision look misguided in hindsight, and the conclusion that is drawn from those events could dramatically shape choices going forward. For example, if a particular type of incident of concern or attack does *not* occur, it can never be definitively known whether the actual risk was lower than assumed or whether efforts to prepare successfully deterred adversaries from making the attempt. When resources are constrained, the former might be particularly attractive.

Given the immediacy of the personnel constraints that the AFGSC SF faces, approaches that enable the consideration of alternative ways to reduce different types of risk are required. Even if additional personnel would be the preferred option to respond to a particular risk, a fully effects-based

³⁷ Retention issues at one location we visited have led to undermanning in key mid- and senior-level positions that forced responsibilities down to lower-ranked and less experienced SF personnel. At another location we visited, the SF faced challenges retaining and attracting personnel in mid-grade positions (e.g., E-4 to E-6).

planning approach should allow consideration of procedural, technological, or infrastructural changes that could reduce risk if manpower is not available.³⁸ Analytic and planning approaches that capture the potential contributions of risk management tools could allow for a more explicit assessment of how much manpower, technology, and infrastructure portfolios could reduce risks to facilities, personnel, and key assets and enable a more nuanced consideration of what adequate manpower for base security means in practice.³⁹

Although existing tools support the significant use of risk assessment in security planning, the fact that concerns about SF staffing and the capability to meet all security requirements persist suggests the need for additional approaches to support risk-based security decisionmaking by planners and installation commanders. Installation commanders are also tasked with making implicit risk acceptance or tolerance decisions in the course of such assessments.⁴⁰ This is difficult, given the significant uncertainties in the nature of threats, as is anticipating how the consequences of different types of incidents could affect base function and military capability.⁴¹

As a result, complementary tools that can support commanders in their assessments of the risk associated with threat or incident types and security postures could contribute to planning and support trade-offs and inevitable risk acceptance. Tools that can provide a more broadly overarching or top-down view of risk—and of how much of that risk existing security measures might *not* be able to address—can complement the bottom-up analysis provided by asset-based risk analyses. This complementary view can help advance the doctrinal goal of integrated base defense across different types of security threats and inform the choices that base commanders are required to make within existing security and base defense planning processes.

To address this need, our work explored the utility of using a different lens for risk-based security planning. Rather than start at the individual asset level, we used simple and transparent models representing SF personnel and their capabilities to perform their roles. The models we developed and used are called the POLICE Law Enforcement Incident Response Risk Model, the DEFENSE Base

³⁸ For example, Briar argues that the SF might shift some routine law enforcement missions to civilian security staff or law enforcement or might have to “accept more risk in functions such as law enforcement, resource protection, crime prevention, administration and entry control” (Briar, 2014, p. 275).

³⁹ These types of approaches echo the needs that were raised regarding manpower planning in other studies: “Our experience with Air Force manpower processes suggest that there is a lack of objective, systematic feedback loops for determining whether adequate manpower requirements are being established” (Robbert et al., 2020, p. xiv).

⁴⁰ However, although accepting risk might be required, Air Force Doctrine Publication 3-10 advocates against doing so (because it can limit options under resource constraints):

Upon completion of the criticality, threat, and vulnerability assessments, commanders should have the information they need to make decisions regarding what level of risk they are willing to accept. However, risks to the most critical Air Force assets should be mitigated or eliminated whenever possible. If risks cannot be eliminated, commanders should implement measures to mitigate them to the greatest extent possible (Air Force Doctrine Publication 3-10, 2019, p. 31).

⁴¹ For example, other analysts of AFB defense approaches have highlighted concerns about the ability to perform locally focused threat assessments:

Under the ID concept, installation commanders would bear the brunt of accepting the increased risk based on a realistic assessment of the local threat. A key critique of ID is that threat is still largely a postulation of what we think we might face versus a transparent review of the actual local threat (Briar, 2014, p. 275).

Similarly, interdependencies among the different elements of a facility and its key systems is a challenge in risk assessment and analysis.

Incursion Response Risk Model, the SILO PROTECT Missile Field Security Forces Risk Model, and the CATAPULTA Indirect Fire Model. Those models are then challenged with various security demands or incidents with the goal of exploring how the risk of negative outcomes—including routine law enforcement incidents, to which the SF cannot respond effectively, successful base penetrations, or other attacks that threaten key assets—changes with different levels of SF manning, training, and equipping or with different assumptions about base infrastructure or technical protections. The models were informed by existing analysis of USAF SF concerns, our interviews with SF subject-matter experts (SMEs) at multiple AFGSC bases, and our analysis of quantitative and other data provided to us. We interviewed SF SMEs in person and over video call between March 2022 and June 2022. The interviews are non-attributional and no names are provided. The prototype models we developed demonstrate how such tools could be used to inform SF planning and to highlight the specific challenges faced by AFGSC planners, given the nature of their security missions.

Organization of This Report

In the rest of this report, we describe the development and illustrative application of a set of top-down risk analysis tools that explore different risks posed to air bases and key assets and the effectiveness of the SF in managing those risks. We describe in Chapter 2 the prototype models and how they represent security risk for different types of threats. In Chapter 3, we examine a specific set of scenarios and example analyses, informed by our interviews with AFGSC SF SMEs, to demonstrate the application of the tool in a way that is relevant to AFGSC concerns. Chapter 4 presents our conclusions. The appendixes describe each of the models in greater detail.

Prototyping Tools to Better Support Effects-Based Security Planning

To effectively complement existing risk assessment approaches, models for assessing SF operations and the effectiveness of those activities must focus on *outcomes* (i.e., the negative consequences that could occur if SF manpower, training, technology, or supporting infrastructure are insufficient to fully address the risk from a particular type of incident). Models that allow for the exploration of how potential consequences change with increases or decreases in personnel, different equipment, or other changes allow for a clear comparison of alternatives. In addition, because such estimates provide transparency into the residual risk to assets or personnel that arises from resource or capability constraints in the SF, the models can also clarify—at the installation level and for higher echelons—the potential consequences of risk-tolerance decisions and, therefore, allow for a more systematic consideration of whether the level of residual risk is acceptable.

Across the different types of incidents that the SF are charged with preventing or responding to, as defined in Chapter 1, the risk management effects the SF are seeking to achieve are distinct and quite different. For the AFGSC, we discuss those risks in five groups for modeling purposes:

1. **Incursion threats that target key assets or locations within the base perimeter:** For these incidents, the goal of the SF is to halt or attrit an attacking team to prevent them from success if they reach their target. Increased risk equates to a higher probability of the attackers reaching their target or reaching their target with more combat capabilities intact. For the AFGSC, the central concern in this category is risk to nuclear assets.
2. **Missile field security threats:** Risk managed by the SF in the missile field combines time-critical response to potential incidents (when alerted by security alarms) with the ability to provide security coverage when maintenance or other operations must be performed at silo sites. Risk in the former case is the ability to respond within response time standards to incidents and the speed with which the SF can support initial units on the scene to repel an attack. In the latter case, risk is related to having sufficient manpower to cover needed maintenance; shortfalls that result in the deferral of required upkeep undermine missile force readiness.⁴²
3. **Active shooter threats that target personnel:** The goal of the SF in an active shooter incident is to neutralize an attacker as rapidly as possible because incident duration is related to the number and seriousness of the casualties that an attack can cause. Increased risk is equated to

⁴² Among these five risk groups, only missile field security threats are specific to AFGSC SF requirements. Other USAF forces face and are tasked with managing the other four types of risk, although the types of key assets that are protected differ.

increasing the average duration of incidents and, therefore, potentially higher levels of injury or death.

4. **Outside-the-perimeter IDF or sUAS threats:**⁴³ During an IDF attack or kinetic sUAS incident, increased risk is related to a higher probability of weapons reaching their intended targets and causing damage. Capability limitations that delay the ability of the SF to find the originating location of such an incident to intercept the incoming weapons (e.g., by deploying counter unmanned aerial systems [cUAS]) or that delay the ability of the SF to respond against the launch point could allow attackers to stage larger and more accurate attacks (e.g., more shells or sUAS). More directly than some of the other threat types, the nature of any infrastructure protecting the intended targets (e.g., hangars sheltering aircraft) will have a significant effect on the risk of damage to assets in such attacks. Jurisdictional restrictions, concerns about collateral damage, and the consequences of using lethal or other force during a response complicate SF options to manage and respond directly to threats that originate outside base perimeters in either domestic or expeditionary locations.⁴⁴
5. **Routine law enforcement incidents:** In a policing context, incident outcomes are often related to the speed with which officers can arrive on the scene and the response options the officers have when they are there. According to the SF SMEs we interviewed, a rapid response to alarms protecting sensitive facilities (e.g., facilities containing classified information) is required to meet defined response time standards. Although broader standards for other incidents, such as the standards used in civilian policing, do not apply to the SF, a fast arrival time is desirable to limit the potential for poor incident outcomes (e.g., domestic violence or other life safety incidents). Risk for this class of incidents, therefore, relates to whether personnel or other constraints increase average response times, increasing the probability for both time to *initial* response (first officer on the scene) or time to *sufficient* response (having enough officers arrive to allow for a safe response) that exceeds acceptable thresholds.⁴⁵


These different incident groups represent the risk posed to different USAF assets: personnel, aircraft and other systems or sites, and nuclear assets. Figure 2.1 illustrates those three types of assets at risk and how our five incident groups explore that risk. We developed analytic approaches using different types of models to demonstrate how basic simulations can explore the trade-offs and potential consequences of security planning choices across risk to these different assets.

⁴³ Note that outside-the-perimeter threats could also include directed energy and electronic warfare or jamming threats as discussed previously, but we did not model those types of threats. sUAS could also be used as surveillance to support incursion incidents, which could increase the probability of success for attackers in such scenarios. In our analysis of UAS, we do not include that use case explicitly, although it could be included as a factor that affect the relative capability of Blue (defender) and Red (attacker) forces in incursion simulations.

⁴⁴ Non-lethal options could expand the ability of the SF to respond in these contexts, but actions that take place outside base perimeters will likely always be complex for a variety of reasons.

⁴⁵ Particularly for incidents with the potential for violence (e.g., domestic disputes), having enough officers on the scene to manage the people involved in the conflict can reduce the potential for injury to either the officers or the participants involved.

Figure 2.1. Different U.S. Air Force Assets Protected by Security Forces and the Risk Scenarios Included in the Model Prototyping Effort



SOURCES: The icons (clockwise from top left) are sourced from Microsoft, Lance Weisser at The Noun Project, and Luisa Iborra at The Noun Project.


NOTE: The different asset classes that are protected by the SF are represented by the blue circles, and the names of the scenarios that were simulated in the prototyping effort are shown in red.

Developing Risk Analysis Tools Using a Hypothetical Air Force Global Strike Command Base

To minimize the need to use significant amounts of base-specific sensitive information during this model prototyping effort, we elected to use a hypothetical AFGSC base, named Notional AFB, as the basis for our work. The design of Notional AFB (Figure 2.2) was informed by the general geographic layout and size of other AFGSC bases and AFBs and includes a variety of base entry points, conditions around the base perimeter, and a set of key assets, including munitions (e.g., for an AFGSC, a nuclear weapons storage area or a weapons storage area [WSA]) and flight line locations. Groups of silos and missile alert facilities (MAFs) in the area surrounding the base reflect the infrastructure at bases with missile field operations.

The size of Notional AFB within the main base perimeter is between 10.5 square miles (core built area) and 13 square miles (including an unbuilt soft-protection area, an area that is outside the base fence line but where the SF maintains situational awareness), which makes this an intermediate-sized base. Different subsets of the main base area include residential (single family-style housing units) and housing areas (barracks and group housing adjacent to the base operations area), remote training and support areas, and central operations areas. The commissary and other personnel-focused facilities are located in the interface between the housing and operations areas. Operations occupies the core military area of the base, adjacent to the runway and the ramp area.

Figure 2.2. Map of Notional Air Force Base



As is the case for actual bases, the capability to respond to different security risks at Notional AFB is treated as a function of the combined contributions of manpower (and how SF personnel are trained and equipped), base infrastructure, and technology to complement, augment, or even substitute for SF personnel to reduce risk.

Security Forces Personnel

Manpower is the central component of most physical security strategies, and the protection of air bases and their assets and personnel is no exception. Constraints in manpower availability and concern about the personnel demands for future security requirements are the primary driver of this analysis. The level of personnel capability is shaped by their training (greater training, in principle, would correspond with greater policing effectiveness or greater readiness to respond to higher-intensity

intermediate and military threats). The potency of SF manpower in reducing risk to bases is driven by the total capability the SF can bring to bear on incidents, driven by numbers of personnel and the level of capability brought by each SF member.

In our modeling, we directly examined manpower numbers. The personnel numbers that we used in our analyses were informed by the post priority charts of multiple AFGSC bases and then averaged and normalized for the intermediate size of Notional AFB. Because we focused on modeling responses to different incidents, our main concern was SF personnel who were involved in that task; individuals who were involved in supporting roles (e.g., logistics, controllers managing dispatch) were not directly modeled. In our analysis, personnel could be mobile (e.g., SF personnel engaged in routine law enforcement activities who also would be called to respond to other types of incidents) or fixed (e.g., SF personnel who manned checkpoints along the base perimeter or protected specific facilities). The latter category included substantial guard forces tasked with protecting individual high-value locations or assets, including WSAs, armed platforms, and other key facilities.

Each application of manpower to security roles contributes to risk management in distinct ways across different types of threats. This allocation of manpower supports the ability of the SF to concentrate capability on responding to individual incidents and on the protection of particularly important assets for base functioning or USAF combat capability. How manpower is allocated also shapes how SF personnel can reinforce or backfill for other categories of tasks and their effectiveness when they do (e.g., SF personnel who are in fixed positions or who patrol distant areas of the base perimeter might have a longer response time to an incident in the base core than personnel who are dedicated to routine law enforcement.)⁴⁶

Although we do not model it, the allocation of the SF to these different defensive tasks versus supporting activities (e.g., badging, visitor management) has the potential to affect base operations in ways that are distinct from risk to assets and personnel. If manpower constraints mean that the SF personnel who normally would be delivering supporting security service activities must be devoted to risk reduction tasks, wait times for those services could increase. Similarly, staffing constraints at gates or checkpoints could affect throughput rates at those locations, making it more difficult to maintain smooth base operations. In our discussions with SF SMEs at some bases, reductions in these functions have already been made in response to current manpower levels. Although our focus is solely on the risk reduction activities of the SF, these other effects from how forces are allocated should also be a consideration.

Infrastructure and Technology to Augment, Complement, or Substitute for Personnel

Security personnel play their defensive roles within a structure defined by base infrastructure and the security technology that surrounds them. Physical security measures that constrain adversary approach paths, provide strong points where defensive forces can concentrate for greater effect, and

⁴⁶ Note that depending on manpower levels, the need to allocate staffing across different categories and provide enough personnel to fulfill all risk management roles can result in limits in how the force can sustain capability and personnel health readiness over time. Without sufficient slack personnel to backfill for staff who are engaged in training or in non-duty status (to meet physical or mental health treatment needs), readiness can erode over time.

require OPFOR expenditure of time or resources to penetrate provide advantages to defenders. Improved personal gear or weapons can increase individual SF effective proficiency against opponents who threaten the base (relative to the attackers' skills, equipment, and armament). Infrastructure can shield assets from attack, including hangars and shelters that protect aircraft from IDF threats and existing hardened silos that protect missile systems. But technologies intended to augment security can also *increase* demands on SF personnel. Although perimeter or building intrusion alarm systems provide a layer of detection capability and increase the probability of detecting an attack in progress, they also can result in false alarms that must be investigated and resolved.⁴⁷ Additionally, technological systems have their own staffing and maintenance requirements, and when those systems fail, they can create substantial backfill requirements on personnel.

Types of technological and infrastructure improvements can include the following:

- **Improvements to technologies that currently create SF response requirements**, such as the deployment of more-sophisticated intrusion or other alarm systems that reduce false alarms, would reduce such incident response requirements.
- **Infrastructure improvements that concentrate SF capability or reduce the consequences of SF incidents**, such as improvements to physical infrastructure that can augment or complement the efforts of the SF, can help balance manpower limits. Improvements can vary from options that limit the effectiveness of challenging IDF threats, to those that support smaller numbers of personnel in response.
- **Technologies that can perform tasks currently done by SF personnel could allow those personnel to be retasked.** The automation of some functions (e.g., automated gates, speed or other traffic enforcement cameras, the electronic submission of information for some security or policing incidents) could allow SF personnel to be reallocated to other tasks. Some such substitutions could involve the acceptance of some security risk (e.g., automated access control).

Modeling Security Forces Risk Management Across Different Threats to the Notional Air Force Base

To support command decisions that involve security risk in ID planning, the analysis needs to be able to represent the risk reduction contributions made by personnel in combination with technology and infrastructure. In the following sections, we examine each category of threat, define operational measures for the types of residual risk that we considered for each threat, and show how we modeled those risks using the prototype toolkit.

⁴⁷ In SF incident data from multiple bases, the resolution of such alarms represented an appreciable percentage of incidents that required the dispatch of personnel.

Incursion Threats Targeting Key Assets or Personnel

Base incursion attacks (whether covert or overt) could be opportunistic (e.g., attempting to cause as much damage as possible to any asset the attack team encounters) or could be purposeful and targeted (e.g., seeking access to specific areas of the installation to disrupt operations, strike at personnel or assets, or achieve other goals). All such incidents can be highly consequential (e.g., the attack at Camp Bastion, which damaged multiple aircraft and caused significant numbers of casualties).⁴⁸ However, assaults targeting highly sensitive or critical assets or locations could be exceptionally high consequence (i.e., causing loss or damage to assets and airframes that are critical to combat capability or causing breakdowns of nuclear surety), driving specific security requirements that must be maintained at all times. Because of this focus on AFGSC SF challenges, our efforts concentrated on targeted incursion threats in our prototyping efforts.

Rapid and decisive SF action could halt an incursion before it reaches its intended objective, which would essentially eliminate any risk of attacker success. However, even if it is not defeated en route to its goal location, the Red OPFOR must reach its objective with sufficient capability intact to complete the planned action at the target. As a result, SF action during the course of the incursion—even if it does not stop the attackers before they reach their goal—could cause mission failure if the SF action attrits or ties down enough of the attacking force. For a nuclear asset that has a dedicated SF presence tasked *only* with defending that asset, the remaining capability an attacker would have to possess to achieve success would be much higher than if the attacker's objective was to damage, for example, aircraft parked in the open that did not have a significant dedicated protective force. Risk arises from the potential for the OPFOR will reach its target with sufficient capability to achieve its goals.

As discussed in Chapter 1, there are other, potentially lower risk, incidents that could fall into this category (e.g., attempted base incursions by individuals or groups for demonstrative rather than destructive purposes). In such cases, just reaching its goal could constitute success from the perspective of the incursion force, although what that force is able to do when it gets there (e.g., damage property or platforms) is of greater concern from a base security perspective.

DEFENSE Base Incursion Response Risk Model

Incidents such as base assaults are difficult to assess using statistical approaches to risk analysis; their rarity means that simulation cannot be informed by observed data in the same way as more frequently occurring security risks.⁴⁹ As a result, rather than attempting to use datasets to define distributions of incidents and outcomes, other modeling techniques have to be applied. To do this, we built a prototype model to analyze a base incursion from a starting location (usually the base perimeter) to a target location within Notional AFB.

The level of risk from an attempted incursion is related to the size (i.e., capability) of the attacking force, and the proficiency of each attacker is related to their training, armament, and other variables.⁵⁰

⁴⁸ These types of opportunistically targeted threats are similar to our modeling of active shooter incidents.

⁴⁹ National Research Council, *Review of the Department of Homeland Security's Approach to Risk Analysis*, National Academies Press, 2010, Figure 5-1, p. 94.

⁵⁰ In security planning, one or more example OPFOR numbers and capabilities must be defined to support assessment. The terms used for such a defined OPFOR include *design basis threat* or *postulated threat*.

The level of threat that is posed by an attacking team will be the product of the number of attackers and the proficiency of each individual attacker. The OPFOR will move onto the base toward its goal at a speed defined by how the force is traveling (e.g., using vehicles would be faster than traveling on foot but might limit the attackers to the existing roads on the base).

The ability of the SF to respond is similarly related to its numbers and proficiency, but also its ability to detect an incoming OPFOR and to coordinate movement to close with and engage the threat. As defenders, SF personnel occupy fixed positions (e.g., at checkpoints or gates to enter the base, at locations where dedicated SF teams protect key assets) and are mobile (e.g., responses to incursions are made by SF personnel otherwise engaged in law enforcement operations or by other SF personnel recalled to duty as an incursion is unfolding).

Existing planning constructs define the numbers and capabilities of the SF in some locations. For example, drawing on Robbert and colleagues (2020), “the fixed number of security posts at a given location plus the training currencies required are the primary drivers of the Security Forces manpower needed to man those posts.”⁵¹ But ensuring that there are enough personnel to staff every checkpoint says nothing about whether those numbers are too few, too many, or the right amount for there to be a high likelihood of being able to protect the base against any specific level of threat. This emphasizes the need to consider how force size and capability would be sized *on the basis of desired risk management effectiveness* versus on the basis of standard approaches to job analysis and man-hour task estimation.⁵²

For this threat, we built a highly simplified model that allowed for the exploration of different balances between Red OPFORs and Blue SF defensive capabilities, which showed how the results of such incursions could be presented in terms of risk measures. Given the responsibilities of the AFGSC SF, the modeling and measures were framed to be able to make comparisons of incidents at key assets (that have their own dedicated protective forces) and at other base assets or locations. The intent in building the model was not to reproduce other and more-detailed small-unit engagement models;⁵³ such models could be used in more-detailed explorations of base defense scenarios. However, the scope of this prototyping effort necessitated a more-simplified approach that could be used rapidly to explore different risk concerns and distinct ways of mitigating those risks.

⁵¹ Robbert et al., 2020, p. 34.

⁵² This framing (and our modeling effort) considers only action by *base* SF responding to an incursion, which runs counter to the line of argument in base defense literature that every airman could and should contribute to defending a base:

Finally, if the ID concept is truly designed to integrate and harness the full capacity of the USAF, it must follow through with its stated goal that base defense is a ‘fundamental battle competency for all Airmen, whether garrison or deployed’ and that ‘every Airman is a sensor.’ Without fundamental change to USAF culture and mandated roles in base defense, these slogans are hollow (Briar, 2014, p. 277).

That approach, along with the multifaceted risks and benefits of adopting it, could be modeled. Even in recent expeditionary operations, defensive roles were reportedly largely limited to the SF:


No Airmen outside of security forces at [Joint Base Balad] had any defense responsibilities, even if it was their own compound or sector, and there was resistance to taking on security responsibilities within all other organizations on the base. If ID is to be truly transformational, all Airmen must play a role, even if only a secondary support role (Milner, 2013, p. 40).

⁵³ Read more about the One Semi-Automated Force (OneSAF) and Joint Conflict and Tactical Simulations (JCATS) at United States Military Academy West Point, “Combat Simulation Lab,” webpage, undated.

Figure 2.3 (depicted at the top of the figure) shows an example scenario on Notional AFB to illustrate the modeling approach. The attacking Red OPFOR enters the base at ECHO gate, the access gate that leads to the residential portion of the base. The OPFOR follows the base road network in vehicles, crossing the operational area and eventually reaching the parking area to attack aircraft that are parked there. For the purposes of our simplified modeling, we abstract the attackers' path through the base into a straight line that they travel along from the base perimeter to their intended target (depicted at the bottom of Figure 2.3). Fixed SF units that are present at locations along that line are encountered by the attackers as they move through the base. Mobile SF move from their locations to the line of travel to engage Red when they get close enough.

The strength of the attacking force is determined by a combination of the force's numbers (e.g., five attackers) and the proficiency of each attacker (e.g., level of skill, training, equipment, armament). The overall strength of the OPFOR can, therefore, be expressed as a product of those two values. For example, an OPFOR with a capability of 20 could be made up of a team of ten individuals, each with a net proficiency of 2; 20 individuals, each with net a proficiency 1; or any other similar combination. This is a *highly* simplified representation of attacker and defender relative capability that neglects how various factors could advantage attackers and defenders in different ways over the course of an engagement (e.g., the element of surprise could give attackers an initial early advantage, and the defenders' tactical plans, which take into account geography and infrastructure, could give them an advantage once they have time to begin to execute those plans). This representation also neglects ways that larger numbers of lesser-trained personnel might have advantages over single, more-highly trained individuals through disaggregation or coordinated fires or ways that particular types of technology, surveillance, or other capabilities could provide asymmetric advantages.

Figure 2.3. Example Red Inursion into Notional Air Force Base via Roadways from ECHO Gate to the Aircraft Parking Apron



NOTE: The top portion of the figure depicts an incursion scenario on Notional AFB. The bottom portion of the figure depicts the attackers' path through the base as a straight line from the base perimeter to the intended target.

For our modeling, we only use this overall capability value to represent the attacking force and vary it over a range (evaluating outcomes versus a set number of SF personnel who can respond individually with a proficiency of 1).⁵⁴ The model, therefore, explores the consequences of relative balance between attacker and SF capability, with the added effect of the time required for the SF to travel from their duty station to engage the attacking force.

Because the goal of this effort was to not duplicate things that small-unit engagement models are already very good at, engagements were treated very simplistically. When the Red force encountered

⁵⁴ For example, in his analysis of historical base incursions, Vick found that team sizes generally varied from two to nine individuals but with outliers on the high end for very large operations (Vick, 1995, p. 58). In an example from previous conflicts, Buontempo and Ringer described that

In the 2012 ground attack on Camp Bastion in Afghanistan, a team of 15 heavily armed and well-trained—but not to the level of special forces—Taliban insurgents successfully infiltrated the base boundary and perimeter defenses to destroy six Marine Harrier aircraft with antipersonnel grenades. They also damaged ten other aircraft along with support facilities and assorted equipment. Furthermore, 2 friendly forces were killed and 17 individuals wounded (Buontempo and Ringer, 2020, p. 115).

the Blue SF (e.g., the personnel at ECHO gate), the model treated the engagement as neutralizing the same amount of capability from the Red force as represented by the Blue SF. This was intended to represent the generally larger Red force having to peel off one or more units to tie down the Blue SF, while the remainder of the Red force continued toward its tactical goal. As a result, if the two SF at ECHO gate each represented 1 unit of proficiency, the attack team of 20 would be reduced to 18 that would continue toward the parking apron (with some delay, reflecting the need to divert from the engagement).⁵⁵

An attacking force could seek to enter the base clandestinely in an effort to avoid having to engage the SF until it is closer to the tactical goal. For that scenario, the model includes detection probabilities that can be set at different levels for different segments of the Red force's line of travel. For example, detection probability might be high going through a checkpoint but much lower for a simple movement through base residential areas. If Red engages Blue, detection is assumed and all Blue units become aware of the attack.

Because of the way the model was constructed, single parameters can be used to represent a wide variety of potential SF options. Increases in SF manpower would correspond to different numbers of units at different locations, either fixed on Red's path of attack or able to respond and engage. Training or personal equipment that makes an individual SF member more effective is reflected in proficiency. Improved mobility options would correspond with a change in speed (or, for platforms such as Blue UAS that would not be restricted to the road network, shorter distances to close in on Red and, therefore, quicker engagement times.) Technologies that are intended to detect threats (e.g., improved sensing platforms in outer base areas) would change the probabilities of detection along some or all of the path of attack.⁵⁶

Measures of Risk Explored

Given the objectives of an incursion attack, the measures of residual risk, on the basis of a set of SF capabilities and a defined OPFOR, that we explored were: (1) the probability that the attacking group was neutralized before reaching its target, since sufficient SF capability to do so would fully address the associated risk, and (2) the residual capability of the Red force when it reached its target, if the group was not neutralized en route.

For an otherwise undefended target, the Red OPFOR reaching its goal with any capability remaining could equate to it successfully achieving its tactical goal (e.g., damaging parked aircraft). For a defended asset, the remaining capability of Red would engage the asset's defense force, and success would depend on whether Red's capability was sufficient to overcome the defense.⁵⁷

⁵⁵ Given the wide variety of factors that can affect engagement outcomes in small-unit encounters, this simplification is clearly significant. Developing better ways of representing unit engagement outcomes (ideally, while preserving the benefits of simplified modeling) should be a priority in future work of this type.

⁵⁶ Although the detection of an attempted base penetration could occur before the attackers reach the perimeter, for our analysis, we simplified the consideration to detection at the perimeter and the base boundary as the earliest positions where the Red force could be engaged.

⁵⁷ As we will discuss in more detail in Chapter 3, the difference between the attacker's initial capability and the residual capability with which it engages the asset's defense force also includes the contribution that all base SF made to the defense of that asset, separate from the efforts of the dedicated defense.

Caveats and Potential Extensions of the Prototype

The intent of examining incursion risk using a highly simplified model was to make it possible to readily explore SF actions and to do so in a very transparent manner. The consequence of that choice, however, is the level of abstraction that is required to do so. For example, unlike more complex small-unit engagement models, in which fields of view, obscured terrain, and specific parameters of skill or armament shape engagement results, the very simplified engagement in our model meant that another key measure of risk—Blue casualties incurred while responding to an incursion—could not be meaningfully modeled. Elaborating the engagement logic in the model (or supplementing the analysis with more detailed models) could allow casualty tracking as a complementary measure of risk to attack outcomes.

Missile Field Security Threats

The primary risks of concern in the missile field are essentially incursion threats: individuals breaching the (very near) perimeter around a silo and damaging or interfering with the silo's operation in some way. Although technically classified as an incursion threat, the response to such incidents has more in common with day-to-day law enforcement operations because the SF personnel who are responsible are based at MAFs that are near, but not at, the silos under their protection. As a result, the SF units that respond to a potential silo incursion incident (detected by onsite security systems or by other means) are dispatched from a nearby SF facility. The number of response teams available is determined by the staffing levels deployed to the missile field and by any additional reinforcement forces that are available. Like responses to routine law enforcement incidents, time is of the essence. The rapid arrival of response forces can prevent damage or other negative outcomes, reducing the overall risk from this threat. The ability to respond rapidly is driven by how quickly an incident is detected, the rapid dispatch of responding units, and other factors, such as the mode of transport, road surfaces, and weather conditions. In the case of a true silo attack (i.e., not a false alarm), which would be more analogous to an incursion threat, there is also the question of whether the Blue SF response engages the threat successfully. That success is driven by all the same variables of capability, situational awareness, and tactical mobility that are relevant in other contexts.


Unlike response activities in routine law enforcement, missile field SF have another set of manpower-intensive activities that address a different type of risk: the SF personnel who cover the maintenance activities that are required to keep the weapons and the systems protecting those weapons ready and operational. The most basic of these activities is the need for the SF to cover maintenance related to the technical security systems. If the systems are taken offline to be maintained or upgraded, supplementary SF are posted at the missile site as part of the maintenance operations. If security and intrusion detection systems at a site go down, demands on the SF to backstop for the technology grow even greater. When that occurs, SF personnel must be posted on site around the clock until the systems are brought back into operation. There are also SF requirements associated with the maintenance of the weapons systems. When a silo must be opened to perform a maintenance operation, a much more substantial SF presence is posted on site until the maintenance is completed.

SILO PROTECT Missile Field Security Forces Risk Model

Because SF operations in the missile field have some commonality with routine law enforcement operations (response units having to travel to respond to incidents, which could occur in various locations), we used the POLICE Law Enforcement Incident Response Risk Model as the basis for a model of missile field activities (which we describe in more detail). However, given the modifications involved, we treated the modified version of the POLICE model as its own model, renamed the SILO PROTECT Missile Field Security Forces Risk Model, for discussion purposes.

To explore the trade-off between SF steady-state requirements and irregular long-term maintenance concerns, the SILO PROTECT model includes incidents of long duration that occupy set numbers of SF personnel. We simulated an entire missile field at Notional AFB that is comparable in size to Minot AFB and contains 15 MAFs where SF were based in the field, with each MAF surrounded by ten silos (illustrated in Figure 2.4). An SF squadron is deployed to the missile field for extended protective operations, with personnel divided between the 15 MAFs to protect the surrounding silos and to reinforce activities in other areas of the field as needed.

Figure 2.4. Notional Air Force Base Intercontinental Ballistic Missile Field



NOTE: ICBM = intercontinental ballistic missile.

We created an incident calendar with tactical security incidents (e.g., responding to an alarm suggesting a possible incursion at a silo) and maintenance incidents (e.g., covering the long-term opening of a silo for weapon upkeep with large numbers of SF personnel). The frequency of incidents and the number of SF personnel required to respond to each type of incident were based on

approximate estimates provided by SF SMEs during project interviews. For the prototype model, we assumed a response would come from either the MAF in its unit of ten silos or an adjacent group of missiles (i.e., the travel time distribution for the model was created based on a distribution within two adjacent MAF/silo complexes rather than a full geographic simulation of the entire missile field).

Table 2.1 summarizes the incident types and their characteristics. All numbers are approximate and do not correspond directly to security standards, but they are informed by discussions with SF SMEs to support the relevance of relative comparisons among the different activities and the duration of those activities.

Table 2.1. Simulated Missile Field Security Forces Activities

Activity	SF Personnel Involved	Duration at Silo
Alarm response	4	$\frac{3}{4}$ – 12 hours (1 $\frac{1}{2}$ hours, most likely)
Routine launch facility checks	2	$\frac{1}{2}$ – 3 hours
On-site posting because of alarm failure (“camper”)	3	1 – 3 days
Short-term maintenance activity (e.g., fixing alarms or sensors at a silo)	3	3 $\frac{1}{2}$ – 8 hours
Long-term maintenance activity	18	2 – 7 days

NOTE: Duration at silo includes only on-scene response time. Travel time from the dispatch location is treated separately in the modeling.

The incident types in Table 2.1 are listed from highest to lowest priority from top to bottom. Responding to alarms at silos was given the highest priority in the model and those incidents are accordingly staffed first to minimize any delay in sending a team to investigate. Routine launch facility checks were ranked as the second-highest priority, followed by posting SF personnel to cover silos with alarm system problems, smaller-scale systems maintenance, and larger-scale missile maintenance activities.⁵⁸ This effectively meant that in the simulations, tactical protective missions (i.e., alarm response) were strongly prioritized over maintenance operations. Therefore, in situations with limited personnel numbers, it might not have been possible to complete all maintenance requirements. If a maintenance event arose when there were insufficient SF personnel to meet the requirements of that event, the maintenance was deferred. The number of times maintenance was deferred in a simulation run was tracked to reflect how shortfalls in SF personnel were creating readiness risk for the missile force.

⁵⁸ SF personnel who are posted to cover silos with alarm system problems were termed *campers* in our project discussions. For these tasks, SF personnel stay on site when a silo’s alarm system is not functioning properly. In our model, the result of prioritizing malfunctioning alarms below routine launch facility checks (which do not involve large numbers of personnel or extended periods to complete) could have resulted in some simulated delay before SF staff were able to be present at the site experiencing alarm anomalies.

Measures of Risk Explored

For our analysis, we used two primary measures of risk for missile field SF operations: (1) the risk of long response times because of constraints on personnel to respond and (2) the risk from instances when maintenance efforts had to be deferred because of insufficient personnel to secure the operation.

As emphasized by the SF SMEs with whom we spoke, these are critical measures of risk for the missile force for which SF manpower shortfalls create tensions that pull in opposite directions. Prioritizing a rapid reaction to events such as alarm response or the deployment of campers to cover failures in technological security systems manages direct, adversarial risk to the force. However, if doing so means that critical maintenance cannot be completed, systems can fall out of ready status, which is also a risk for the effectiveness of the nuclear forces. In essence, shortfalls in either measure are not acceptable, but—when personnel are short—one measure might have to be covered at the expense of the other.

Caveats and Potential Extensions of the Prototype

Because of how the SILO PROTECT model scheduled units to cover incidents, once a member of the SF was committed to a specific task (e.g., monitoring a silo with an improperly functioning alarm system), they would not be reallocated to another task. For missile field security operations, this is realistic in many cases (e.g., personnel protecting an open silo would not be pulled away from that post to respond to an alarm). However, in reality, SF staff *would* be redeployed from routine launch facility checks to other tasks. This presumably created a small distortion in the frequency of delayed coverage of silo sites with malfunctioning alarms and the number of deferred maintenance events. It also meant that we did not *directly* model response times to alarm incidents, because the model lacked the logic to redeploy personnel from routine launch checks to an alarm occurring elsewhere in the missile field.

In addition, the way that we modeled missile field operations in this prototype focused on SF *coverage* of the types of time-sensitive and long-term incidents and tasks and did not include incident engagement success for responses to true incursion attempts at silos. Although our incursion model did demonstrate how the results of engagement modeling could be used in a risk and effects-based approach for assessing SF requirements, those concerns are nonetheless important in missile field security operations. As a result, a more complete risk-based examination of silo security would also include tactical outcomes as part of the analysis, particularly given the challenges that distance, mobility, and travel time can pose for the ability of the SF to amass force at an incident scene in the event of an attack.⁵⁹

Our modeling also does not capture the effects that changes in complementary forces could have on SF requirements, most notably maintenance capacity. The duration of long-term demands on the SF, including personnel covering malfunctioning technology systems and system and weapons maintenance, is driven in part by the availability and capacity of maintenance personnel. USAF investments in maintenance personnel capacity—to reduce wait times for maintenance, increase the

⁵⁹ In our discussions with SF SMEs, several issues were highlighted that would have been reflected in this type of modeling, including concerns about the off-road mobility of potential attackers versus defenders, the consequences for engagement success of the Blue units arriving over time versus being able to more effectively amass for response, and the potential for equipment and armament asymmetries between attackers and defense.

speed of maintenance, or both—would be another way to allow for the reallocation of SF personnel to other risk reduction activities.

Active Shooter Threats Targeting Personnel

An active shooter threat is a threat that is not theoretical for military installations. As discussed in Chapter 1, both completed attacks and planned incidents that were disrupted by law enforcement have taken place. The involvement of insiders in such attacks can make detection a challenge and the SF potentially become aware of the threat only when the first shots are fired. The involvement of trained military members in such attacks could also mean that the proficiency of the attacker might be more substantial than shooters in a civilian context.

Addressing risk from an active shooting incident is essentially a matter of fast response: The longer a shooter is able to continue their attack, the higher the potential casualty count. Rapid and decisive action by the SF or the police to stop the attack contains the number of people who can be injured or killed. An active shooter who seeks to cause additional casualties must move in search of new targets the longer a shooting incident continues. As a result, such an incident is similar to an incursion, for which the ability to rapidly locate and closely engage the attacker will define the attack outcome.

Ending an active shooting incident also depends on the relative capability of the attacker and the defenders who engage them; a group of SF engaging simultaneously will have a better chance of ending an incident than a one-on-one engagement between a defender and an attacker.

DEFENSE Incursion Response Model Applied to Active Shooter Threats

Because of the similarities between an active shooter threat and a base incursion (i.e., an attacker moving through the base seeking to achieve a goal), we modified our DEFENSE Base Incursion Response Risk Model to explore an active shooter incident. Instead of a defined path to a specific, targeted site on base, the active shooter route moved through the base from one side to another to simulate an attacker's continuous movement in search of targets.

In a departure from our base incursion model, we needed to modify how we simulated the engagements of an active shooting incident; otherwise, the first SF responder on the scene would effectively end the incident, which is inconsistent with the observed course of historical active shooting incidents. As a result, we implemented a basic probability-based engagement resolution (described in more detail in Appendix C) in which two equally matched (i.e., equal proficiency) opponents each had a 50 percent chance of neutralizing the other when they engaged. Differences in the opponents' proficiency could then be represented as bias in the coin flip (e.g., a more proficient SF responder would have a 60 percent chance of success in an engagement rather than a 50 percent chance).

Measures of Risk Explored

The main measure of risk tracked in this modeling effort was the average duration of the shooting incident, given different SF conditions, with the assumption that longer incidents would have greater

numbers of associated casualties. As a result, the risk reduction between two SF options would be reflected in a lower average incident duration.

Secondarily, because this simulation included the basic coin flip engagement logic, a measure of risk that is related to direct injury or fatality to SF personnel is the average number of engagements that are required before the incident is ended. Each coin flip encounter the attacker wins corresponds with an SF defender who was neutralized (e.g., killed, injured, or otherwise overcome) by the attacker, so the total number of such events (minus the final engagement that ends the attack) would be related to the number of SF casualties that occur during the scenario.

Caveats and Potential Extensions of the Prototype

Although adding a basic engagement logic to our base interdiction model made it possible to explore an active shooter scenario, that logic is nonetheless very simplified. Even though basic probabilistic outcomes for engagement are used in the modeling, specific circumstances can change engagement outcomes considerably in real-world situations. As a result, although the prototype makes it possible to examine differences in SF posture or approach based on risk for these types of scenarios, more-detailed models with more parameters to capture real-world outcomes are needed for more-refined results.

Outside-the-Perimeter Indirect Fire or Small Unmanned Aerial System Threats

Considering risk with respect to IDF or kinetic sUAS attacks is qualitatively different from incidents that require physically entering the base. Originating outside the base perimeter, IDF attacks are countered by a wider range of capabilities than just SF closing in to engage. Those strategies can include preemptive efforts to deny safe launch locations, long-distance counterfires targeting those locations, and rapid response in an effort to catch and engage firing teams before they are able to disperse.⁶⁰ Putting pressure on launch locations reduces the number of weapons that could be launched in an attack and, assuming that the pressure forces the attackers to hurry, disrupts the attackers' ability to aim successfully, reducing the probability that launched weapons will accurately hit targets and the likelihood of damage.⁶¹ However, carrying out these types of operations requires action within the BSZ.⁶²


Figure 2.5 illustrates the much longer distances involved, given the range of a medium-sized mortar, a short-range sUAS, and a moderate-range sUAS.

⁶⁰ "Within the BSZ, efforts of security forces, or other base defense forces assigned area security duties, to suppress IDF threats consist of physical presence, aggressive patrolling, and limited active defensive measures designed to deny adversaries access to the standoff footprint. Intriguingly, however, in what amounts to a significant omission for joint security operations planning, the BSZ is recognized as a planning construct that is used only by the air component" (Buontempo and Ringer, 2020, p. 116).

⁶¹ An analysis of the Balad Joint Base in Iraq explored similar risk metrics (including average miss distances of attempted attacks) to assess base defense operations (Milner, 2013).

⁶² Air Force Doctrine Publication 3-10, 2019.

Figure 2.5. Base Security Zone of Notional Air Force Base Relevant to Indirect Fire and Small Unmanned Aerial System Threats



NOTE: CLOS = command line of sight. The range of a medium-sized mortar is depicted in the top green circle enclosing the base, with the aimpoint centered on the flight line. In the middle of the figure, the outer blue circle depicts the range of a short-range sUAS compared with the range of the inner green mortar circle. In the bottom right corner, the much larger blue circle depicts the range of a moderate-range sUAS compared with shorter-range weapons in the inner green circle.

For UAS threats, technological countermeasures can be a component of risk reduction, reducing the chance that an attacking device reaches its intended target. Counterfires aimed at incoming ordinance (e.g., mortars) are also available that would reduce the number of attacking devices that reach their targets, although those could be challenging to use in a domestic peacetime context, given the potential for collateral damage. The hardness of the target (e.g., the vulnerability of aircraft in shelters, hangars, or other infrastructure) could also lower the risk associated with these threats by reducing the chance of damage even from weapons that reach their intended target.

Although the SF has control over some of these for risk reduction approaches, other approaches (e.g., infrastructure choices for protecting aircraft) are outside its purview. In the domestic peacetime context, the SF is also limited in its ability to operate outside the boundaries of AFBs, meaning that the detection and response to threats that originate outside the base perimeter may rely on local law enforcement.⁶³ Adversary action outside base perimeters also will limit counterforce options, especially actions with the potential to inflict fatalities or collateral damage, not just during peacetime and in domestic locations but even during wartime, depending on the rules of engagement and the requirements of host countries.

CATAPULTA Indirect Fire Model

Building a customized model for SF ground operations in response to IDF threats was unfortunately beyond the scope of this effort. However, given its importance and prominence in both historical attacks on air bases and literature on air base security, we applied an existing Project AIR FORCE (PAF) IDF model to make it possible to illustrate how SF activities and their effect on this threat could be represented in a risk-based way.

The CATAPULTA Indirect Fire Model is a model that was developed in earlier PAF research and that has been used in a variety of studies.⁶⁴ Using the characteristics of weapon systems (e.g., accuracy, number of weapons fired, damage radii for each weapon, how the attacker is attempting to arrange the impact pattern, weapon guidance) and the characteristics of any submunitions, the model calculates the extent of damage caused to specific aircraft types arranged in an open parking area and captures factors such as aircraft size and position on the parking apron.⁶⁵ Appendix D includes additional detail about the model.

Using this model to consider SF risk management requires linking actions to changes in the variables included in the simulation, such as the following:

- **Number of attacks staged:** SF actions to detect or respond to potential attackers could act as a deterrent and reduce the number of attempted attacks. The extent to which the SF (or local law enforcement contributing to situational awareness in the BSZ around domestic bases) could deter attacks would be related to the extent of presence and the perceived risk of staging such attacks.
- **Accuracy:** SF actions that put pressure on attackers could affect accuracy of an attack and reduce the probability of fires hitting by denying the attackers time to aim or, for guided

⁶³ Other authors have also questioned whether these capabilities are sufficiently included in base protection planning, even during wartime:

However, the Air Force's base defense inventory does not include organic counter-rocket, -artillery, and -mortar capabilities or the associated threat early warning alert systems. This capability must be coordinated with the Army or host-nation forces, if available. . . . Undoubtedly the Army's capacity to support air bases with counter-IDF systems and associated threat early warning alert systems will be further stressed by the Air Force's force structure expansion plans and emerging concepts for distributed operations (Buontempo and Ringer, 2020, p. 116).

⁶⁴ The primary developer of CATAPULTA is Scott Savitz, and the tool has been used in studies of air base defense under a variety of different scenarios and conditions, predominantly documented in classified literature.

⁶⁵ The model can also be used to estimate casualties from IDF attacks on personnel tenting areas but that option was not used in our analysis.

rounds or sUAS, interrupting guidance to target if the attackers are using datalinks. The extent that SF action could reduce the accuracy of shots fired or the number of sUAS launched during an attack would be proportional to how quickly the SF could identify the launch location and put pressure on the attackers, either through the launch of counterfires or a response to the attacker's position (on the ground or in the air).

- **Number of weapons fired:** Putting pressure on a firing location could similarly reduce the number of weapons fired if the attackers sought to escape before a response force or counterbattery fire arrived. As with the SF effect on attack accuracy, the effectiveness would depend on how quickly the SF could fix and respond to the launch location.

Systems to intercept incoming weapons (e.g., cUAS or defensive systems for mortars or missiles) would also have an analogous effect by reducing total weapons effectively fired by destroying a subset of those weapons en route. The amount by which such measures would reduce overall attack risk would be related to the ability of the standard performance parameters of the technology systems to fix and engage the incoming weapons.

Other infrastructure changes could reduce the probability of damage or weapon damage radii, even if weapons strike at or near their intended target (i.e., impact a shelter roof rather than the plane inside).

Measures of Risk Explored

Analyses of historical IDF threats to bases have shown that those threats pose a variety of risks. The most easily recognized risks are the potential for casualties among targeted personnel and damage to combat critical platforms or facilities that are housed on the base. However, even an attack that is unsuccessful (e.g., does not cause damage because the attack missed its intended targets) results in substantial disruption to base operations to make the area safe and to clean up after the incident. Depending on the circumstances, the time needed to resume flight operations could involve considerable risk to military outcomes. Given our use of the CATAPULTA model to illustrate risk analysis for this threat, our measure of risk is narrowly focused on the expected damage to planes in the example attacks, in which changes in SF action can reduce the average number of planes damaged or the seriousness of the damage.

Caveats and Potential Extensions of the Demonstration Effort

Unlike the other models that were developed specifically for this project, in this case, the effect of SF actions are not being directly simulated (e.g., we are not modeling an SF team searching for and then closing in on a launch point for an IDF attack). Instead, the link between SF action or the application of technology that might reduce this risk is done parametrically by simulating residual risk across a range of possible SF effects on attacker accuracy or the size of the attack. This approach allows for a structured discussion about the extent to which different options might manage this threat (e.g., the potential scope of local law enforcement effectiveness in detecting a launch point for mortars or the addition of Blue SF UAS that could enable more-rapid detection and response), but it is not direct simulation of the response.

Routine Law Enforcement Incidents

Addressing the risk from routine law enforcement incidents is a dispatching and response problem, for which force planning is needed to ensure that units will be available when needed to respond to an uncertain stream of events (Figure 2.6). Risk comes from incident occurrence and the consequences of those incidents, and response activities—ideally rapid and effective response activities—manage that risk.

Figure 2.6. Incident Response at Notional Air Force Base



For specific base SF, the distribution of incidents will vary over time. At a small base (meaning fewer facilities and a smaller population to protect), incidents might occur one at a time and multiple simultaneous incidents would be a rare event. At a larger base, multiple simultaneous incidents might be more common. Fewer expected incidents with fewer response demands could be successfully addressed with fewer SF personnel dedicated to law enforcement operations.

Individual incidents can vary in intensity, potential consequences, and, therefore, seriousness and response requirements. A routine traffic stop or the investigation of a minor traffic accident might require only one security officer and have a very low chance of serious negative outcomes. On the other hand, a domestic violence incident, which can be extremely dangerous to the individuals involved and the officers responding, might require multiple security officers to gain control of the situation and to resolve safely. Depending on the local conditions, bases experience a mix of incident types (e.g., a base with many sensitive facilities protected by alarms will have a much greater number of responses to

alarms going off than one with fewer such facilities, and differences in base population will affect the distribution of other types of incidents).

POLICE Law Enforcement Incident Response Risk Model

To model routine law enforcement, we compiled a list of different incident types and the frequency of those types based on police blotter data provided by multiple AFGSC bases. For our simulations, we grouped more-specific types of incidents from base data into composite types for simplicity. During base visits, we elicited estimates from the SF SMEs about the number of needed personnel, the required response time to resolve each type of incident, and the extent of post-response administrative and other work associated with fully resolving an incident. For example, at one base we visited, the SMEs estimated that resolving a major vehicle accident that occurred on base (a relatively uncommon event) would likely require four members of the SF to manage the scene for approximately two hours until towing resources could remove the involved vehicles. We combined those estimates with averages from blotter data to create the hybrid schedule for possible incidents and their response requirements occurring at Notional AFB.⁶⁶

We used that data in our POLICE Law Enforcement Incident Response Risk Model (depicted in Table 2.2 and described in detail in Appendix A) and as the basis to create a calendar of incidents that occurred during the model run. The model was run through the calendar, assigning available response units to each incident as they occurred. If no response units were available (e.g., because of low manpower levels), there was a delay until resources could be dispatched. The total response time to an incident included any dispatch delay, travel time (drawn from a distribution based on distances on the base), time spent on the scene, and—for some analyses—time when the responding unit was occupied after the incident with the completion of administrative and other requirements to fully resolve the incident.⁶⁷ This basic process is shown in Figure 2.7.

Table 2.2. Simulated Routine Law Enforcement Incident Types, Priorities, and Required Manpower

Incident Type	Modeled Priority	Manpower	Incident Type	Modeled Priority	Manpower
Security alarm or alarm fail	Highest	3	Abandoned vehicle	Low	1
911 emergency services call	High	2	Assistance rendered	Low	1
DUI	High	2	Criminal	Low	2
Death	High	8	Drug	Low	2
Domestic	High	2	Eagle eye	Low	1


⁶⁶ This is essentially a statistical approach to estimating risk on the basis of the probabilities of various incident types and their requirements. Day-to-day incidents are more common (higher probability) and, therefore, provide more data to characterize risk from different levels of personnel availability and other risk reduction investments. As a result, day-to-day incidents are more amenable to statistical approaches for estimating risk than rarer incidents with more uncertain consequences (National Research Council, 2010, figure 5-1, p. 94).

⁶⁷ The analysis was performed with and without administrative work occupying responding units, which allowed us to bracket the effect of those requirements on available manpower for response.

Incident Type	Modeled Priority	Manpower	Incident Type	Modeled Priority	Manpower
Fire response	High	3	Patrol response	Low	1
Ground emergency	High	3	Property damage	Low	2
Medical emergency	High	3	Vehicle accident	Low	2
			Warrant	Low	3

NOTE: DUI = driving under the influence.

Figure 2.7. Modeling Cycle and Times Simulated



This workload-based approach to considering how numbers of SF personnel affect the ability to meet the demand of different levels of incidents is consistent with manpower analyses used in civilian policing.⁶⁸ In such analyses, the ability of different levels of manpower to meet varying call-for-service demand are examined by time of day and day of the week and often reflect seasonal variation from month by month. Matching the number of officers on duty at different times of day with hourly demand is intended to ensure sufficient staff are available for rapid police response. Response goals are set on the basis of call priority. For example, the standard for a high priority call (e.g., an active, violent confrontation or other life-safety incident) might be that an officer arrives on the scene within a few minutes, while for a lower priority call (e.g., a property crime or the investigation of a petty criminal), the standard for officer arrival might be much longer.⁶⁹

The potential role of technology and base infrastructure factors is more limited in routine law enforcement than in risk management related to other threats. There are technologies that can affect incident response operations, including communications, command and control (C2), vehicles, personal protective equipment (PPE) and SF duty weapons. Because our prototype model does not measure whether incidents are responded to successfully (e.g., analogous to an active shooter being neutralized, ending that type of incident) or directly model risk to responding officers during individual incidents, technologies such as PPE or armament are not currently reflected.⁷⁰ Communications, C2, and vehicle effects could be modeled, because each of those technologies could affect the dispatch and travel of responding officers to an incident. However, given the size of our modeled base, travel time (shown in the bottom half of Figure 2.7) represents a relatively small portion of the total incident response time, so the effects of such changes on model outputs would be modest.

Measures of Risk Explored

In our modeling, we focused on two main measures of response-related risk:

- **Response time:** Some SF incidents (notably sensitive facility alarm responses) have defined response time requirements, but other types of incidents have the potential to escalate to more-negative outcomes the longer a response takes. For a crime that is in progress, the longer it takes for officers to arrive on the scene could mean that more are people injured or a higher probability that the perpetrator escapes apprehension.⁷¹ The risk associated with longer response times goes up with personnel constraints, since multiple simultaneous incidents or units stationed far from the incident location could extend dispatch or travel times,

⁶⁸ See, for example, the review in Jeremy M. Wilson and Alexander Weiss, *A Performance-Based Approach to Police Staffing and Allocation*, Office of Community Oriented Policing Services, 2014.

⁶⁹ When faced with personnel constraints, some civilian law enforcement departments have deprioritized certain incidents to the point where officers will not respond to the scene, and victims are directed to file online reports for insurance or other purposes. Furthermore, responses to residential burglar alarms have been stopped in some areas or fines implemented if a particular location produces a high rate of false alarms that occupy police resources.

⁷⁰ Given the small numbers of such incidents happening on base in a given year, on the basis of SF blotter data, these factors would also only affect a small percentage of law enforcement response activities.

⁷¹ For example, it is much more critical for responders to arrive on the scene quickly when an incident's outcome is still evolving (e.g., injured individuals with worsening conditions, a perpetrator who is continuing to hurt people, an intruder who has triggered an alarm and is still inside the building) than one for which the outcome is static (e.g., responding to a property crime when the perpetrator is already gone or to a minor traffic incident).

respectively (as depicted in Figure 2.7). Given a specific number of officers, deployed response units, and other factors, the residual risk in achieving that mission can be specified as *the probability that the SF cannot get an officer on the scene within a specified number of minutes*.

- **Risk considered from the perspective of the officers responding to an incident:** Although many routine incidents pose only minor risk to the responding officers, incidents that involve violence or the potential for violence can pose a significant risk of injury. In such cases, the simultaneous presence of multiple officers can be critical for controlling the situation and reducing the chance of negative outcomes. If constraints in the number of officers or other facets of response capacity slow the ability to gather sufficient officers at a scene, the probability of injury to the first responding officers could increase. Given assumptions about how different response postures affect the chance of officer injury, a complementary risk measure is *the time required to reach a sufficient number of SF personnel on the scene for the incident type*, for which sufficiency is defined based on the response requirements from our SF SME interviews. The longer the time it takes before sufficient manpower on the scene is reached, the higher the risk potential to the initial personnel on scene.

These measures illustrate how risk could shift if manpower levels change, either temporarily to respond to a short-term requirement (e.g., increases in demand for personnel deployment or in staffing requirements associated with fielding new weapon systems) or long-term if new weapon generation facilities begin operations without proportional increases in SF manning.

Caveats and Potential Extensions of the Prototype

In considering routine law enforcement, there are additional potential measures of risk that could be examined beyond the simple service-delivery/workload approach that we used. Analogous analyses in civilian policing often include variation in demand and therefore variation in manpower on different shifts, therefore producing more nuanced measures of response risk for day versus night shifts. Modeling could also include measures for *successful* response. In our model, an incident is resolved simply by allocation of a responding unit spending the appropriate period of time at the incident and doing post-incident administrative work. However, adding additional nuance to what makes a response to an incident more or less successful could provide a way to reflect concerns raised in our SF SME interviews about levels of experience of the force and training effectiveness (e.g., a higher chance of a negative outcomes if a newly enlisted airman responds to a complex or demanding incident than if a more highly trained or experienced SF member responds).

Our analysis for routine law enforcement operations focused *only* on police services or direct response activities. Another facet of SF law enforcement operations is the investigation of criminal behavior after the fact, a function performed by SF members (potentially assisted by members of the USAF Office of Special Investigations). Just as the case for incident response, various characteristics of the SF (e.g., training, capability, experience) and available manpower could affect the success of investigation activities. Furthermore, the potential consequences of increases or decreases in those

factors can be expressed in terms of risk, although for this activity, risk would be framed in terms of investigative outcomes or using variables such as case clearance rates.⁷²

Chapter Summary

After we looked across the different threat types addressed by SF operations, we developed a set of prototype models (and adapted some existing models focused on base defense) to make it possible to show how different manpower levels, tactics, or technology and infrastructure changes affect risk to nuclear assets, aircraft, other systems or sites, and personnel. The four prototype models are as follows:

- The DEFENSE Base Incursion Response Risk Model explores targeted attacks inside the base perimeter and opportunistic incursion threats, such as active shooting incidents.
- The SILO PROTECT Missile Field Security Forces Risk Model (derived from the POLICE Law Enforcement Incident Response Risk Model) explores the SF response to incidents and security requirements in the missile field.
- The CATAPULTA Indirect Fore Model explores direct attack risk outside the perimeter.
- The POLICE Law Enforcement Incident Response Risk Model explores day-to-day law enforcement risk management and response.


Because management of the full variety of risks is the ultimate goal of the SF, effects-based SF and base defense planning requires the ability to perform such analyses so that the risk consequences of decisions are clear and risk tolerance judgments can be made in a transparent and systematic way. The Venn diagram in Figure 2.8 shows these different classes of USAF assets and how the different threat scenarios that were explored in this chapter overlap with those classes.⁷³

As described in Chapter 1, different risks fall across a spectrum that is defined by the likelihood of incident occurrence (which includes both the chance that those incidents will be attempted and the probability of success if they do) and the consequences associated with those incidents if they do occur. In our prototyping effort, we included our measures of risk and the ways that SF action can reduce those risks. For example, our measures related to incursion threats focus on the probability that an attempted incursion will be successful, while measures for IDF and sUAS attacks focus on the consequences that are caused by an attack. Different SF strategies (and complementary technological and infrastructure) can reduce risk by affecting one or both of those components of risk. Figure 2.8 summarizes the different risk reduction effects that are relevant for each of our modeled threats and scenarios.

⁷² There is some literature that examines civilian policing criminal investigation that could provide a foundation for this type of analysis for the USAF SF (for example, David L. Carter and Jeremy G. Carter, "Effective Police Homicide Investigations: Evidence from Seven Cities with High Clearance Rates," *Homicide Studies*, Vol. 20, No. 2, May 2016), although recent literature that is focused specifically on investigative manpower is limited.

⁷³ Note that Figure 2.8 documents how our scenarios aligned across asset types. For example, IDF attacks could be targeted at populated areas of the base with the intention of posing risk to personnel rather than aircraft, but that scenario was not included in our prototyping effort.

Figure 2.8. Different U.S. Air Force Assets Protected by the Security Forces, Risk Scenarios Included in the Prototyping Effort, and Types of Security Forces Risk Reduction Explored



SOURCES: The icons (clockwise from top left) are sourced from Microsoft, Lance Weisser at The Noun Project, and Lluisa Iborra at The Noun Project.

NOTE: The different classes of assets that are protected by the SF are represented by the blue circles. The names of the different scenarios that were simulated in the prototyping effort are shown in red within those circles.

In Chapter 3, we use the prototype models to explore several example scenarios with respect to each threat and the associated risks. The scenarios were informed by long-term planning challenges the SF is facing and key questions and concerns raised in our discussions with SF SMEs during the study.

Exploring Example Scenarios and Challenges Using the Prototype Tools

In our discussions with SF SMEs, we explored a variety of current challenges that are affecting SF effectiveness and the sustainability of the force. Current manpower constraints (including manning SF units at less than their authorized levels) coupled with rigid requirements for nuclear security are straining the SF and challenging its ability to be effective across the full scope of its missions. Basic equipment concerns (e.g., the functionality of vehicle or C2 technology) are also having an impact on effectiveness. These accumulating demands on force personnel create additional manpower stresses (e.g., stress concerns that result in personnel placed in do-not-arm status and unavailable for duty) and concerns about whether personnel performance and effectiveness are being degraded over time. The consequences of all of these different factors and constraints will play out as increased risk, which must be considered in planning. Our discussions also explored a variety of opportunities to reduce risk or to increase the effectiveness of the SF in managing risk. Changes in policies, improvements in infrastructure, and the use of new technologies could all provide paths to address security and mission risks, even if constraints on SF manpower continue to be an issue.

Those concerns and opportunities provided the roadmap for our modeling efforts. We applied the prototype models to each threat type to explore both the risk consequences of the challenges that were raised in our discussion and the potential benefits of some of the opportunities for risk reduction. In the following sections, we examine each threat type and the potential effects of excursion scenarios that explore different approaches to risk management.

Incursion Response Risk

To illustrate our risk analysis for base incursion, we examined two main scenarios: (1) an overt attack by an OPFOR targeting a key asset (a flight line with nuclear-armed planes and, therefore, a posted and dedicated defensive force) and (2) an attempted covert infiltration to reach the base WSA from a near entry point that has the SF focusing on detection and interdiction before the OPFOR reach the WSA. Each scenario illustrates different facets of SF capability and capacity for managing risk to key assets within the perimeter of the base.

In addition, both scenarios highlight—in somewhat different ways—a point made in our interviews with SF SMEs that, for nuclear or other key assets *within the perimeter of an AFB*, all base SF on duty contribute to the protection of those assets. Despite the importance of nuclear missions, the standards for the protection of weapons and platforms do not currently take into account that additional protection is provided by generally tasked SF (i.e., the SF personnel who normally conduct law enforcement operations but who would immediately pivot in the event of an incident that

threatens key assets). As a result, the standards require the same level of protection for the transport of such weapons on base as they would in an off-base situation. If tools such as the prototype modeling approaches we used can be developed such that the added protection by generally tasked SF can contribute to meeting asset protection standards, the SF could gain more flexibility in managing manpower while still providing necessary security levels across the critical assets they protect.

Direct Assault Scenario

One of the most challenging scenarios for base SF, and a major driver for the high level of protection afforded to nuclear and other key assets, is a direct assault on a base by a sizeable and well-trained OPFOR. Box 3.1 details the scenario we simulated to show how risk from this type of threat could be explored. We varied the size of the attacking team but kept the deployment of SF personnel and capabilities set. Figure 3.1 presents the scenario visually.

Box 3.1. Direct Assault to Reach a Key Asset

Scenario

A Red team of 5 to 20 individuals breaches entry control at the FOXTROT gate and attacks the flight line (a distance of 1,660 meters).

Assumptions

- **Blue posture** (each Blue SF member was assumed to be equally proficient as each Red team member):
 - Entry control: three individuals (200 meters from the Red team entry point)
 - Near Red's objective: ten individuals (two to three fire teams, 1,500 meters from the Red team entry point)
 - Operations area (near ramp): six individuals (two fire teams)
 - Housing area and operations area boundary: six individuals (two fire teams)
 - Residential area: six individuals (two fire teams)
 - Southern operations area (near the WSA): six individuals (two fire teams)
- **Detection:** Red has a high probability of being detected near the entry gate (half the time) or is detected somewhere else along the incursion route with equal probability.
- **C2:** Blue units move independently to engage Red once Red is detected.
- **Engagement:** When opposing units are within 50 meters, engagement occurs; all units within 100 meters participate. The proficiency of each Red attacker is set to equal each Blue defender. When the forces engage, the more numerous force loses the same number of people it engages (e.g., 20 Red engaging two Blue results in two Red neutralizing two Blue and 18 Red continuing).
- **Speed:** Red is 75 percent as fast as Blue units are, which simulates greater uncertainty operating in a less familiar environment.

43

Figure 3.2. Probability of at Least 1 Red Opposing Force Member Reaching Their Target by Red Force Size





Figure 3.3. Average Number of Red Opposing Force Members Reaching Their Target by Red Force Size



NOTE: The graph presents the number of remaining Red force members after neutralizing all engaged SF during the simulation, including the ten SF terminal defense at the target area. The error bars are 95 percent confidence intervals that are based on the replication mean and estimated standard errors over a number of replications ($n = 30$).

Clandestine Infiltration Scenario

Complementing the large direct assault scenario is a scenario in which the attackers attempt to enter the base stealthily to closely approach a sensitive site before acting. For this scenario, we compare two cases: one in which the attackers are only definitively detected when they reach the WSA because of heightened security in the area and a second in which perimeter detection systems provide warning rapidly when the attackers come onto the base. In the baseline case, in which the attackers are detected when they reach the WFA, the shorter amount of time that is available for Blue units to close in and engage the attackers limits their response. In the increased detection case, more Blue units have more time to move and engage the OPFOR. Box 3.2 details the scenario we simulated to show how risk from this type of threat could be explored.

Box 3.2. Clandestine Infiltration and Assault on Key Asset

Scenario

A Red fire team of four individuals surreptitiously breaches the outer fence at the BRAVO gate to attack the WSA (2,900 meters). Red is deemed successful if it reaches its objective with at least one person.

Low-Detection Variant Assumptions

- **Blue posture** (each Blue member was assumed to be equally proficient as each Red team member):
 - Operations area (near ramp): six individuals (two fire teams)
 - Housing area and operations area boundary: six individuals (two fire teams)
 - Residential area: six individuals (two fire teams)
 - Southern operations area (near the WSA): six individuals (two fire teams)
- **Detection:** Red is not detected until it approaches the WSA.
- **C2:** Blue units move independently to engage Red once Red is detected.
- **Engagement:** When opposing units are within 50 meters of each other, engagement occurs; all units within 100 meters participate. The proficiency of each Red attacker is set to be equal to each Blue defender. When the two forces engage, the more numerous force loses the same number of people that it engages (e.g., four Red engaging two Blue results in two Red neutralizing two Blue and two Red continuing).
- **Speed:** It was assumed that Red is traveling on foot and would take 54 minutes to reach its objective unimpeded. The Blue units were motorized and can travel 25 to 40 mph.

High-Detection Variant Assumptions

- Same as above, but Red can be detected at the perimeter fence using surveillance technologies with high probability.

The low-detection scenario is shown graphically in Figure 3.4, which has the small Red team entering at the southern BRAVO gate of Notional AFB and progressing the relatively short distance to the WSA along the road network. For Red to maintain a clandestine posture, it was assumed the team would be on foot. When detected, the Blue units closed in at vehicle speeds along the base road network to engage the Red attackers.

Figure 3.4. Modeled Clandestine Infiltration Scenario Schematic (Low Detection)




Figure 3.5 shows a comparison of the outcomes for the two example scenarios. In both cases, the small Red team is likely to be neutralized before reaching the WSA. Even with limited detection capability, there is still an approximately 60 percent chance that the Red team will not reach its target. Improved detection increases the amount of time that is available for response, which further reduces the risk of Red reaching its target and increases the chance of Blue defeating Red to above 80 percent. For cases in which portions of the Red team do reach the target, the increased Blue response time that is provided by the detection technology reduces risk by shrinking the expected number of Red attackers that reach the target (see Figure 3.6). Without improved detection, the average number of attackers that reach the target is, on average, one person higher than in the scenario with the technology.

Figure 3.5. Chance of Red Neutralization Before Reaching Weapons Storage Area





Figure 3.6. Average Surviving Red Personnel Reaching Weapons Storage Area



NOTE: Error bars are 95 percent confidence intervals that are based on the replication mean and estimated standard errors over a number of replications ($n = 30$).

Key Takeaways from the Example Scenarios

The demonstration simulations were focused on attacks targeting different key assets: The first was a substantial attack on armed aircraft (and, therefore, a more exposed location), and the second was on a fortified storage area. In both cases, significant levels of protection were provided for those assets by SF personnel who were not in nuclear protection roles at the time of the attack. In the second scenario, which focused on the WSA, no defensive force was included in the simulation and the results of the infiltration focused only on the number of remaining individuals who would

encounter the dedicated defensive force. In the former scenario, a terminal defensive force of ten SF personnel is included in the simulation, and the number of remaining Red forces (depicted in Figure 3.3) was above the number required to neutralize that SF defense. Illustrating a dynamic that was raised in our interviews with SF SMEs, in these simulations, the non-nuclear SF made substantial contributions to the protection of assets within the base perimeter. As a result, in a system- and effects-based approach to security planning, the effectiveness of those SF forces could be considered as satisfying some of the requirements for the protection of nuclear assets, which suggests a potential for increasing SF protective flexibility for AFGSC SF planners.

Missile Field Security Risk

When responding to routine law enforcement incidents on base, the distance that the SF had to travel to get from their starting location to an incident was often extremely short, which means that travel time had a relatively small effect on risk management effectiveness. However, missile field operations were the exact opposite: Because of the significant distances involved (sometimes with speed constraints imposed by vehicle type or road surfaces), travel time can be a central driver of effectiveness. As a result, personnel constraints that add additional dispatch time before travel begins can be a more significant limitation on risk management. Using the SILO PROTECT model, we explored the following two scenarios that looked at the different elements of risk the SF manage:

1. Scenario one looked at different levels of SF personnel in the squadron assigned to protect the missile field that created potential constraints for security alarm response time and the coverage of maintenance operations.
2. Scenario two looked at simple infrastructure improvement that showed the effect of the mix of gravel versus paved roads (each with different travel speeds) in the missile field on response time risk.

The SILO PROTECT model responds to a key issue that was raised in our discussions with SF SMEs: the challenge of managing the very distinct risks of adversary action against missile silos and the disruption of scheduled maintenance activities for those systems (because of insufficient number of SF to protect the repair activities) which risks the readiness of those weapons. Managing adversary action demands a rapid response to potential incursions (e.g., intrusion detection alarms at sites) which, because of the significant travel distances, require SF teams be placed as close as possible to the sites they are protecting. Scheduled maintenance requires committing significant SF personnel to sites that require or are undergoing maintenance and repair, making those personnel unavailable for rapid response activities, sometimes for long periods. The simultaneous simulation of personnel deployment for both activities can show where, under personnel constraints, the requirements of one type of risk management begin to hinder the ability to manage the other type effectively.

Because the base POLICE model that we adapted to model missile field SF operations did not include an option for preempting an ongoing task for a higher priority activity, alarm response times (the highest priority and most time-sensitive incidents, as presented in Table 2.1) were not modeled

directly.⁷⁴ Although many SF activities would not be preempted by an alarm response (e.g., campers, coverage of maintenance activities), since those personnel cannot disengage, other SF activities most certainly would be (e.g., routine launch facility checks). In addition, personnel who are in the missile field but on rest rather than active status could respond (although with some delay). Because we were adapting the POLICE model for missile field operations, our simulation did not have the initial capability to directly track rest status for personnel. As a result, all SF personnel were treated as available to respond, and total manning to produce the simulated response capability in the graphs would be higher by a multiplier to account for personnel in rest status.

For a security alarm response, we assumed that there will always be a unit available somewhere in the missile field to respond (i.e., there will be no delay time before an SF team can be dispatched, although the team's starting location could be further from the alarm location than the closest MAF to that silo). As a result, for our example measures of risk, we focused on the other types of modeled activities to show the effects on risk from manpower or infrastructure changes.


Incident Response Time Risk from Manpower Constraints

As manpower devoted to missile field security is reduced, there will be fewer SF personnel standing by to respond to time sensitive incidents and a smaller pool of personnel available to cover long-duration tasks, such as campers or maintenance security coverage. To demonstrate this risk dynamic, we simulated SF operations in our model missile field at four different manning levels—between 200 and 350 personnel available to respond to incidents in the 15 MAF/silo complexes.⁷⁵ Figure 3.7 shows the result of those staffing constraints, with the amount of time before enough personnel are available to staff each type of incident (i.e., variations in dispatch delay) increasing at lower levels of manpower. Although the tasks can still be staffed, the delays in being able to do so show increasing risk, particularly for events such as campers, for where staff are deployed to compensate for breakdowns in technological security systems. At high levels of staffing, the amount of time it takes to arrive on the scene becomes dominated by travel time, because the personnel constraints causing delays in dispatch are no longer a contributor.

⁷⁴ These activities could be modeled directly with some modification to the model, but time and scope constraints within the project, given the range of risks explored, meant it was not possible to do so in this prototyping effort.

⁷⁵ This range corresponds with between 13 and 24 people on duty and available to respond at each MAF.

Figure 3.7. Effect of Security Forces Squadron Size on Time Before Different Incident Types Can Be Fully Staffed (Including Dispatch and Travel Time)



NOTE: Error bars are 95 percent confidence intervals that are based on the replication mean and estimated standard errors over a number of replications ($n = 30$).


Readiness Risks Created by Maintenance Coverage Shortfalls

The most personnel-intensive incident in the simulation was the coverage of major system maintenance, which requires nearly 20 SF personnel to cover the activity for an extended period. These are the activities that contribute to the management of nuclear readiness risk because maintenance cycles must be kept up to maintain the capabilities of the systems. Because of the number of people involved, coverage of these maintenance activities is the most difficult under manpower constraints. In our interviews with SF SMEs, these were the operations that were flagged as the most “at risk” in situations when personnel numbers dropped, either temporarily (e.g., because of personnel who become unavailable from being placed on the do-not-arm list or for other reasons) or as a result of longer-term manning at below authorized force levels.

In our model, we tracked the number of times that these more-significant maintenance activities could not be staffed because of manpower shortfalls at different levels of force numbers deployed in the missile field. We looked at force levels from 100 personnel available to respond and cover incidents, to up to 350 personnel available to cover a wide variety of possible force levels. We tracked two events: (1) maintenance operations that were delayed and (2) maintenance operations that were

so delayed that they were effectively postponed indefinitely.⁷⁶ The results at different force levels are presented in Figure 3.8.

Figure 3.8. Major Maintenance Events Delayed or Postponed Because of Manpower Constraints



NOTE: Error bars are 95 percent confidence intervals that are based on the replication mean and estimated standard errors over a number of replications ($n = 30$).

As the number SF personnel devoted to missile field protection decreases, the number of scheduled missile maintenance cycles that have to be postponed, either temporarily or indefinitely, increases rapidly.⁷⁷ Given the numbers of SF personnel who are required to cover such activities, that dynamic is not surprising, since prioritizing coverage of these maintenance events would consume increasingly larger percentages of the SF as total numbers decrease.


⁷⁶ A maintenance event was counted as *delayed* if it could not be covered with available SF when it came up in the incident calendar. During many simulation runs, a delayed maintenance event could be subsequently picked up when the SF finished other high priority incidents and enough personnel became free to cover the event. The total duration of the simulations that were run for this analysis was one calendar year. If maintenance events were still backlogged at the end of the simulation run (i.e., the events had not been covered by freed up SF by then), the events were treated as *cancelled* or *postponed indefinitely*.

⁷⁷ At high numbers of available SF personnel (more than 300 in Figure 3.8), delays occur only rarely as a result of the chance occurrence that many incidents happen at the same time.

Reducing Response Time Risk via Infrastructure Improvement

Because travel time is a significant contributor to the total response time for missile field security incidents (for which the total response time includes any dispatch time that results from units being unavailable and the time for those units to reach the scene), we used a basic infrastructure improvement scenario to show potential risk reduction by physically increasing the potential speed of SF response. Because missile fields are serviced by a combination of paved and unpaved roads, we varied the extent of road paving for that scenario.⁷⁸ In the model, the speed at which the SF units can travel is related to the road surface. On the basis of the distribution of distance measurements in our missile field around Notional AFB, travel times varied by different combinations of road surfaces, as shown in Figure 3.9.

Figure 3.9. Variation in Expected Travel Time from Dispatch to Incident Locations Within Two Adjacent Ten Silo/One Missile Alert Facility Portion of the Missile Field by Road Surface




NOTE: The maximum travel distance for this simulation was 65 miles between dispatch and response location. The error bars show 1 standard deviation in the travel time.

Using these different distributions of travel times, Figure 3.10 shows the effect on how rapidly the SF can be on the scene for different types of incidents or tasks. Although these effects would be most critical for alarm responses (which were assumed to be dispatched immediately and, therefore, were constrained only by travel time), going from a fully unpaved to a fully paved road surface can produce

⁷⁸ Other scenarios could effectively increase the speed of the SF response to a site, including different vehicles or aerial options (e.g., Blue UAS) that could reach silos much more rapidly. Other conditions could also reduce response speed, including weather or road damage.

an approximately one-hour difference in time to arrive on the scene on the basis of the modeled travel speeds.

Figure 3.10. Change in Security Forces Time to Scene for Different Percentages of Road Surface in the Missile Field



NOTE: Increasing the amount of paved road surface does cause decreases in expected *travel* time for all incidents. However, the marginal returns from improved travel time in Figure 3.9 and the variability in dispatch time (as a result of backlogged events of different levels of priority) means that the differences in *total time* to the scene for *non-alarm response incidents* are not significant after the amount of paved road reaches 50 percent. Error bars are 95 percent confidence intervals that are based on the replication mean and estimated standard errors over a number of replications ($n = 30$).

Key Takeaways from the Example Scenarios

Simulations of missile field SF operations show the trade-off between two critical types of risk: (1) immediate security incidents, exemplified by alarms, or situations in which technological security measures are not functioning and (2) readiness risks from the delay or the deferral of missile maintenance activities. Although the former is potentially more *salient* when security roles and functions are being considered, it is risk that is only realized when attacks occur. The reality that maintenance is a constant requirement means that risk is more *certain* for the latter. If there are insufficient numbers of SF to cover those requirements, the missile force will accumulate readiness risk over time. Although true incursion attempts at silos are fortunately comparatively rare, maintenance requirements are constant.

In our prototype models, we simulated manpower levels as steady state force levels (i.e., what would performance look like with a set number of personnel on duty during a rotation in the missile field?). This approach does not fully reflect the dynamics of risk consequences from shorter-term SF manpower constraints (e.g., surge requirements to deploy new systems or to support generation operations, the transfer of nuclear weapons from storage onto planes) or even manpower shortfalls that result from personnel who are unable to be on duty because of illness. As such surges occur, they could cause significantly heightened risk in the short term, which, depending on the threat environment at the time, could pose consequential concerns for SF readiness or SF ability to execute nuclear missions.

Active Shooter Risk

One central driver for successfully resolving active shooting incidents rapidly and, therefore, minimizing casualties is the ability of the police or the SF to concentrate forces and effectively engage the attacker. As a result, the main scenario we explored with respect to an active shooter threat was the effect of manpower levels and the number of responders available to rapidly respond. However, in our interviews with SF SMEs, one issue that was raised was how existing demands on the force (e.g., long shifts to cover requirements, challenges fully meeting training schedules, a stressed force with a significant number of members on the do-not-arm list) would affect SF performance when an incident occurred. This was framed in terms of concerns about the level of force *proficiency* (even if units were in *compliance* with training and other requirements and, therefore, formally viewed as ready). Although this issue of proficiency could be incorporated into any of the prototype models we developed, we chose to use our adapted incursion model for the active shooter scenario to illustrate this concern by varying the probability of success when the SF engages the shooter in an encounter.⁷⁹

Box 3.3 presents the details of the scenario we used for the simulation, and Figure 3.11 shows the path of the attacker through Notional AFB. We simulated SF units in each main area of the base traveling to close with the shooter with a maximum incident timeline of one hour. If the shooter was not successfully engaged within one hour, the simulation ended. The duration of the incident was tracked with the assumption that the incident duration would be proportional to the casualty numbers that resulted from the attack. Although the initial response to a shooter would rely on the SF personnel on duty and their ability to respond rapidly, as an incident continued, there would be a surge of personnel from other roles (e.g., SF personnel responding after the entrance gates were closed, SF personnel active in other roles, and off-duty SF). As a result, our use of one hour as the maximum

⁷⁹ In the more general incursion model, the issue of proficiency or stress on the force would be integrated into the capability variable (i.e., a less ready SF member might have a capability of 0.8 rather than 1.0). As written, the routine law enforcement response model does not include a successful response as an outcome, although the effects of force proficiency could be incorporated as changes in how long it takes the SF to respond to incidents (a less proficient force would be slower to resolve a situation) or other factors. Similarly, the success of the SF in addressing IDF threats could similarly be decremented at lower levels of proficiency, although our application of an existing model for that threat means that there is not a direct anchor point in the model to do so. The shooter's probability of surviving an engagement is described by $\mathbb{P}(X > B)$, where B is the number of Blue units participating in the engagement and X is a binomial random variable that corresponds with B trials and success probability: $1-p$, where p represents proficiency (or accuracy). In other words, each trial represents whether a Blue unit misses the shooter, and the engagement is resolved in favor of the shooter if all Blue units miss (and it is assumed that the attacker neutralized the involved Blue units, so they do not get additional engagement opportunities).

incident duration was likely unrealistic, given that SF forces could be concentrated to end the incident more rapidly. We did not simulate that component of the scenario, however, because doing so would have obscured the effect of different levels of personnel proficiency that this case was intended to demonstrate.

Box 3.3. Active Shooter on Notional Air Force Base


Scenario

An armed individual or insider enters the installation from the DELTA gate and initiates the incident while moving toward the ramp near flight line (4,500 meters from the gate). Police units and patrols respond and engage Red with varying proficiency.

Low-Detection Variant Assumptions

- **Blue posture:**
 - Operations area (near ramp): two individuals (randomly located)
 - Housing area and operations area boundary: two individuals (randomly located)
 - Residential area: two individuals (randomly located)
 - Southern operations area (near the WSA): two individuals (randomly located)
- **Detection:** Not applicable—Red is deemed detected when the incident begins.
- **C2:** Blue units move independently to engage the shooter once the incident begins.
- **Engagement:** When a unit is within 20 meters, engagement occurs. Units engage with the shooter individually. Proficiency is described as the likelihood that a Blue individual can neutralize the shooter during the engagement. The duration of the engagement is assumed to be instantaneous.
- **Speed:** The shooter moves approximately 3 mph (on foot). Blue is motorized and can travel 25 to 40 mph.

Figure 3.11. Modeled Active Shooter Scenario Schematic



Security Forces Personnel Capability and Readiness Effects on Casualty Risk

Figure 3.12 shows the results of the active shooter simulations. The results are bookended by cases of zero proficiency (i.e., no SF engagement will be successful, so the incident goes to the maximum modeled duration) and 100 percent proficiency, where incident duration is determined by the modeled time for the first SF unit to close in on and engage the shooter. A combination of high proficiency and numerical overmatch results in a relatively flat central part of the curve but with increasing slope for proficiency levels below 40 percent (where 50 percent would be approximate parity in proficiency between the attacker and the defenders). Below the 40 percent level, the average incident duration begins increasing with assumed higher levels of casualties.

Figure 3.12. Effect of Modeled Security Forces Proficiency on the Average Duration of Active Shooter Scenario





Figure 3.13 shows the average numbers of engagements between the SF and the attacker in the active shooter simulations. Because a successful engagement between the SF and a single shooter will end the attack, the number of engagements greater than 1 represents a measure of Blue casualties as a result of the incident. The number of engagements trends to 1 at high SF proficiency (i.e., once SF personnel close in and engage, they end the incident) but at lower proficiency levels, engagements and, therefore, Blue SF casualties increase.

Figure 3.13. Expected Number of Engagements Between Attacker and Defenders with Increasing Modeled Security Forces Proficiency



NOTE: The error bars represent 95 percent confidence intervals based of 30 replications. Some error for proficiency values greater than 80 percent might be negligible. The estimates and errors are sensitive to assumptions regarding the incident. The maximum allowable duration of the incident is 60 minutes.

Key Takeaways from the Example Scenario

Although a narrower scenario, the active shooter simulation shows how differences in proficiency between the SF and the attacker (which in our model could be driven by training, other force readiness and health drivers, and equipment or armament) can affect risk management effectiveness. At high proficiency, outcomes were constrained by dispatch and travel time (which, similar to other scenarios, could have been reduced with higher numbers of SF personnel on duty to respond).

But more important than the cases above 50 percent proficiency (essentially parity between the attacker and the defenders) are the cases on the other side of the scale. Even in our simplified simulation, durations began increasing with small relative differences and by 35 percent (effectively a 1.4 to 1 ratio between the proficiency and effectiveness of the attacker to each defender), the durations began jumping more significantly. Such a difference is certainly within the realm of possibility, assuming a well-armed and well-prepared attacker is engaging defenders who might already be stressed from meeting routine security demands (which was the concern raised by SF SMEs in our project discussions). The effects of such demands might not be obvious until an actual incident occurs (particularly modest reductions in proficiency and effectiveness), supporting the value of using simulation to explore the potential effects of those demands on risk.

Indirect Fire and Small Unmanned Aerial System Attack Risk

Because we did not create a custom model for an SF response to off-base IDF and sUAS kinetic risks or the practical limits of SF options to respond to risks, our exploration of this threat was the most limited. To bracket how our modeling could explore how security strategies could affect these

risks, we looked at (1) a small mortar attack scenario, in which SF pressure on the attackers limited the accuracy of the attack, and (2) an sUAS kinetic attack scenario that showed how technical measures or infrastructure changes that reduce the effective numbers of sUAS used in an attack shape risk exposure. In both scenarios, however, the actions of the SF to affect attacker behavior are only notional, since those actions are not directly modeled in the existing CATAPULTA model.

Pressure on Mortar Launch Position to Reduce Indirect Fire Risk

The use of mortars in attacks on military facilities, and on airfields in particular, has been a common feature in war and in irregular conflict. As described in Chapter 1, large numbers of mortar attacks on bases have necessitated off-base responses and counterfires capability. In irregular and terrorist conflicts, police activity in areas around bases has been a core element of reducing the risk of IDF attacks. Although the prevention of attacks is the fundamental goal, such pressure can force attackers to adopt tactics that limit their accuracy, thereby reducing risk, even if some number of attacks are still attempted. A prominent example of this dynamic was the extended conflict between the Provisional Irish Republican Army (PIRA) and the British military. PIRA used a variety of mortars in its attempts to strike bases or aircraft and had notable tactical successes. However, in many cases, the group was forced by SF action to fire mortars in a non-traditional way, such as concealing the mortar in a parked vehicle and triggering it with a timer rather than using a fire team to launch it, who could observe initial shell impacts and adjust aiming to increase the chances of tactical success.⁸⁰ The resulting accuracy degradation significantly reduced the risk posed by PIRA's mortar attacks.

To demonstrate this risk reduction outcome of SF efforts to detect and respond to off-base IDF threats (and, for this threat, potentially local law enforcement or community collaboration with the SF), we used the CATAPULTA model to simulate the outcomes of a set of basic small mortar attacks. Each simulation used five unguided shells with three levels of accuracy defined by circular error probability (CEP): (1) a low CEP, high accuracy case of 108 meters (which is within the range attainable by a trained crew); (2) a high CEP, low accuracy case of 324 meters (which represented more approximate aiming by a fire team under pressure); and (3) a medium CEP, medium accuracy case of 216 meters.⁸¹ The five shells were simulated as being fired in a row along the lengths of aircraft parked in an open area.


We modeled two scenarios: (1) a high-density parking area with 40 large aircraft—based on the dimensions of a KC-135—parked in 48 available parking spaces arranged in a 6-by-8 grid and (2) a low-density parking area with only 20 parked aircraft out of 48 available spaces. Five hundred simulations were run for each combination of weapon accuracy and target set to estimate the number of aircraft that were hit directly, severely damaged (but not hit directly), and moderately damaged; lightly damaged aircraft were excluded.

⁸⁰ See discussion in Brian A. Jackson, John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, and Horacio R. Trujillo, *Aptitude for Destruction: Vol. 2, Case Studies of Organizational Learning in Five Terrorist Groups*, RAND Corporation, MG-332-NIJ, 2005.

⁸¹ The CEP is a measure that is used to describe the precision of a standoff attack mode. Given the characteristics of the system and the skills of the personnel using that system, CEP is the radius of a circle around the intended aim point within which one-half of the rounds fired would be expected to land.

Figure 3.14 shows the results of one example run for the highest and lowest accuracy cases for each target array, illustrating how the CATAPULTA model calculated the impact points for each of the mortars and then determined how close the impact was to a plane. In Figure 3.14, the blue points represent the bounding points for each aircraft, including its nose, wingtips, and tail. Mortar impact points are indicated by the concentric circles, with the center circle indicating a direct hit or severe damage, the middle circle indicates moderate damage, and the outer circle indicates the boundary of the region where light damage would occur. The horizontal red bars indicate direct hits or severe damage to the corresponding plane, the horizontal yellow bars indicate moderate damage, and the horizontal white bars indicate light damage.

Figure 3.14. Representative Model Runs of Mortar Impact Points in Simulations for High and Low Accuracy, Dense and Sparse Target Sets




NOTE: m = meters. The blue points represent the extremes of each KC-135 dimension aircraft, including the nose, wingtips, and tail. The areas of the mortar impact are indicated by the concentric circles, with the center circle indicating a direct hit or severe damage, the middle circle indicating moderate damage, and the outer circle indicating light damage. The affected planes in this run are indicated by red horizontal bars for a direct hit or severe damage, yellow horizontal bars for moderate damage, and white horizontal bars for light damage. For runs that show five impact points, the missing impacts occurred outside the plotted area.

Figure 3.15 presents the simulation results graphically, showing the reduction in risk (of damaged aircraft) if SF action reduces the accuracy of the mortar attack. For the scenario in which the aircraft are densely parked, almost every mortar shell in the high accuracy simulation at least moderately damages an aircraft; reducing attack accuracy cuts the expected damage by more than one-half. In the lower density parking scenario, the dispersal of the aircraft means that there is a greater chance of a complete miss by one or more shells, SF action that reduces accuracy brings expected damage down a comparable amount from an already lower baseline.

Figure 3.15. Average Aircraft Receiving Direct Hits and Severe and Moderate Damage by Mortar Accuracy and Target Set

Average
Aircraft
Damaged



NOTE: Lightly damaged aircraft are omitted from the graph to show the differences in numbers of more seriously damaged planes.

However, as discussed briefly in Chapter 1, SF action to *achieve* reduction in the accuracy of such attacks is very challenging, particularly in the domestic U.S. environment. The ability to do so may rely on action by local police or alternative technical or response options that are not readily available and whose implementation would be quite complex. Even in expeditionary contexts, SF action within the BSZ to achieve this goal is still challenging (and manpower intensive) and requires its own type of coordination with host country forces and domestic internal security or law enforcement organizations.⁸²

⁸² See, for example, the discussion of operations in Iraq in Milner, 2013.

Technical Countermeasures or Hardened Infrastructure Example for Reducing Risk from Small Unmanned Aerial Systems Attack

The sUAS threat has been recognized as a potentially serious threat to both air base operations and SF personnel. The nature of the threat to large aircraft is such that substantial off-base attacks could pose serious prelaunch survival challenges for any aircraft that are not capable of being removed from the base or protected by a combination of active and passive defenses. The physical area from which an sUAS threat could arise is quite large (even when the range is limited to sUAS or loitering weapons with data links) compared with other off-base attack systems, such as guided or unguided mortars that might be used to attack targets in the continental United States.⁸³ DoD's Armed Forces (the USAF in particular) and the U.S. Department of Homeland Security (DHS) writ large have been making substantial efforts to counter such systems. Many of these efforts are associated with rendering sUAS and loitering weapons ineffective by directly killing the air vehicle or interfering with its radio control and data links. The SF already has some of these capabilities, but large attacks by autonomous and semiautonomous sUAS will remain a serious challenge.

There is, however, another opportunity to further disrupt sUAS operations, at least for systems that make use of some of the shorter-range variants for which line of sight to the sUAS is required for SF operations. Direct SF action to harass, mark, and track those who are responsible for launching the base attack might be feasible and could decrease the effectiveness and magnitude of the attacks. The main operation for the SF in this context would be to launch their own sUAS when threats approach the base that would attempt to localize the control elements and personnel responsible for the control of the group. This effort would entail localization of the threat emitters and launching SF sUAS into the area to search, tag, and possibly disrupt operations.

If the sUAS control signal can be reliably detected, the localization process could use time difference of arrival systems to geolocate the sUAS controlling transmitter and generate a small search area. That narrowed search area could then be passed to the dispatched Blue sUAS system and, through proper channels, to local law enforcement.⁸⁴ Because of the relative isolation of many AFGSC bases, it might be possible to do this without overwhelming the intercept receivers and with less perceived impact on local population privacy through the use of fixed intercept receivers that limit signal collection to areas within the line of sight of the approaching sUAS.⁸⁵ For short-range systems (10 km or less), there are scenarios in which the control element could be somewhat exposed to detection. Unfortunately, within a modest range of most bases, population centers of a few tens of thousands of people can be found, which increases the number of emitters and creates many opportunities for remote antennas from control elements to make localization problematic.

Such operations would fall into a gray area between the DoD homeland defense mission, DHS, and law enforcement. The likely areas for signal collection could incorporate medium-sized cities,

⁸³ Brian A. Jackson, David R. Frelinger, Michael J. Lostumbo, and Robert W. Button, *Evaluating Novel Threats to the Homeland: Unmanned Aerial Vehicles and Cruise Missiles*, RAND Corporation, MG-626-DTRA, 2008.

⁸⁴ Kimberly N. Hale, *Expanding the Use of Time/Frequency Difference of Arrival Geolocation in the Department of Defense*, RAND Corporation, RGSD-308, 2012.

⁸⁵ A notional range of 40 to 50 km might be reasonable in terms of dealing with a transmitter on the ground that is communicating with an sUAS approaching an air base. This range would address sUAS making a dive toward a target, but it would lose contact as the sUAS dives below the radio frequency horizon in the last few seconds of flight.

where excluding the bulk of the signals would be required. Other tips, such as launch reports or unusual activities, might also spur action that interferes with operations.

The risk reduction payoff associated with these SF operations might be thought of in terms of the number of effective sUAS missions that might arrive at a target area. To highlight the effect on the number of planes that could potentially be damaged, the CATAPULTA model was used to assess the likelihood of directly striking an aircraft-sized target using a vertical attack profile. This sort of diving attack would be typical for a higher altitude sUAS or loitering weapon that descends rapidly on the target, such as the Switchblade munition. We assumed that an attacker could strike anywhere on a large aircraft-sized target (we used a KC-135 as a surrogate for the surface area of a large aircraft and an F-16 to represent hits that would produce high levels of damage to the key regions of a larger plane. See Figure 3.16).

Figure 3.16. Illustration of a Smaller, Critical Cross Section of a Larger Simulated Aircraft



For the smaller cross section and the CEPs for incoming targeted sUAS of 1 and 3 meters, the simulation resulted in a single-shot probability of a hit of approximately 0.8 and 0.5. As a matter of comparison, if we adopted the larger presented area of the target—the KC-135—then the single-shot probability of a hit trended up, as expected, with weapons of both 1- and 3-meter CEP producing a fairly uniform 0.8 probability of a hit. Further improvements in accuracy would increase the probability of a hit for the attacking weapon.

Figures 3.17 and 3.18 show the results of an attack that was launched at parked aircraft locations, with one weapon allocated to each target. Each set of bars shows the expected number of damaged planes if a certain number of sUAS make it to their targets (i.e., countermeasures that intercept weapons or render them ineffective and protect planes effectively will reduce the number of weapons reaching their target).

Figure 3.17. Average Expected Direct Hits by Direct Attack Small Unmanned Aerial Systems at 1-Meter Circular Error Probability





Figure 3.18. Average Expected Direct Hits by Direct Attack Small Unmanned Aerial Systems at 3-Meter Circular Error Probability



A salvo of ten reliable 1-meter CEP sUAS or loitering weapons generated approximately eight direct hits on the designated high-value section of a large aircraft (depicted by the furthest red bar on the right-hand side of Figure 3.17). If 3-meter CEP weapons were used, then approximately five aircraft received direct hits on the high-value portion of the aircraft from a salvo of ten (depicted by the furthest red bar on the right-hand side of Figure 3.18).

For each weapon that is prevented from launching or successfully reaching its intended target by jamming or other approaches, the expected return would be a reduction in the expected number of

aircraft that are struck by 0.5–0.8, depending on the final CEP of the weapon and the defined area of the target (depicted by the blue bars moving right to left in Figure 3.17 and Figure 3.18). The effects of infrastructure-based risk reduction approaches—including hangars for planes with doors that could be closed when an attack is detected or measures such as sUAS defensive nets—could also be viewed in terms of the expected number of sUAS that such measures would prevent from successfully striking their intended targets. Having all aircraft in shelters, with the assumption that one-half of the doors could be closed fast enough on the basis of expected attack detection timelines, would be the equivalent in our ten sUAS scenario of five systems failing to reach their target or a net risk reduction of approximately four directly hit planes (in the 1-meter CEP simulation).

Key Takeaways from the Example Scenarios

Although a more direct simulation of SF efforts to address IDF and off-base sUAS attack risks could allow for trade-offs among different tactical, technological, or infrastructure options, using the CATAPULTA model illustrates how SF actions, even actions that do not fully prevent an attack from being staged, could reduce risk. Our simulations posited the *effects* of SF action—a reduction in accuracy and in the number of platforms launched or that successfully reached their targets—that translated into reductions in the number of aircraft damaged in an attack. For scenarios in which aircraft carrying high-value assets had to be exposed during arming or before deployment, even relatively small reductions in the number of damaged aircraft could be extremely consequential. Furthermore, combinations of risk reduction approaches could be potentially even more effective (e.g., SF or domestic law enforcement efforts to identify and pressure launch locations to cut strike accuracy were combined with infrastructure changes that reduced the likelihood of successful strikes).

Law Enforcement Response Risk

Reflecting concerns that were raised by SF SMEs in our interviews, we used the POLICE model to explore three law enforcement response scenarios:

1. Response risk associated with constraints in manpower that can be allocated to law enforcement response functions, which could result from an increased need to cover nuclear-related SF missions.
2. A policy change intended to reduce demands on SF personnel who are responsible for law enforcement operations by identifying a subset of on-base law enforcement incidents that might be delegated to local civilian law enforcement.
3. An example technology change focused on alarm response demands that significantly improves security systems protecting on-base sites and reduces the demand to respond to high priority incidents.

The following sections discuss each of these scenarios.

Change in Risk Associated with Security Forces Manpower Constraints

In our interviews with SF SMEs, several interviewees emphasized that SF staffing levels are challenging the ability to meet all mission requirements. Because securing sensitive assets is understandably prioritized over routine law enforcement operations, the effects of current staffing levels are potentially already manifesting in increased incident response risk. We demonstrate how workload-type modeling can be used to explore these risks by running a set of simulations in which the numbers of SF in law enforcement response roles at Notional AFB are reduced, increasing the risk of slow or insufficient responses. Because the SF SMEs with whom we spoke expressed some concern that using recent incident data risked biasing the results of the analysis because of reduced activity during the coronavirus disease 2019 (COVID-19) pandemic, we also ran an excursion in which we doubled the rate of all incident types for comparison.

The SMEs similarly cautioned us not to focus only on time spent on the scene, because the time required for administrative follow-up and documentation after an incident requires a significant amount of SF responders' time. For the purposes of modeling, we used double the amount of time spent on the scene as the time required for administrative documentation to reflect longer time for more-complex incidents.⁸⁶ Because occupying a potential responder for that full amount of time likely overstates the effect of administrative work on response capacity (for high priority incidents, such work would be deferred), we ran simulations with and without the administrative tail for each incident, and the two results bracketed estimates of response risk.


Panels A and B of Figure 3.19 show the overall results for the baseline simulation using incident frequencies that were calibrated on the basis of multiple AFGSC base datasets, including the tail administrative work that occupies one responder after each incident. The results show that, as the number of units dedicated to response decreases, the risk of longer response times (modeled as greater than ten minutes to arrive on scene) increases and the percentage of incidents without the number of desired responders available to cover those incidents increases. The average amount of time until sufficient responders are available is quite long at low staffing levels, a result driven by SF units being occupied to administratively complete past incidents before becoming available to respond again.

⁸⁶ This simplification for modeling purposes was suggested by an SF SME during the project interviews.

Figure 3.19. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels


Probability of long dispatch time and an insufficient number of immediately available personnel by SF response unit

Panel A



Error bars denote 1-standard deviation of the performance measure. Otherwise omitted for clarity. Based on steady-state simulation run for over 100 years.

Panel B




Error bars omitted for clarity. Only includes incidents for which insufficient units were available after initial dispatch. Based on steady-state simulation run for over 100 years.


Risk measures *without* the effect of post-incident administrative requirements are shown in Figure 3.20. As suggested, the occupation of response units with post-incident administrative requirements is a central driver of both the risk of long response times and insufficient numbers of response units being available for larger response activities. Removing the administrative time (i.e., turning it off in the simulation) results in the lower risk curves in Figure 3.20 (depicted by the gray traces). However, neither risk drops to 0. Given that the priority of administrative activities will differ by incident type, the types of incidents that responders might be diverted from to complete administrative work could also differ. However, the analyses with and without the administrative time should bracket the true risk in simulated conditions. A curve representing when responders could divert from low priority post-incident activities for higher priority incident response actions would fall between the grey and blue (baseline) traces in Figure 3.20, Panels A, B, and C.

Figure 3.20. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels, Ignoring Post-Incident Administrative Requirements


Panel A



Panel B




Panel C




To address the concern that demands on the SF for incident response were significantly reduced by COVID-19, we also ran a comparative excursion case in which we doubled the frequencies of all incident types to show how increases could affect response risk. Panels A, B, and C in Figure 3.21 present the results of that case compared with the baseline from the panels in Figure 3.19. Increases in demand can result in significant increases in response risk, particularly at low levels of manpower devoted to response activities. Increases in demand also increases the risk that insufficient numbers of units will be immediately available to respond to incidents that require multiple responders for safety or effectiveness.

Figure 3.21. Estimates of Response Risk for Notional Air Force Base at Different Security Forces Police Staffing Levels, Doubling Incident Frequencies from Recent Base Police Blotter Data


Panel A



Panel B



Panel C




Example Policy Change Intended to Reduce Response Risk

For AFBs in the domestic United States, the jurisdiction for law enforcement can be shared between military and civilian law enforcement organizations.⁸⁷ As a result, in some cases, civilian law enforcement organizations could be involved in responding to criminal offenses that occur on a base. As an example of a *notional* policy change that could reduce the workload for USAF SF members and, therefore, reduce the risk for incidents those members respond to, an option to cede responsibility for some types of incidents to local civilian law enforcement was explored. We did this for the simulation by dropping a set of incident categories from the list of possible incidents that can occur. We dropped serving warrants and responding to vehicle accidents, property damage cases, and criminal cases from the model, since those incident categories seemed credible to group together for the demonstration. Figure 3.22 presents the results compared with the baseline of all incident types. Because these incidents represent a relatively small portion of the demand on the SF, delegating those incidents to local law enforcement results in only a small reduction in risk and even then, only for the cases with the tightest manpower.

⁸⁷ Judge Advocate General's School, *The Military Commander and the Law*, 15th ed., Air University Press, 2019.

Figure 3.22. Effect of Delegated Response, Percentage of Incidents with Dispatch, plus Travel Time Greater Than Ten Minutes



Example Technology Change to Reduce Response Risk

Because of the security requirements that are associated with the critical and sensitive facilities present on some bases, responding to alarms warning of a potential unauthorized entry to those facilities is an extremely high priority for the SF. Responding to such alarms is the only type of incidents for which there are defined response time requirements for SF law enforcement. The number of such incidents can pose a challenge for the SF.⁸⁸

To illustrate this dynamic and how a hypothetical change in technology—in this case, an improvement in the alarm systems that considerably reduces the number of false alarms—could affect the risks of slow or insufficient responses to other incidents, we ran a set of simulations in which the number of alarms was cut in half (i.e., a 50 percent reduction from the baseline frequency). Figure 3.23 and Figure 3.24 compare the results of the baseline simulation with this excursion case. Particularly at low numbers of available police units for response, reductions in the number of (false) alarms significantly reduced the risk of incidents with long response times and reduced the probability of insufficient numbers of units being available for incident response at dispatch.

⁸⁸ In our data, alarms represented more than 50 percent of all incidents responded to by the SF.

Figure 3.23. Percentage of Incidents with Total of Dispatch and Travel Time Greater Than Ten Minutes





Figure 3.24. Percentage of Incidents Without the Desired Number of Response Units Available at Dispatch



Key Takeaways from the Example Scenarios

Response time to alarms protecting sensitive facilities is a defined performance measure for the SF. Even though similar performance standards are not in place for other incidents, the same considerations that motivate response time requirements in civilian policing still apply: For life-safety incidents, having trained law enforcement on scene quickly matters a great deal. The prototype POLICE model showed how manpower constraints can make rapid response very difficult, creating response risk that must be managed with increasingly smaller numbers of SF devoted to this role. However, strategies for managing that risk go beyond increasing manpower. Our modeling illustrated how improvements that reduce false alarm response demand could significantly reduce risk. Equally important is that this approach to modeling can also show when the effects of alternative approaches—illustrated with the example of delegating some incidents to civilian police—might only have a modest risk reduction effect.

Chapter Summary

The prototype tools that were demonstrated in this chapter show the potential value that outcome-focused SF risk assessment can contribute to effects-based security planning. Although the initial framing of the study was focused on manpower, the outcome-focused analyses enabled effects-based planning that brings together the effects of not just constraints on personnel numbers but also how security risks can be shaped by other technological and infrastructure strategies.

For example, the results of analyses such as our incursion scenarios demonstrate how residual risk can be estimated, making it possible for an installation commander to make an informed risk tolerance decision. The difference between a 20 percent chance of an attacker reaching a sensitive target differs considerably from a 50 percent chance. Similarly, having eight attackers make it far enough on base to engage a nuclear asset's dedicated defense is a more serious concern than only two attackers reaching the terminal defense. These types of results can also show where the points of diminishing marginal returns for risk reduction start to occur with additional investments in manpower or other capabilities. Such points where the risk curves produced from the modeling begin to flatten out suggest where limits might be reached and where additional investment might not substantially reduce risk.

However, in the course of security planning for nuclear assets, these approaches also provide a way to show how the broader base SF contribute to the protection of those assets, making it possible to explore how those weapons and platforms could have higher levels of protection than defined in current standards and requirements. Assessing that contribution is consistent with treating all base SF as a system, with all SF potentially able to respond and protect all assets on base. And these approaches can demonstrate important, if not always intuitive, trade-offs: Although the readiness risk to nuclear systems from delays in maintenance operations might not traditionally be viewed as part of security operations, looking at missile field SF operations in a unified way makes it possible to explicitly show the link between managing the delayed maintenance risk and expected security risk management tasks.

Although our approaches generally were implemented in a parametric rather than in an explicit way, our top-down approach for looking at the outcomes of SF risk management also enabled looking at more-complex trade-offs, the effects on personnel readiness and proficiency, and the interaction of technological and other measures with personnel-based security activities. For example, the simulation of the response to routine law enforcement incidents demonstrated changes that could reduce risk (e.g., improvement in security alarms that reduced the frequency of alarm response) and changes that had minimal effects (e.g., the delegation of some types of incidents to civilian law enforcement). When security plans are being built, the ability to distinguish between promising and less promising options could help minimize costs and disruption as the SF continue to innovate.

Conclusions

The primary impetus for this study was concern that personnel shortfalls, notably an unwillingness to increase manning levels to accommodate new SF requirements, could put the safety of USAF personnel and assets at risk. Although personnel constraints and the push to do more with less people is not unique to the SF, the manpower intensity of many security roles and the need to respond to incidents with a sufficient number of people for safety and effectiveness makes such constraints potentially more difficult. Furthermore, the AFGSC has a set of SF requirements that are not obligatory for the remaining parts of the USAF. The requirements for protecting U.S. nuclear assets and related systems are not flexible, and the need to provide that required protection for nuclear facilities and the ability to surge to cover transport, training exercises, and nuclear mission support operations stresses the SF. In our discussions with SF SMEs, it is clear that, at current manning levels, covering all security requirements is already a challenge, and the approaches that are required to do so (including 12-hour shift operations) are stressing the force.

Trends that could increase those requirements could make it even more challenging for the SF. Although less of a concern for the AFGSC than for some other parts of the force, increases in the need for the SF to deploy during contingencies would make protecting domestic bases and assets more difficult. For the AFGSC, the rollout of new weapon systems (i.e., Sentinel) and new weapon generation facilities would significantly increase both short- and long-term requirements for nuclear asset protection. Meeting such demands could draw forces away from other base security roles, increasing the risk associated with day-to-day law enforcement incidents but simultaneously weakening protection against on-base incursion threats. Increasing demands could also further strain SF surge capacity, making it more difficult to sustain generation or other operations that require significant increases in on-duty personnel to meet security requirements.

Outcome-Focused Analysis Can Support Effects-Based Security Forces Planning

Without tools that look at the effects of different manpower levels and other security planning choices on SF outcomes, the consequences of such shortfalls might be unappreciated until a damaging incident occurs. Although resources are always constrained, which means that risk will never be eliminated, risk acceptance decisions should be made deliberately. Placing fewer personnel in a specific security role or accepting some level of fatigue in the SF, reducing proficiency and effectiveness, might not be avoidable under certain operational or other constraints. However, decisionmakers need ways to consider whether the consequences of these types of decisions are large or small and whether the

associated risk they are accepting is tolerable. This is particularly important, given the wide range of risk-reducing outcomes the SF are tasked to achieve, such as the following:

- preventing harm to nuclear systems or costly aircraft from direct attack or from IDF incidents
- protecting personnel from violent threats
- protecting sensitive facilities and information by providing a rapid response to potential security incidents
- supporting nuclear forces maintenance operations to maintain their readiness
- contributing to the day-to-day operations of installations by addressing crime- and personnel safety-related incidents, including medical emergencies, accidents, and interpersonal conflict or violence.


Across these varied risks, the prototype models demonstrated the value of top-down analytic tools to support decisions in four ways. The models can

- show the outcomes of SF actions as different types of risk reduction
- demonstrate the tangible effects of key SF concerns
- show the trade-offs among different risks when SF capability is constrained
- explore how the SF in one role can reduce multiple risks simultaneously.

The Models Show Outcomes of Security Forces Actions as Different Types of Risk Reduction

The prototype tools developed in this project demonstrated approaches to estimate the effects of the SF on different risks through either the probability of damaging events occurring or by reducing the consequences of those events. These effects are summarized in Figure 4.1. The ability to relate changes in SF personnel, resources, or capabilities to risk reduction outcomes is critical for command decisionmaking and for assessing trade-offs among different strategies or resource allocation options. Being able to estimate that the consequences of a decision would result in a 5 percent decrease in the probability of the SF stopping a base incursion or a 10 percent improvement in response times to serious incidents on base makes decisionmaking more tangible. That tangibility and the ability to trade off risk reduction in one area against greater risk in another reduces the chance of a commander making implicit risk tolerance decisions with unexpected downstream consequences.

Figure 4.1. Different U.S. Air Force Assets Protected by the Security Forces, Risk Scenarios Included in the Prototyping Effort, and Types of Security Forces Risk Reduction Explored



SOURCES: The icons (clockwise from top left) are sourced from Microsoft, Lance Weisser at The Noun Project, and Luisa Iborra at The Noun Project.

The Models Can Demonstrate Tangible Effects of Key Security Forces Concerns

In our interviews with SF SMEs, the issue of personnel proficiency—and the effect that stress on the force might have on the ability of the SF to respond to critical incidents—was raised as a serious concern. Even if forces are meeting formal training requirements, if long shifts or other stresses are degrading the *effectiveness* of the forces, it is logical that the ability of the SF to respond to incidents and reduce risks to personnel, platforms, or nuclear assets would be reduced. However, without a way to connect that logic to potential outcomes, proficiency losses remain an abstract challenge—a challenge that might be neglected in the face of other immediate challenges and priorities. Modeling that estimates risk reduction outcomes—even with proficiency represented very simplistically, as done in this effort—provides a way to explore how significant those losses might be, and, therefore, the costs of not addressing those risks.

The Models Can Show Trade-Offs Between Different Types of Risk Under Security Forces Capability Constraints

When SF actions are involved in addressing multiple risks to critical assets, shortfalls in personnel can require trade-offs between which risks are reduced and which are accepted. In the case of the missile field simulation, the prototype demonstrated this trade-off directly, showing how there are tensions between allocating staff for rapid incident response and having enough personnel to cover maintenance requirements. But the reality is that additional such trade-offs exist, particularly with how requirements for nuclear security are defined. Specific and rigid requirements for the number of personnel who must be placed at on-base nuclear facilities (e.g., the WSA within our notional base's perimeter) mean that those personnel cannot respond to base incursions that target other locations, to active shooter incidents, or other threats. Without ways to estimate risk reduction, such trade-offs are more difficult to frame and consider in decisionmaking, which leads to the risk that choices are made without full knowledge of the risk consequences.

The Models Can Show How Security Forces in One Role Can Reduce Multiple Risks Simultaneously

Despite the significant requirements for nuclear protection and force trade-offs for other risks, the types of models developed in this study also suggest that the harshness of those trade-offs could be lessened, at least to some extent. This would be in line with concepts of effects-based security and the treatment of base protection as a single system, which has all SF personnel contributing to the management of security risks. Within the perimeter of USAF installations, assets are protected by multiple layers of security that are made up of SF in a variety of individual roles. Personnel placed at perimeter checkpoints are easy to view from that perspective. However, when the whole base security effort is viewed as a system, SF personnel conducting day-to-day policing are also providing protection for critical facilities or nuclear assets on-base: In the event of a threat to those assets, those SF personnel will pivot to engage the threat.

Bringing together the results of our models, SF personnel in routine law enforcement roles might be increased *if the role is explicitly recognized*. Those personnel would simultaneously

- reduce risk measures associated with responses to routine law enforcement incidents (including critical alarm responses to sensitive facilities)
- represent a larger force of mobile responders for active shooter threats, reducing the casualty risk from those incidents
- create a larger mobile force to respond to base incursions that target sensitive or nuclear assets, reduce the risk of incursion attackers reaching their intended targets, and attrit the attackers that did reach their targets.

Although more-detailed simulation and analysis is clearly needed, our results suggest that SF engaging a Red OPFOR well before Red reaches its target can make a real and substantial contribution to the protection of facilities and assets, which could be considered in the development of requirements. Given existing and expected resource and personnel constraints, the identification of opportunities to reduce multiple risks simultaneously is likely to become more important over time.

Considering Future Steps and Applications

Because of the scale of this study and the variety of SF roles and responsibilities for risk management, our scope was necessarily constrained. We simulated different pieces of SF operations at our Notional AFB in isolation. Although we set the personnel numbers in those simulations in a way that was informed by actual practice, the timeline of the project did not allow us to take the logical next steps: define the entire SF complement of our base, allocate those forces among positions and tasks as a whole force, and then simulate the effectiveness of that laydown of SF personnel against the various threats. Taking those steps would more clearly demonstrate the trade-offs that do exist and explore strategies to either manage the consequences of those risks or show what levels of risk must be tolerated when one defensive strategy is chosen over another. Such an integrated effort would also more clearly show how SF personnel and other investments could reduce multiple risks simultaneously and support efforts to increase the risk management effects within operational, technical, and personnel constraints. The fact that even the basic prototype models developed here could successfully examine multiple key SF planning challenges suggests that continued effort to improve tools for effects-based planning is warranted and could make a tangible contribution to manpower, planning, and resource allocation decisionmaking.

Furthermore, although our decision to use a notional base was useful for several reasons, our base visits and discussions with SF SMEs drove home the point that the circumstances and requirements of each base are unique. Although some differences are quite obvious (e.g., bases with missile fields versus those without), other differences—such as in terrain, local environment, nature of base populations—can be more subtle. Such differences would necessarily cause distinctions in the threat and risk environment and could mean that some approaches for managing risk would be viable in some locations but not in others. Existing bottom-up planning tools (e.g., tools supporting the *Enterprise Protection Risk Management*) allow installations to plan against their local laydown of facilities and local threats and so, too, should complementary top-down tools for evaluating how plans for protecting installations might perform if challenged.

In addition, our discussions with SF SMEs flagged several other issues that could be explored with modeling and simulation. Although our example analyses touched on the potential for base infrastructure changes to reduce pressure on the SF, additional security-focused design shifts (including potential ways to improve the protection of nuclear assets with lower manpower demands) could be explored. Similarly, innovation in security technologies features prominently in thinking about ways to improve base protection at reduced manpower levels. For example, the evaluation of efforts to use robot dogs to augment SF activities has been done at several installations, and innovations in sensor technology also have potential.⁸⁹ However, our discussions emphasized that new technologies have their own manpower tail for operations and maintenance that must be considered when viewing them as a substitute for uniformed SF personnel.⁹⁰

⁸⁹ See discussion in Anabel Del Valle, “Robot Dogs Arrive at Tyndall AFB,” 325th Fighter Wing Public Affairs, March 29, 2021.

⁹⁰ In addition, the requirements for the use of technologies in the protection of sensitive assets were also flagged as an innovation challenge for the SF.

Finally, the application of such models as those we developed could be used to refine the way in which certain SF manpower requirements are determined. For example, rather than simply setting SF requirements based on the number of gates at a facility and how long those gates are going to be open, manning requirements could be adjusted to reflect the level of threat that those personnel are expected to be ready to repulse while they are at those gates. Similarly, as new technologies that have the potential to augment or enhance SF operations emerge and are adopted, models such as those presented here could be used to refine manning factors and inform investment decisions.

Across these areas, the prototype tools we developed, although sufficient to show key trade-offs and opportunities, can potentially be viewed as a foundation on which future efforts to support effects- and outcomes-based security planning for USAF installations can be built.

POLICE Law Enforcement Incident Response Risk Model

This appendix summarizes the POLICE Law Enforcement Incident Response Risk Model that we used in previous chapters to estimate the risks associated with changes in SF manpower, available technology, and duties.⁹¹ We first describe the architecture of the discrete event simulation model, followed by a description of the input data.

Model Architecture

We employed discrete event simulation to model law enforcement incident response.⁹² Discrete event simulation is a well-established method in operations research to study queueing systems and has been used to simulate civilian police department calls for service to obtain statistical estimates of performance measures, such as response delay and officer use.⁹³ During a discrete event simulation, a computer models a sequence of events, herein considered as the commencement and resolution of incidents, and tracks changes in system variables (e.g., available responders, workload, average response delay).

The architecture for the model is shown in Figure A.1. Given a starting number of police units at time $t = 0$, a schedule of incidents or incident history from time $t = 0$ until the simulation end time t_F and assumptions regarding incident responses at an installation (such as the recommended manpower for a given type of incident), the simulation computes various time (e.g., dispatch delay, time to the scene, etc.), risk, and use measurements that describe the overall performance of the base's law enforcement response.

Description of a Model Run

The generation of the incident history and assumptions regarding incident responses at an installation are discussed in the section titled "Input Data and Assumptions" in this appendix. A model iteration proceeds according to the flowchart in Figure A.1. First, assuming that the simulation


⁹¹ See, for example, the review in Wilson and Weiss, 2014.

⁹² Barry L. Nelson, *Foundations and Methods of Stochastic Simulation: A First Course*, Springer, 2013; Averill M. Law and W. David Kelton, *Simulation Modeling and Analysis*, McGraw-Hill, 1982.

⁹³ J. Paul Brooks, David J. Edwards, Toni P. Sorrell, Sudharshana Srinivasan, and Robyn L. Diehl, "Simulating Calls for Service for an Urban Police Department," *Proceedings of the 2011 Winter Simulation Conference*, Institute of Electrical and Electronics Engineers, December 2011.

clock t is less than t_F , the simulation handles the next event on the calendar. There are two types of events: the commencement of a new incident or the resolution of an ongoing incident.

Figure A.1. POLICE Law Enforcement Incident Response Risk Model Architecture



The commencement of a new incident marks the time that a request for law enforcement response is initiated. Before a unit is dispatched to a new incident, the simulation advances the simulation clock to the current time and determines the incident type and priority. We discuss the categories and priorities for incidents in the “Input Data and Assumptions” section later in this appendix (see also Table A.1), but different types of incidents can be broadly categorized as criminal justice–related, emergency response–related, SF-specific, or miscellaneous. Categories can also be further subdivided.

We also assumed that the type and priority of the incidents were not modified after the initial determination. Next, after determining the nature of the incident, the simulation checks whether a police unit is immediately available to respond. If no units are available, the incident is added to a

dispatch queue. The dispatch queue contains all incidents that are awaiting the assignment of a police unit during the simulation run. Alternately, if at least one unit is available, the simulation assigns a unit to that incident. On the basis of assumptions about incident responses at the installation, including travel time, recommended manpower, time at the scene, and time processing paperwork, the simulation determines (stochastically) the resolution of the incident and schedules the resolution event in the calendar. The simulation then processes the next event on the calendar, assuming that the simulation clock t remains less than t_F .

The second type of event on the calendar is the resolution of an ongoing incident. The simulation advances the clock to the current time and removes the resolution event from the calendar. Next, it removes the police unit that was assigned to the resolved incident, which frees that unit to respond to other incidents. The simulation then checks whether the dispatch queue is empty. If the dispatch queue is not empty, which signifies that at least one incident is awaiting a responding police unit, the highest priority incident is then assigned to an available unit and the simulation determines (stochastically) the resolution of that incident and schedules the corresponding resolution event in the calendar. The simulation then processes the next event on the calendar, assuming that the simulation clock t remains less than t_F .

Simulation Outputs

During a run, the simulation tracks the following metrics for each incident:

1. Dispatch time: The amount of time that elapses between the start of the incident (i.e., the initiation of a request for a unit) and when a unit is first assigned to the incident. This is equal to the amount of time that an incident remains in the dispatch queue and is also known as the queuing time.
2. Travel time: The amount of time that elapses between when a unit is assigned to the incident and when the unit arrives at the incident location. As we will describe, this metric is dependent on an exogenous input model for travel time and assumptions regarding incident responses on the installation.
3. Time on the scene or service time: The amount of time that elapses between when a unit arrives at an incident location and when a unit finishes processing the incident at that location. This metric is dependent on an exogenous input model for service time and assumptions regarding incident responses on the installation.
4. Administrative time: The amount of time that elapses after the incident has been processed at the scene until all remaining administrative and other requirements to fully resolve the incident are completed. This metric is dependent on assumptions regarding incident responses on the installation.
5. Additional available units: The number of additional units that are available at the time the first unit is assigned to the incident. This metric represents an additional buffer or margin of safety for incidents that is available to reduce the risk of harm to personnel who respond.
6. Total available units: Whether the number of additional available units, along with the assigned unit, exceed the recommended manpower for that incident at the time of assignment.

7. Time until a previously assigned unit becomes available: The expected amount of time from when a unit that was assigned to process a different incident prior will become free and able to assist with the incident under consideration. This duration is calculated only for incidents that do not have the recommended number of additional units available at the time the first unit is assigned to the new incident. The recommended number of units is an input that we will describe in the “Input Data and Assumptions” section below. The duration is a measure for the amount of time a unit is at risk when responding to an incident with a margin of safety that is below the recommended number of available units. This is an easily computable upper bound for the amount of time for when the next, rather than previously assigned, unit becomes available.

Software

We implemented the POLICE Law Enforcement Incident Response Risk Model using Julia 1.7.1, a high performance, scientific computing language that was developed using Visual Studio 1.70.2, an integrated development environment for software development.

Input Data and Assumptions

This section describes how the input data used in the POLICE Law Enforcement Incident Response Risk Model simulation was obtained or modeled. Types of input data include incident blotter data to estimate incident frequency and service time and SME data on incident priorities, recommended manpower, service times, and other assumptions.

Incident Blotter Data

We received incident blotter data directly from one anonymous USAF installation, partial incident blotter directly from another installation, and data about the two installations and others from our discussions with SMEs for calendar year 2021. Table A.1 presents the incident data from one installation that provided blotter data that summarized the types, start times, and end times for each incident. The data were incomplete for February 2021 and excluded breaches, gate runs, unauthorized entries, and munitions escorts because those events were not considered law enforcement incidents. Priority values for each incident were determined during PAF’s interviews with SMEs. Security alarms constituted the highest priority and most frequently responded to incident, followed by types of emergency incidents and serious criminal incidents.

Table A.1. Summary of Received Incident Data

Incident Type	Incident Category	Priority (low, high, highest)^a	Recommended Manpower^a	Frequency	Percentage
Security alarm or alarm failure	Alarm response	Highest	3	802	57.78
Patrol response	Other criminal justice	Low	1	201	14.48
Medical emergency	Criminal justice or life safety	High	3	132	9.51
Vehicle accident	Other criminal justice	Low	2	93	6.7
Fire response	Criminal justice or life safety	High	3	50	3.6
Warrant	Other criminal justice	Low	3	28	2.02
Criminal	Other criminal justice	Low	2	24	1.73
911 emergency services call	Criminal justice or life safety	High	2	12	0.86
Domestic	Criminal justice or life safety	High	2	10	0.72
Eagle eye	Other criminal justice	Low	1	10	0.72
Property damage	Other criminal justice	Low	2	10	0.72
Assistance rendered	Other criminal justice	Low	1	7	0.5
Abandoned vehicle	Other criminal justice	Low	1	3	0.22
Death	Criminal justice or life safety	High	8	2	0.14
Ground emergency	Criminal justice or life safety	High	3	2	0.14
DUI	Criminal justice or life safety	High	2	1	0.07
Drug	Other criminal justice	Low	2	1	0.07

NOTE: The incident blotter data cover incidents that took place from January 1, 2021 to December 31, 2021, excluding February 2021. Frequency and incident data were from one anonymous USAF installation and priority and manpower data were from discussions with SMEs. The data exclude breaches, gate runs, unauthorized entry, and munitions escorts.

^a Priority values and recommended manpower values were elicited during SME interviews.

Modeling, Generation, and Calibration of Incident Times and History

Interarrival times, which represent the time between successive incidents of the same type, are sampled from probability distributions fitted to the data that were received from one USAF installation (summarized in Table A.2). Interarrival time equivalently determines the relative frequency of the type of incident and the incident start time. Our choice of time-independent probability models instead of time series (e.g., autoregressive moving average [ARMA] or autoregressive integrated moving average [ARIMA]) or nonstationary stochastic process methods reflects the limitations of the data for each incident category.⁹⁴ Exploratory data analyses using linear models were insufficiently powered to detect significant seasonal, weekly, or monthly fluctuations for each incident type. Such effects are likely present and could be revealed by additional data, however, as we will discuss, the use of time-independent probability models generates synthetic data that match the original data with respect to the relative frequency of the incidents.

For incident types with a frequency greater than 10, we fit exponential, Weibull, and lognormal distributions using the method of maximum likelihood and proc univariate in SAS 9.4. We rejected the fit if the Kolmogorov-Smirnov statistic yielded p -values less than 0.05.⁹⁵ For incident types with a frequency less than 10, we assumed and fitted an exponential distribution, which is a widely used distributional assumption for queuing systems.⁹⁶

⁹⁴ ARMA and ARIMA are well-established time series regression methods to statistically estimate trends and are used widely in such fields as finance and econometrics. Another well-known reference is George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung, *Time Series Analysis: Forecasting and Control*, John Wiley & Sons, 2015. For our application, the use of time series was not best suited for the data and the relatively low number of incidents in some categories. Lawrence Leemis, "Input Modeling Techniques for Discrete-Event Simulations," in Brett A. Peters, Jeffrey S. Smith, D. J. Medeiros, and Matt W. Rohrer, eds., *Proceedings of the 2001 Winter Simulation Conference*, Institute of Electrical and Electronics Engineers, December 2001.

⁹⁵ The Kolmogorov-Smirnov test is a well-known statistical test for goodness of fit that compares univariate empirical data with a known theoretical probability distribution. The null hypothesis indicates that the empirical data were consistent with data sampled from the known theoretical distribution and a rejection of the hypothesis (indicated by small p -values) indicates that the proposed distribution might not be a good fit. For the original reference, see N. Smirnov, "Table for Estimating the Goodness of Fit of Empirical Distributions," *Annals of Mathematical Statistics*, Vol. 19, No. 2, June 1948.

⁹⁶ The exponential distribution with rate parameter λ , given by the cumulative distribution function $1 - e^{-\lambda x}$, is a probability distribution that is commonly used to model the amount of time between successive events that occur regularly at a constant average rate and independently. The distribution retains several attractive mathematical properties that make it ubiquitous in probabilistic modeling. However, the distribution is sometimes too rigid to fit empirical data and other distributions, such as Weibull and Gamma, which relax some of the assumptions of the exponential distribution. The uniform distribution, with bounds $b > a$, and the probability density function, given by

$$\frac{1}{b - a}$$

for $a \leq x \leq b$ and 0, otherwise describes univariate data that is arbitrarily and equally probable in between the two bounds. It is used widely and commonly as a prior to indicate lack of knowledge. The triangular distribution is an extension of the distribution with an additional mode parameter of c with $a \leq c \leq b$. The inclusion of the mode parameter allows greater specificity than the uniform distribution and, oftentimes, both the bounds and the most likely value are easily obtainable from various data sources, including SME responses. For additional information, please refer to David Johnson, "The Triangular Distribution as a Proxy for the Beta Distribution in Risk Analysis," *Journal of the Royal Statistical Society: Series D (The Statistician)*, Vol. 46, No. 3, 1997; and Samuel Kortz and J. René Van Dorp, "A Novel Method for Fitting Unimodal Continuous Distributions on a Bounded Domain Utilizing Expert Judgment Estimates," *IIE Transactions*, Vol. 38, No. 5, 2006.

Table A.2. Summary of Interarrival Distributions

Incident Type	Incident Category	Priority (0 = highest)	Interarrival Time Distribution (days)	Kolmogorov-Smirnov (<i>p</i>-value)
911 emergency services call	Criminal justice or life safety	1	Exponential (19.71622) (scale parameter)	$p > 0.500$
Abandoned vehicle	Other criminal justice	2	Exponential (121.67)	Sample size too small
Assistance rendered	Other criminal justice	2	Exponential (52.4)	Sample size too small
Criminal	Other criminal justice	2	Exponential (10.74722)	$p > 0.500$
DUI	Criminal justice or life safety	1	Exponential (365)	Sample size too small
Death	Criminal justice or life safety	1	Exponential (182.5)	Sample size too small
Domestic	Criminal justice or life safety	1	Exponential (34.96)	$p > 0.500$
Drug	Other criminal justice	2	Exponential (365)	Sample size too small
Eagle eye	Other criminal justice	2	Lognormal (2.121718, 2.064783) (mu sigma)	$p > 0.150$
Fire response	Criminal justice or life safety	1	Weibull (4.7023, 0.7003674) (scale, shape)	> 0.250
Ground emergency	Criminal justice or life safety	1	Exponential (182.5)	Sample size too small
Medical emergency	Criminal justice or life safety	1	Exponential (2.235)	$p > 0.075$
Patrol response	Other criminal justice	2	Weibull (1,3707, 0.8347)	$p > 0.100$
Property damage	Other criminal justice	2	Exponential (25.635)	$p > 0.500$
Security alarm or alarm failure	Alarm response	0	Weibull (0.327, 0.787)	$p > 0.100$
Vehicle accident	Other criminal justice	2	Exponential (3.06575)	$p > 0.500$
Warrant	Other criminal justice	2	Exponential (9.61)	$p > 0.075$

NOTE: Exponential distribution was parameterized by scale parameter λ^{-1} , lognormal distribution was parameterized by μ and σ , and Weibull was parameterized by scale parameter λ and shape parameter k .

At the beginning of each model run, the selected distributions are sampled to generate a history or calendar of incidents that begin at simulation time $t = 0$ until t_F . To validate the input models, we generated a calendar of incidents to compare the frequencies obtained by the input models with the actual data using a chi-squared test. The result of the test indicated no statistically significant difference in the frequencies between the generated and the actual data.

Modeling of Travel Times

Travel time to the incident location was estimated using a mixture of probability distributions for Notional AFB, which was based on an analysis of an anonymous AFGSC installation with similar dimensions and features. Notional AFB was divided into different regions, such as base operation areas, residential areas, and other densely populated areas. We assumed that for interregional travel, a unit would travel 5 to 10 mph and take a uniformly distributed amount of time that corresponded with the origin and destination. To preserve the anonymity of the installation, we did not identify or describe any of the regions. We assumed that a police unit began in the same region that corresponded with where that unit is typically stationed. If the unit traveled between regions on the base, we assumed that the unit would travel 15 to 20 mph and that the intraregional travel time would follow triangular distribution. A triangular distribution generalizes the uniform distribution and allowed us through its parameterization to incorporate the most likely, best case, and worst case intraregional travel times into the probability model (instead of just best case and worst case, as for the uniform distribution). The most likely, best case, and worst case travel times were assessed with data sampled from Google Maps. Finally, for each incident, we determined the most likely region where that incident would occur and weighted the intraregional travel times accordingly using mixture distributions. For example, 50 percent of security alarm incidents were expected to occur in the unit's starting region and 50 percent in another region. The mixture distribution weights

Triangular(0.2,13.9,2.1) (representing travel in the starting region)

the same as

[Triangular(0.01,4.9,0.4)+Uniform(3,6)] (representing an equal chance of travel to a different region from the start).

In another example, one-third of 911 emergency services calls could occur in the region where the police units are initially assigned; another one-third in a different region, requiring both interregional travel and intraregional travel; and the last one-third in another region, also requiring interregional travel and intraregional travel. Thus, the travel time distribution for a 911 emergency services call would be described by the mixture

$1/3 * \text{Triangular}(0.2,13.9,2.1) + 1/3 * [\text{Triangular}(0.01,4.9,0.4) + \text{Uniform}(3,6)] + 1/3 * [\text{Triangular}(0.06, 5.7,1.3) + \text{Uniform}(4.5,6)].$

Table A.3 shows the travel time distributions that were used for the simulation.

Table A.3. Summary of Travel Time Distributions

Incident Type	Travel Time Distribution (minutes)
911 emergency services call	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Abandoned vehicle	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Assistance rendered	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Criminal	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
DUI	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Death	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Domestic	$[\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Drug	$1/4 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/4 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/2 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Eagle eye	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Fire response	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Ground emergency	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Medical emergency	$1/4 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/4 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/2 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Patrol response	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Property damage	$1/3 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/3 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/3 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Security alarm or alarm failure	$1/2 * \text{Triangular}(0.2, 13.9, 2.1) + 1/2 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)]$
Vehicle accident	$2/4 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/4 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/4 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$
Warrant	$1/4 * \text{Triangular}(0.2, 13.9, 2.1)$ $+ 1/4 * [\text{Triangular}(0.01, 4.9, 0.4) + \text{Uniform}(3, 6)] + 1/2 * [\text{Triangular}(0.06, 5.7, 1.3) + \text{Uniform}(4.5, 6)]$

NOTE: Exponential distribution was parameterized by scale parameter λ^{-1} ; lognormal distribution was parameterized by μ and σ ; Weibull was parameterized by scale parameter λ and shape parameter k ; triangular distribution was parameterized by minimum, maximum, and mode; and uniform distribution was parameterized by minimum and maximum.

Modeling of Time on the Scene

Similar to the modeling of incident times and history, we used time-independent probability models to model the amount of time that a unit was on the scene. We had multiple sources for this data, including the blotter incident data we received from an anonymous installation and SME input from two installations. Table A.4 summarizes the distributions for time on the scene by source. For the incident data, we fitted exponential (using the method of maximum likelihood), Weibull, lognormal, and uniform distributions and selected the first distribution for which the p -value yielded by the Kolmogorov-Smirnov test was greater than 0.05. If the number of incidents of a given type were fewer than 10, we selected a uniform, triangular, or exponential distribution on the basis of a visual inspection of the data. For incidents for which we could not find a distribution, we opted for a triangular distribution. The triangular distribution is widely used in simulations for its ability to generalize the uniform distribution, to incorporate nonuniform priors held by SMEs, and to retain flexibility.⁹⁷

Two SMEs were asked to provide ranges of time on the scene by incident type. For the analyses in this report, we used uniform distributions with parameters that corresponded with the minimums and maximums elicited from the SMEs. If the SMEs did not evaluate an incident type, we used the incident data as a source. Additional sensitivity analyses (not reported) indicated that there were no major qualitative differences between using the SME data or the blotter incident data that we received directly from the installation for most incident types. However, simulation runs were sensitive to the duration of security alarm incidents, which were the most frequent incident and sometimes required more than 100 hours on the scene in the incident data (these were potential data entry errors). For the stability of the results, we prioritized use probability distributions based on the SME data.

⁹⁷ Johnson, 1997; Kotz and Van Dorp, 2006.

Table A.4. Summary of Time on the Scene Distributions

Incident Type	Source 1: Incident Data (hours)	Kolmogorov- Smirnov (<i>p</i>-value) (source 1)	Source 2: SME (hours)^a	Source 3: SME (hours)	Selected Distribution
911 emergency services call	Lognormal (−0.92617,0.481545)	$p > 0.150$	1.5–8		Uniform (1.5,8)
Abandoned vehicle	Triangular (0.383,0.467,0.383) (min, max, mode)	Sample size too small			Triangular (0.383,0.467,0.383) (min, max, mode)
Assistance rendered	Exponential (0.943)	Sample size too small			Exponential (0.943)
Criminal	Exponential (2.2034)	$p > 0.100$	1.5–8		Uniform (1.5,8)
DUI	Uniform (7.082,7.083)	Sample size too small	1.5–8	1–3	Uniform (1.5,8)
Death	Uniform (2.82,3.77)	Sample size too small	1.5–8		Uniform (1.5,8)
Domestic	Exponential (2.763)	$p > 0.150$	1.5–8		Uniform (1.5,8)
Drug	Uniform (3.182,3.184)	Sample size too small			Uniform (3.182,3.184)
Eagle eye	Exponential (1.10333)	> 0.250			Exponential (1.10333)
Fire response	Lognormal (0.333,1.5)	$p > 0.150$	0.33–0.5	0.5–1.5	Lognormal (0.333,1.5)
Ground emergency	Uniform (0.49,0.51)	Sample size too small	0.33–0.5		Uniform (0.333,0.5)
Medical emergency	Triangular (0.07,3.57,0.317)	Rejected other distributions	0.33–0.5	0.5	Uniform (0.333,0.5)
Patrol response	Triangular (0.083,25.23,0.17)	Rejected other distributions			Triangular (0.083,25.23,0.17)
Property damage	Exponential (1.828)	$p > 0.500$	1.5–8		Uniform (1.5,8)
Security alarm or alarm failure	Triangular (0.0167,120.32,0.333)	Rejected other distributions	0.333–0.5	0.5–1.5	Uniform (0.333,1.5)
Vehicle accident	Triangular (0.1667,168,1.233)	Rejected other distributions	1.5–8	0.5–2	Uniform (0.5,2)
Warrant	Lognormal (0.398209,0.646449)	$p > 0.150$	1.5–8		Uniform (1.5,8)

NOTE: Exponential distribution was parameterized by scale parameter λ^{-1} ; lognormal distribution was parameterized by μ and σ ; Weibull was parameterized by scale parameter λ and shape parameter k ; triangular distribution was parameterized by minimum, maximum, and mode; and uniform distribution was parameterized by minimum and maximum.

^a Source 2 indicated that some incidents rarely take up to 8 hours and reported lower figures. We used the maximum number of hours provided.

Other Inputs and Assumptions

The three remaining inputs were the number of police units, the length of the simulation run (t_F), and the assumptions for administrative, post-processing, and paperwork time following an incident. We received post-priority charts for multiple AFGSC bases and found that they had 3 or more police units posted at any given time. In our analyses, we varied the number of units as a key parameter of interest. Moreover, we assumed that this number was fixed during any particular simulation run. Additionally, administrative time (representing post-processing, paperwork, etc.) was added to the amount of time on the scene. On the basis of SME input, administrative time was usually some multiple (most commonly 1) of time on the scene. Incidents with longer times on the scene often had commensurately longer times spent on reporting, post-processing, and other administrative tasks. For the baseline analysis, we assumed that additional administrative time is equal to time on the scene (i.e., we added 1 times the amount of time on the scene to the total response time). Moreover, we assumed that the police unit was a committed resource during this time and not available to respond to other incidents. We ran an additional analysis without administrative time in addition to the baseline analysis that did include administrative time. Finally, for the simulation effort, we executed a steady-state simulation and set $t_F \gg 0$ (10–100 years) to obtain the replication means and confidence intervals of the simulation outputs listed previously. A visual inspection of the estimates showed stability in the performance measures in less than a year for most estimates.

Model Runs and Limitations to Be Addressed in Future Work

Using the inputs in Tables A.1 through A.4, adding 1 times the amount of time on the scene to the total response time and varying the number of total police units yielded the baseline analysis and estimates of response risk for Notional AFB at different SF police staffing levels, as previously reported. A similar, subsequent analysis that doubled the frequency of the incidents (halved the interarrival times) and retained other baseline assumptions was performed to yield estimates of risk at different staffing levels. Our analysis of a delegated response was performed by omitting from the model incidents that served warrants or responded to vehicle accidents, property damage cases, or criminal cases to represent delegation to civilian law enforcement. Finally, our analysis of the response risk with improved alarms was performed by reducing the frequency of alarms, thereby representing technology improvements that potentially save SF police units time.

The model architecture has one significant limitation: It does not allow for preemption (units in middle of an incident are moved to another, higher priority incident before the current incident is resolved) and assumes that priorities and available personnel are fixed. Because preemptive queueing systems are well studied, these limitations can be overcome with additional simulation logic for unit reassignment and the disruption of ongoing work. Also, more-sophisticated methods are available to estimate steady state means and to manage input uncertainty that were not pursued in this research. Other limitations surround the incomplete or lack of representative input data. Incident data were collected from only one AFGSC installation for part of one year (during the COVID-19 pandemic) and recommended manpower, service times, and administrative times were based on SME inputs and subject to the typical limitations associated with SME data, such as various forms of bias.

Transportation times and other factors used in the simulation study might not be representative of all AFGSC installations.

DEFENSE Base Incursion Response Risk Model

This appendix summarizes the DEFENSE Base Incursion Response Risk Model that we used to estimate risks associated with changes in force ratios, available technology, and individual proficiency. First, we describe the architecture of the model and provide a high-level description of its logic. Second, we detail the mathematical and software details. Finally, we summarize how the simulation runs were generated.

Model Architecture


We employed a continuous-time simulation with a mix of continuous-state variables and event-based logic. The simulation described herein is a prototype specialized for this research on AFGSC installations. More sophisticated simulation architectures, such as OneSAF or the Advanced Framework for Simulation, Integration and Modeling (AFSIM), could be adapted to analyze response risks that are associated with defending against base incursions or active shooter incidents.⁹⁸

Figure B.1 depicts a flowchart of the DEFENSE Base Incursion Response Risk Model. We considered a covert or overt attack on an installation by the OPFOR, which could be either a single team or a single active shooter. The simulation initializes at time $t = 0$ with a starting position, velocity, and capability for the OPFOR. The simulation continues until the simulation clock reaches the end time of t_f or if another termination condition is achieved (the OPFOR is destroyed or reaches its target location with sufficient capability or force).

Additionally, SF units are initialized at time $t = 0$ and stationed throughout the base with initial velocities (typically 0), starting positions, and capabilities. There are two types of SF units: stationary and mobile. Stationary units are assigned to a particular location and will only engage the OPFOR if the OPFOR is near. It was assumed that these units would not move far from their initial posts to pursue the adversary or would pursue the adversary only if the units were along the incursion path. Examples of stationary units include gate control personnel and WSA guards. Mobile units, assuming that the OPFOR is detected, are expected to move to intercept the adversary or to move toward a rally position. Examples of mobile units include fire teams, patrol units, or mobile law enforcement units.

⁹⁸ Doug Parsons, John Surdu, and Ben Jordan, "OneSAF: A Next Generation Simulation Modeling the Contemporary Operating Environment," *Proceedings of the Euro-Simulation Interoperability Workshop*, June 2005; Peter D. Clive, Jeffrey A. Johnson, Michael J. Moss, James M. Zeh, Brian M. Birkmire, and Douglas D. Hodson, "Advanced Framework for Simulation, Integration and Modeling (AFSIM)," *Proceedings of the 2015 International Conference on Scientific Computing*, July 2015.

Figure B.1. DEFENSE Base Incursion Response Risk Model Architecture



Description of Model Iteration

The objective of the adversary is to move to a preestablished target location and retain sufficient capability or force at that location. During the simulation time interval $[t, t + \delta)$, where $\delta > 0$ is a sufficiently small timestep, the simulation proceeds as follows (assuming that $t < t_f$ and that no other termination condition has been achieved): If the OPFOR is not detected before time t , then the simulation queries a detection model (described in the “Model Formulation” section later in this appendix) to determine whether the adversary would be detected during the time interval based on its position and other assumptions. If the OPFOR is not detected, the OPFOR’s position is adjusted closer to its target based on its base speed. If the OPFOR is detected, then the mobile SF, assuming they have full knowledge of the OPFOR’s location and trajectory, adjust their velocities to intercept or to move to a previously established rally position. Stationary forces might also move to intercept the OPFOR.

If the OPFOR is detected by time t , the simulation will check if the OPFOR is within a prespecified encounter radius of an SF unit. If not, the simulation will update the positions of the OPFOR and all SF units and proceed to the next time interval. If the OPFOR is within the encounter radius of an SF unit, then an engagement occurs. During the engagement, all SF units within the engagement radius (typically greater than the encounter radius) concentrate on the adversary. Engagements are resolved instantaneously in the simulation and attrition is assigned according to a model, with differing rules for cases involving an active shooter incident or a more conventional engagement. The resolution of the engagement sometimes will lead to the destruction or the removal of the units by the model and to the modification of the units' capabilities. Assuming that no termination condition is achieved, the simulation will update all units' positions and continue to the next time interval.

At the beginning of the next time interval, the simulation will check the terminal conditions. If the OPFOR is destroyed or has insufficient capability, the run ends with a victory for the SF. If the OPFOR reaches the target location with sufficient capability, the run ends with a victory for the OPFOR. Finally, the simulation will terminate, possibly with neither side victorious once the clock reaches or exceeds t_f .

Model Formulation

This section presents the mathematical formulations that underly the simulation, including the movement dynamics, detection model, collision detection (for velocity changes), and attrition model for engagements that arise from encounters. Before detailing these topics, we define notation and present some preliminary notions and assumptions.

Preliminaries

As already stated, let $t \geq 0$ denote time, as represented on the simulation clock, and $\delta > 0$ be a sufficiently small timestep. Let $\mathcal{B}$ denote the set of all SF (blue) units, and let disjoint subsets $\mathcal{S}$ and $\mathcal{M}$ represent the stationary and mobile units, respectively. Let $\mathcal{R}$ denote the singleton set that corresponds to the OPFOR (Red) unit. We denote the set of all units as $\mathcal{U} = \mathcal{R} \cup \mathcal{B}$.

Within the simulation model, we described unit $u \in \mathcal{U}$'s position at time t using the standard two-dimensional coordinate system. By letting $x_u(t) \in \mathbb{R}^2$ denote unit $u \in \mathcal{U}$'s position at time t , we imposed the following restrictions:⁹⁹

- Assumption 1: For $u \in \mathcal{R}$, we require $x_u(t) \in \left\{ \begin{pmatrix} z \\ 0 \end{pmatrix} \mid z \in [0, a] \right\}$ for all t and some $a > 0$.
- Assumption 2: For $u \in \mathcal{S}$, we require $x_u(t) \in \left\{ \begin{pmatrix} z \\ 0 \end{pmatrix} \mid z \in [0, a] \right\}$ for all t and some $a > 0$.
- Assumption 3: For $u \in \mathcal{M}$, we impose no restrictions and allow $x_u(t) \in \mathbb{R}^2$.

⁹⁹ We summarized additional set notations as follows: The relation $\in$ indicates set membership, $\mathbb{R}$ denotes the set of real numbers, $\{\}$ defines a set using set builder notation, and $X \cup Y$ denotes the union of sets X and Y , the set that contains elements of both.

The coordinate system employed by the simulation model does not necessarily correspond to the actual two-dimensional representation of the installation. Assumption 1 limits the OPFOR's positions and movements over time to a subset of a one-dimensional subspace of $\mathbb{R}^2$ in the model. This assumption is not as restrictive as it seems. Consider the Notional AFB incursion scenario discussed in Chapter 2 and shown in Figure B.2. The path the OPFOR takes from the ECHO entry gate to the aircraft parking apron can be projected (shown at the bottom of Figure B.2) onto a one-dimensional subspace, where 0 represents entry control at the ECHO gate and a represents the OPFOR's objective, the parking apron. Assumption 2 similarly restricts the positions of stationary SF units to the OPFOR's path, representing obstacles (entry gate controllers, WSA posts, etc.) that must be surmounted by the OPFOR. Assumption 3 imposes no restrictions on the coordinates of the mobile units. However, this assumption entails an additional, significant assumption:

- Assumption 4, Mapping of Coordinates: Coordinates for the mobile units in the simulation are obtained by mapping their actual locations to $\mathbb{R}^2$ with the requirement that the intended path of the OPFOR is mapped to the subset $\left\{\begin{pmatrix} z \\ 0 \end{pmatrix} \mid z \in [0, a]\right\}$ for some $a > 0$.


For example, if $a = 100$ meters, then the vector

$$\begin{pmatrix} 50 \\ 100 \end{pmatrix}$$

represents a location 100 meters orthogonal in some direction to the midpoint of the OPFOR's path. In the modeling, we generated the coordinates for the mobile units only approximately in the model by assessing the installation maps. More-sophisticated methods of implementing the projections and the mappings are available. Moreover, Assumption 4 also implies that constant velocities in the simulation coordinates will not correspond to constant velocities in actuality, as actual paths will have turns and directional changes. We ignored this issue and did not consider further the topography or the road network of the installation shown in Figure B.2. Assumption 4 and our implementation, therefore, reduced the movement physics of all units in the model to those of particles in two-dimensions. This is a significant limitation because the intent of the model was to approximate reality and to act as a prototype for future modeling efforts or for the use of more sophisticated tools, such as OneSAF.

Finally, for a unit $u \in \mathcal{U}$, we let $y_u(t) \in \mathbb{R}$ denote the capability of the unit. Capability could refer in practice to the force size (e.g., number of personnel belonging to the unit) or relative firepower. For most of our analyses, y_u was typically interpreted as the force size.

Figure B.2. Example Red Inursion into Notional Air Force Base via Roadways from ECHO Gate to the Aircraft Parking Apron



Movement Dynamics

Let $a > 0$ and

$$\begin{pmatrix} a \\ 0 \end{pmatrix}$$

denote the coordinates of the OPFOR's target location. We now provide the starting positions, capabilities, and velocities for all units: the OPFOR at time $t = 0$ begins its incursion from the origin

$$(\text{i.e., } x_u(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, u \in \mathcal{R}).$$

For $u \in \mathcal{S}$, we require that SF stationary unit starting positions be along the OPFOR's path:

$$x_u(0) \in \left\{ \begin{pmatrix} z \\ 0 \end{pmatrix} \mid z \in [0, a], u \in \mathcal{S} \right\}.$$

For $u \in \mathcal{M}$, $x_u(0) \in \mathbb{R}^2$, all units $u \in \mathcal{U}$ have an initial capability of $y_u(0) > 0$ representing force size, firepower, or similar parameter.

Let $v_u(t) \in \mathbb{R}^2$ denote the velocity of unit $u \in \mathcal{U}$ at time t . The OPFOR unit r moves along the one-dimensional subspace that corresponds with movement along its incursion path. Denoting the OPFOR's base movement speed as $b_r > 0$, we then describe its initial velocity accordingly as

$$v_r(0) = \begin{pmatrix} b_r \\ 0 \end{pmatrix}.$$

All stationary SF units $u \in \mathcal{S}$ always have 0 velocity by assumption. Similarly, assuming that detection of the OPFOR occurs after time $t = 0$, the initial velocities of all mobile SF units $u \in \mathcal{M}$ are also zero. In the next subsections, we describe how and when in the simulation model these velocities can be adjusted. We also assume that each mobile unit m has $b_m > 0$ as a base speed when, in fact, the unit begins to move. For any unit u at time t , and given its current position and velocity, the update in the unit's position is computed as follows:

$$x_u(t + \delta) = x_u(t) + \delta \cdot v_u(t).$$

Detection Model

In this subsection, we present a stochastic model that we used to determine whether the OPFOR is detected in the time interval $[t, t + \delta)$. The model is a segment-based model in which the probability of detection depends on where along the incursion path the OPFOR is located during the time interval.

Consider the one-dimensional subset

$$\left\{ \begin{pmatrix} z \\ 0 \end{pmatrix} \mid z \in [0, 1] \right\} \text{ of } \mathbb{R}^2,$$

which describes the OPFOR's incursion path. We define a probability measure on $[0, a]$ as follows: let $0 < a_1 < a_2 \dots < a_i < \dots < a$ and define function f where

$$\int_{a_{i-1}}^{a_i} f(x) dx = c_i \in [0, 1] \text{ and } \int_0^a f(x) dx = \sum_i c_i = 1.$$

As an example of the detection model, let $a = 1$ represent the OPFOR's incursion path with a length of 1 km. Consider $a_1 = 0.05$ with

$$\int_0^{0.05} f(x) dx = 0.5 \text{ and } \int_{0.05}^1 f(x) dx = 0.5.$$

This detection model assumes that the OPFOR will be detected in the segment along the first 50 meters of the OPFOR's incursion path (perhaps by entry control) in 50 percent of model runs and in the last 950 meters of the incursion path 50 percent of the time.

Without loss of generality, we assumed that the OPFOR will be detected during a simulation run.¹⁰⁰ Let $a > 0$. Let z be a random number drawn from the standard uniform distribution at the initialization of the simulation run. Given the break points $0 < a_1 < a_2 \dots < a_i < \dots < a$ from the detection model, we select the interval $[a_{i-1}, a_i)$, such that

$$\int_0^{a_i} f(x)dx > z \text{ but } \int_0^{a_{i-1}} f(x)dx \leq z.$$

The selected segment $[a_{i-1}, a_i)$ is where detection of the OPFOR will occur during the run. To compute the actual point of detection, we make the following assumption:

- Assumption 5, Calculation of Detection Point: Let $0 < a_1 < a_2 \dots < a_i < \dots < a$ and $[a_{i-1}, a_i)$ be the segment of the OPFOR's incursion path where it is detected during a simulation run. We assume that the point of detection is uniformly distributed on $[a_{i-1}, a_i)$. Let w be drawn from the uniform distribution with parameters a_{i-1} and a_i . Then,

$$\binom{w}{0}$$

is the point of detection.

Suppose the OPFOR unit r 's position in the time interval $[t, t + \delta)$ is

$$x_r(t) = \binom{x_r^1(t)}{0},$$

and let w be as in Assumption 5. The detection check in Figure B.1 will yield that the OPFOR has been detected if $x_r^1(t) > w$.

Velocity Changes

Suppose the OPFOR unit r has been detected during time interval $[t, t + \delta)$. Before discussing the velocity changes for all SF units, we made other assumptions:

- Assumption 6, C2 Assumptions: On detection, all SF units have knowledge of the OPFOR's position and velocity, and all units act independently, if applicable, to intercept the OPFOR or move to a rally position.
- Assumption 7: Velocity changes are near instantaneous.

The lack of coordination among SF units in Assumption 6 could be appropriate in specific, chaotic situations. The modeling of partial or incorrect knowledge of the OPFOR's position and velocity requires more sophisticated approaches that use agent-based logic.

¹⁰⁰ If the OPFOR was not detected, then the logic of Figure B.1. implies that the OPFOR will reach its target if t_f is sufficiently large; there would be no need for a simulation run. Therefore, the simulation results could be interpreted as conditional on the event that the OPFOR is detected. For example, if the OPFOR is not detected 50 percent of the time but according to the simulation, loses 100 percent of the time, then the OPFOR's actual unconditional success probability would be 50 percent and its success probability conditional on detection would be 0 percent.

A stationary unit $u \in \mathcal{S}$ may or may not decide to intercept the OPFOR unit r depending on whether the model user specifies that they may. It could choose to remain stationary (with 0 velocity) or move to intercept the OPFOR along the incursion path with the following velocity for times $t' > t$:

$$v_u(t') = \begin{pmatrix} \text{sgn}(x_r^1(t) - x_u^1(t)) \\ 0 \end{pmatrix} \cdot b_u,$$

where x_r^1 and x_u^1 denote the first components of unit r 's and u 's position, respectively, and $\text{sgn}(\cdot)$, the sign function, and b_u denote the SF unit's base speed. Unit u will move with base speed b_u along the OPFOR unit's incursion path.

A mobile unit $u \in \mathcal{M}$ could move to intercept the OPFOR. We first calculated the amount of time required for u to reach the OPFOR's position. Let $d = \|x_u(t) - x_r(t)\|$ denote the Euclidean distance between the two units and $\Delta = x_u(t) - x_r(t)$ denote the displacement. Let

$$\tau = \begin{pmatrix} a \\ 0 \end{pmatrix}$$

denote the coordinates for the OPFOR's target. We calculated the angle θ between the mobile unit and the OPFOR as follows:

$$\theta = \cos^{-1} \left(\frac{\Delta^T(\tau - x_r(t))}{\|\Delta\| \|\tau - x_r(t)\|} \right).$$

The interception time t^* can then be derived using the law of cosines and the quadratic formula as the non-negative real root of the following equations:

$$t^* = \left(\frac{b_r \cos(\theta) \pm \sqrt{b_r^2 \cos^2(\theta) - (b_r^2 - b_u^2)}}{(b_r^2 - b_u^2)} \right) \cdot \|\Delta\|.$$

Assuming that $b_u^2 > b_r^2$, that the SF mobile unit is faster than the OPFOR unit, guarantees that the interception will occur (although possibly after the simulation ends because of another termination condition). The previous calculation for the interception time can be improved computationally.

The point of interception x^* can be written as

$$x^* = x_r(t) + t^* \cdot v_r(t),$$

and the new velocities for mobile unit $u \in \mathcal{M}$ in times $t' > t$ is written as

$$v_u(t') = \left(\frac{x^* - x_u(t)}{\|x^* - x_u(t)\|} \right) \cdot b_u.$$

If the mobile unit u decides to move to another rally position instead of move to intercept, then in the previous velocity update, the coordinates for the unit's rally position are used for x^* instead. Also,

as shown in Figure B.1., if a unit is sufficiently close to its rally position, then the simulation will assign that unit a 0 velocity that it retains thereafter.

Encounters and Engagements

Suppose there is at least one SF unit $u \in \mathcal{B}$ within a prespecified encounter radius R_e of OPFOR unit r in time interval $[t, t + \delta)$, which is $\|x_u(t) - x_r(t)\| < R_e$. Then, unit u is sufficiently close to OPFOR unit r so that an engagement can begin. Parameter R_e can vary for the simulation run—we assumed that it is something like 50 meters for the SF unit to fire.

During an engagement, nearby friendly units could also concentrate fire so that it is not solely the unit that encountered the OPFOR that participates in the engagement. Let R_n be the prespecified engagement radius and define the subset $\mathcal{N}_t \subseteq \mathcal{B}$ of all SF units participating in the engagement at time t accordingly:

$$\mathcal{N}_t = \{u \in \mathcal{B} \mid \|x_u(t) - x_r(t)\| < R_n\}.$$

Most often, $R_n > R_e$ to reflect that the encountering SF unit will slow the OPFOR's progress sufficiently to allow other SF units to participate, such as all units that are within 100 meters in the case of the 50 meter encounter radius.

When resolving the engagement, we made the following assumption:

- Assumption 8: Engagements are resolved instantaneously in the model.

For resolving the engagement, we presented two possible attrition models in Assumptions 9A and 9B. Recall that $y_r(t)$ is the OPFOR's current capability (or force size) and that $y_n(t)$ is the capability for the engaging SF unit $n \in \mathcal{N}_t$. The standard attrition calculation proceeds as follows:

- Assumption 9A, Standard Attrition Calculation: Let $y_r(t)$ be the OPFOR's current capability and $\bar{y}(t) = \sum_{n \in \mathcal{N}_t} y_n(t)$ be the total SF capability. If $y_r(t) > \bar{y}(t)$, then the OPFOR's new capability is given by $y_r(t + \delta) = y_r(t) - \bar{y}(t)$ and $y_n(t + \delta) = 0$ for each $n \in \mathcal{N}_t$. If $y_r(t) \leq \bar{y}(t)$, then $y_r(t + \delta) = 0$ and $y_n(t + \delta)$ are arbitrary values for each $n \in \mathcal{N}_t$.

Assumption 9A imposes a winner-take-all attrition model. If the OPFOR has more capability or force, then all engaging SF units are effectively destroyed (i.e., capabilities are set to zero) and the OPFOR's capability is decremented but remains positive. If capabilities represent force sizes with equal capability, then Assumption 9A effectively awards the engagement to the group with more forces and assumes the complete destruction of all participating forces in the other group (with ties awarded to the defending SF).

Assumption 9A is only one type of attrition model. For modeling active shooter incidents, we used an alternative, stochastic model to calculate attrition. For the (lone) active shooter model, $y_r(t)$ is assumed to equal 1 to represent the single shooter. Let $\rho \in [0,1]$ represent the proficiency rating or accuracy of the SF forces. Assumption 9B presents the attrition model:

- Assumption 9B, Active Shooter Attrition Calculation: Let n be the number of engaging SF units at time t (i.e., $|\mathcal{N}_t|$). Furthermore, let X be a random variable following the binomial distribution with n trials and success probability $1 - \rho$. If $X < n$, then $y_r(t + \delta) = 0$ and $y_b(t + \delta)$ is an arbitrary value for each $b \in \mathcal{N}_t$. If $X = n$, then $y_r(t + \delta) = y_r(t)$ and $y_b(t + \delta) = 0$ for each $b \in \mathcal{N}_t$.

The random variable X can be interpreted as the number of SF units that the active shooter is able to dodge during the engagement. The lone shooter survives (capability remains nonzero) if they dodge all the engaging units. It is assumed that the shooter is destroyed if they are unable to dodge at least one SF unit and, as in the winner-take-all assumption in 9A, that all SF units are destroyed if the shooter survives. An increase in the number of participating SF units in the engagement or an increase in ρ , which loosely corresponds to some sort of proficiency rating or accuracy, increases the likelihood that the engagement is resolved in favor of the SF.

During a simulation run, units with 0 capability on either side are considered destroyed and removed.

Terminating Conditions

The simulation run ends once the time on the simulation clock t exceeds t_f , the simulation end time. Figure B.1. shows the increasing time increments that guarantee termination. Termination can also occur in two other ways.

First, if OPFOR capability $y_r(t)$ equals 0 at any time t , OPFOR unit $r \in \mathcal{R}$ is considered destroyed and the simulation ends with a loss for the OPFOR and win for the SF. Second, let R_v be a prespecified (victory) radius, y_r^* be the minimum capability the OPFOR is required to retain if it reaches its objective, and τ be the coordinates for the target location as before: The OPFOR wins if at any time t , $\|x_r(t) - \tau\| < R_v$ and $y_r(t) \geq y_r^*$.

Thus, the OPFOR wins and the SF loses if the OPFOR reaches sufficiently close to its target (within a few meters in the implemented runs) and does so with a minimum capability of at least y_r^* (we used $y_r^* = 1$ to represent that at least a force size of 1 was required).

Software

The DEFENSE Base Incursion Response Risk Model was implemented using Julia 1.7.1 and developed using Visual Studio 1.70.2.

Model Runs


We used the DEFENSE Base Incursion Response Risk Model to analyze three scenarios for Notional AFB. Figure B.3 depicts a map of Notional AFB. Boxes B.1, B.2, and B.3 summarize the modeling assumptions for scenarios that involve Notional AFB and describe a direct assault to reach a key asset, a clandestine infiltration and an assault on a key asset, and a lone active shooter incident, respectively.

For the first scenario (detailed in Box B.1), the OPFOR force size and the number of times it was successful was varied, and the OPFOR's final capability was tracked more than 30 replications per force size. The replication means are shown in Chapter 3, along with 95 percent confidence intervals. The results are highly sensitive to assumptions on parameter values.

For the second scenario (detailed in Box B.2), two cases were implemented to illustrate the effect of detection: (1) a baseline model with a low probability of detection along the first segments of the OPFOR's incursion path near the perimeter and (2) a comparator model with all the same parameters but with detection near the perimeter significantly more likely. The number of times the OPFOR was defeated before reaching its objective was tracked more than 30 replications for each case. The replication means are shown in Chapter 3, along with 95 percent confidence intervals. The results are highly sensitive to assumptions on parameter values.

For the third scenario featuring a lone active shooter (detailed in Box B.3), the approximate coordinates for all SF units were randomly generated, the proficiency rating ρ of the SF was varied from 0 to 1, and the simulation end time t_f was set to 60 minutes. Therefore, the maximum durations of the active shooter scenarios in Chapter 3 are censored at 60 minutes. The simulation tracked more than 30 replications for each value of ρ , the number of engagements, and the amount of time until the active shooter's capability value reached 0 and triggered a termination condition because of an engagement (i.e., the shooter was destroyed by law enforcement). The replication means are shown in Chapter 3, along with 95 percent confidence intervals. The results are highly sensitive to assumptions on parameter values.

Figure B.3. Map of Notional Air Force Base



Box B.1. Direct Assault to Reach a Key Asset

Scenario

A Red team of 5 to 20 individuals breaches entry control at the FOXTROT gate and attacks the flight line (a distance of 1,660 meters).

Assumptions

- **Blue team posture** (each Blue team member was assumed to be equally proficient as each Red team member):
 - Entry control: three individuals (200 meters from the Red team entry point)
 - Near Red's objective: ten individuals (two to three fire teams, 1,500 meters from the Red team entry point)
 - Operations area (near ramp): six individuals (two fire teams)
 - Housing area and operations area boundary: six individuals (two fire teams)
 - Residential area: six individuals (two fire teams)
 - Southern operations area (near the WSA): six individuals (two fire teams)
- **Detection:** Red has a high probability of being detected near the entry gate (half the time) or is detected somewhere else along the incursion route with equal probability.
- **C2:** Blue units move independently to engage Red once Red is detected.
- **Engagement:** When opposing units are within 50 meters, engagement occurs; all units within 100 meters participate. When the forces engage, the more numerous force loses the same number of people it engages (e.g., 20 Red engaging two Blue results in two Red neutralizing two Blue and 18 Red continuing).
- **Speed:** Red is 75 percent as fast as Blue units are.

Box B.2. Clandestine Infiltration and Assault on Key Asset

Scenario

A Red fire team of four individuals surreptitiously breaches the outer fence at the BRAVO gate to attack the WSA (2,900 meters). Red is deemed successful if it reaches its objective with at least one person.

Low-Detection Variant Assumptions

- **Blue posture** (each Blue member was assumed to be equally proficient as each Red team member):
 - Operations area (near ramp): six individuals (two fire teams)
 - Housing area and operations area boundary: six individuals (two fire teams)
 - Residential area: six individuals (two fire teams)
 - Southern operations area (near the WSA): six individuals (two fire teams)
- **Detection:** Red is not detected until it approaches the WSA.
- **C2:** Blue units move independently to engage Red once Red is detected.
- **Engagement:** When the opposing units are within 50 meters of each other, engagement occurs; all units within 100 meters participate. When the two forces engage, the more numerous force loses the same number of people it engages (e.g., four Red engaging two Blue results in two Red neutralizing two Blue and two Red continuing).
- **Speed:** It was assumed that Red is traveling on foot and would take 54 minutes to reach its objective unimpeded. The Blue units were motorized and can travel 25 to 40 mph.

High-Detection Variant Assumptions

- Same as above, but Red can be detected at the perimeter fence using surveillance technologies with high probability.

Box B.3. Active Shooter on Notional Air Force Base

Scenario

An armed individual or insider enters the installation from the DELTA gate and initiates the incident while moving toward the ramp near flight line (4,500 meters from the gate). Police units and patrols respond and engage Red with varying proficiency.

Low-Detection Variant Assumptions

- **Blue posture:**
 - Operations area (near ramp): two individuals (randomly located)
 - Housing area and operations area boundary: two individuals (randomly located)
 - Residential area: two individuals (randomly located)
 - Southern operations area (near the WSA): two individuals (randomly located)
- **Detection:** Not applicable—Red is deemed detected when the incident begins.
- **C2:** Blue units move independently to engage the shooter once the incident begins.
- **Engagement:** When a unit is within 20 meters, engagement occurs. Units engage with the shooter individually. Proficiency is described as the likelihood that a Blue individual can neutralize the shooter during the engagement. The duration of the engagement is assumed to be instantaneous.
- **Speed:** The shooter moves approximately 3 mph (on foot). Blue is motorized and can travel 25 to 40 mph.

SILO PROTECT Missile Field Security Forces Risk Model

This appendix summarizes the SILO PROTECT Missile Field Security Forces Risk Model that we used to estimate risks that are associated with changes in SF manpower, available technology, and SF duties in the context of protecting missile field operations.¹⁰¹ We first describe the architecture of the discrete event simulation model, followed by a description of the input data. The architecture is similar to the POLICE Law Enforcement Incident Response Risk Model, and we focus mainly on the key elements of the SILO PROTECT model in this appendix.

Model Architecture

Similar to the POLICE model, we employed discrete event simulation in the SILO PROTECT Missile Field Security Forces Risk Model. The key differences are in how the inputs and outputs are interpreted and in a modification of the simulation logic to assign a required minimum number of personnel to particular incidents.

Description of a Model Run

Figure C.1 depicts the architecture of the SILO PROTECT model. First, assuming that the simulation clock t is less than the ending time t_F , the simulation handles the next event on the calendar. There are two types of events: (1) the commencement of a new incident or squadron action and (2) the resolution of an ongoing incident or action. The commencement of a new incident marks the time that a request for resources (i.e., personnel) from the squadron is initiated. Before the squadron members are dispatched to a new incident, the simulation advances the simulation clock to the current time and determines the incident type and priority. The types of incidents include alarm responses, launch facility checks, on-site posting because of alarm failure (camper incidents), short-term maintenance actions (e.g., fixing alarms or sensors at a silo), and long-term maintenance activities. Similar to the POLICE model, we assumed that priority levels and incident types do not change after the initial determination, and we maintained the overall assumption that the system is not preemptive (i.e., squadron members assigned to units to work a particular action or incident are not reassigned when a higher priority incident occurs).

¹⁰¹ See, for example, the review in Wilson and Weiss, 2014.

The second type of event on the calendar is the resolution of an ongoing incident. The simulation clock advances to the current time and removes the resolution event from the calendar. Next, the simulation frees a number of squadron personnel that is equal to the corresponding minimum requirement for that incident type. The simulation then checks whether the dispatch queue is empty. If the dispatch queue is not empty, which signifies that at least one incident is awaiting assignment, the highest priority incident—if there are sufficient personnel available—is assigned to an available unit and the simulation determines the resolution of that incident and schedules its corresponding resolution event on the calendar. The simulation then continues to process the next event on the calendar, assuming that the simulation clock t remains less than t_F . Like the POLICE model, the simulation only tracks the number of squadron members who are assigned or unassigned to an incident and not the locations and statuses of individuals.

Simulation Outputs

During a run, the simulation tracks the following metrics for each incident:

1. Dispatch time: The amount of time that elapses between the start of the incident (i.e., the initiation of a request for squadron personnel) and when a unit is first assigned to the incident. This is equal to the time an incident remains in the dispatch queue and is also known as the queuing time.
2. Travel time: The amount of time that elapses between when a unit is assigned to the incident and when the unit arrives at the incident location. This metric is dependent on an exogenous input model for travel time, including assumptions about travel speed and the fraction of roads that are paved.
3. Time to the scene: The amount of time that elapses between the start of the incident and when personnel arrive at the incident location. This is equal to the sum of dispatch and travel time.
4. Time on the scene or service time: The amount of time that elapses between when personnel arrive at the incident location and when the incident is resolved.
5. Maintenance events: The effect of whether a maintenance event is delayed (has a positive effect on dispatch time) or is postponed (the delay is sufficiently long that the event remains in the dispatch queue at the end of a simulation run [i.e., one year]).

Software

The SILO PROTECT model was implemented in using Julia 1.7.1 and developed using Visual Studio 1.70.2.

Input Data and Assumptions

This section describes how the input data that we used in the SILO PROTECT model simulation was obtained or modeled. Table C.1 summarizes the data we collected from SMEs at an

AFGSC installation and the input distributions that we used for SF incident and action frequency and time on the scene. The rates of incident occurrence that we elicited from the SMEs were used to parameterize the interarrival time distributions, assuming an exponential distribution. Similar to the POLICE model, the interarrival time distributions were sampled to generate an incident history. SME inputs for duration at the scene were translated into the following probability distributions: Triangular distributions were selected if the SMEs provided a range of values and a most likely value; uniform distributions were used if the SMEs provided a range of values lasting more than one day; and exponential distributions were used if the SMEs provided a range of values lasting less than one day, with the exception of short-term maintenance, which used the lower bound of the SME range for the scale parameter. The reason for these selections was to avoid variance from assuming exponential distribution (with infinite support) for some activities that should normally be completed within one day or shift (e.g., launch facility checks and short-term maintenance) and activities that take multiple days.

Table C.1. SILO PROTECT Key Model Inputs

Activity	Priority (1 = Highest)	Minimum Personnel Required	Selected Interarrival Time Distribution (days)^a	Duration at Silo (service time or time at the scene)	Selected Service Time or Time on the Scene Distribution (hours)^b
Alarm response	1	4	Exponential (0.05833)	0.75–12 hours (1.5 hours most likely)	Triangular (0.75,12,1.5)
Routine launch facility checks	2	2	Exponential (0.117)	0.5–3 hours	Exponential (1.75)
On-site posting because of alarm failure (camper)	3	3	Exponential (1.78)	1–3 days	Uniform (24,72)
Short-term maintenance activity (e.g., fixing alarms or sensors at a silo)	4	3	Exponential (0.233)	3.5–8 hours	Exponential (3.5)
Long-term maintenance activity	5	18	Exponential (0.4667)	2–7 days	Uniform (48,168)

^a Exponential distribution was parameterized by scale parameter λ^{-1} . The rates for interarrival time distribution were computed from SME estimates for incident frequency for a single missile field flight area and assumed 15 missile field flight areas. For example, a long-term maintenance event could occur once a week per missile field flight area, yielding a rate of 2.14 events a day or an average interarrival time of 0.4667 days.

^b Triangular distribution was parameterized by minimum, maximum, and mode; uniform distribution was parameterized by minimum and maximum.


Travel time to the incident location was estimated using the geography of one AFGSC installation. On the basis of that geography, travel distances between two arbitrary missile fields had

values that ranged from 0 to 65 miles along either paved or unpaved roads. We assumed that on dispatch, personnel would have to travel 0 to 65 miles, which was distributed uniformly. The SMEs indicated that they would travel an average of 55 mph on a paved road and only 25 mph on an unpaved road. Therefore, we model the travel time T (in hours), where ρ is the fraction of paved roads and U is a random variable following the uniform distribution with minimum and maximum values of 0 and 65, respectively:

$$T = U / (55\rho + (25(1 - \rho))).$$

Figure C.2 shows how, according to the model for travel time, T (in minutes) varies with ρ .

Figure C.2. Variation in Expected Travel Time from Dispatch to Incident Locations Within Two Adjacent Ten Silo and One Missile Alert Facility Portion of the Missile Field by Road Surface



NOTE: The maximum travel distance for this simulation was 65 miles between dispatch and response location. The error bars represent one standard deviation of the travel time.

Model Runs and Limitations to Be Addressed in Future Work

The baseline size for the squadron was 250 personnel based on SME interviews. This parameter was varied in subsequent analysis, along with the fraction of paved roads ρ . Moreover, for the simulation effort, we executed a steady-state simulation and set $t_F \gg 0$ (one year) to obtain the replication mean estimates and the confidence intervals of the simulation outputs listed previously. A visual inspection of the estimates showed stability in the performance measures in less than one year for all estimates.

Using the inputs listed in the “Input Data and Assumptions” section in this appendix, we assumed for the baseline that $\rho = 1$ (i.e., all roads are paved) and that the squadron size was 250. We then varied the squadron size to estimate the effect on time before different incident types could be fully staffed (including dispatch and travel time) and on the major maintenance events that were delayed or postponed because of manpower constraints.

Like the POLICE model, the SILO PROTECT model has the significant, but surmountable, limitation of ignoring preemption. Also, more-sophisticated methods are available to estimate steady state means and to manage input uncertainty that were not pursued in our research. Other limitations are linked to the incomplete or lack of representative input data. The data were collected from SMEs at only one AFGSC installation and were subject to the typical limitations associated with SME data, such as various forms of bias. Additionally, we performed no calibration or external validation of the results, apart from noting that the results for the baseline squadron size of 250 were consistent with SME impressions and that the travel times seemed reasonable based on the geography for the AFGSC installation that provide the SME data.

CATAPULTA Indirect Fire Model

The description of the model in this appendix is paraphrased and adapted from an unpublished 2013 user guide to CATAPULTA authored by RAND researcher Scott Savitz.

The CATAPULTA Indirect Fire Model is a Microsoft Excel-based model implemented in Visual Basic. For our modeling, we used the following variables in the simulation:

- The number of weapons that were fired.
- The aimpoint distribution that was used (the options are firing at a central aimpoint, the quadrants of the area being attacked, or along a line across the area).
- The level of accuracy (specified as CEP in meters).
- The geometry of the parking area for aircraft and the number of aircraft that were parked (randomly distributed among the available spaces).
- The size of the aircraft that were used (the model included default dimensions for large aircraft based on KC-135s and small aircraft based on F-16s).
- The variables that specified different levels of damage to targets for weapon impacts at different distances.

The model can also simulate personnel targets and the effect of such features as revetments for protection, which were not used in this study.

On the basis of the parameters, the CATAPULTA model employs a Monte Carlo analysis, which lays out the individual arrangements of targets, disperses weapon impact points based on the number fired and specified accuracy (including probabilistic effects in aiming and impact), and calculates the number of aircraft that are directly hit, severely damaged, moderately damaged, and lightly damaged based on the distance of impact from the aircraft. The model averaged results over a specified number of independent runs to calculate the average number of damaged aircraft, given the parameters.

Mortar Cases

For the mortar attack simulations, the modeling parameters were as follows:

- Five weapons were fired.
- The aimpoints were distributed in a line along the axis defined by the nose to tail lines of the parked aircraft.
- The baseline accuracy was 108 meters CEP (the default included in CATAPULTA for a 60 mm mortar), which was then increased twofold (216 meters) and threefold (324 meters).

- The aircraft parking area of 48 spaces (in eight rows) was populated with either 40 or 20 KC-135 sized aircraft.
- Munitions that fell within 2 meters of aircraft inflicted severe damage, within 30 meters of aircraft inflicted moderate damage, and within 67 meters of aircraft inflicted light damage (the CATAPULTA default values).

The results of 500 runs for each accuracy condition, for each number of parked aircraft were averaged to produce the reported results.

Small Unmanned Aerial Systems Cases

For the simulated sUAS attacks, the modeling parameters were as follows:

- The number of weapons fired were varied from one to ten.
- The weapons targeted individual parking spaces for the aircraft.
- Accuracy varied from 1 meter to 3 meters to simulate the high precision ability to be guided of the sUAS.
- The aircraft parking area of 48 spaces (in eight rows) was populated with 40 KC-135–sized aircraft or 40 F-16–sized aircraft (with the latter used as a surrogate core area of the KC-135 for the smaller critical footprint of a larger platform).
- Munitions that fell within 0.5 meters of aircraft inflicted severe damage, within 1 meter of aircraft inflicted moderate damage, and within 2 meters of aircraft inflicted light damage (reflecting the small explosive payloads used on the systems).

The results of 50 runs for each accuracy condition, each number of munitions, and each type of parked aircraft were averaged to produce the reported results.

Abbreviations

AFB	Air Force Base
AFGSC	Air Force Global Strike Command
BSZ	base security zone
C2	command and control
CEP	circular error probability
COVID-19	coronavirus disease 2019
cUAS	counter unmanned aerial systems
DHS	U.S. Department of Homeland Security
DoD	U.S. Department of Defense
DUI	driving under the influence
ICBM	intercontinental ballistic missile
ID	integrated defense
IDF	indirect fire
MAF	missile alert facility
OneSAF	One Semi-Automated Force
OPFOR	opposing force
PAF	Project AIR FORCE
PIRA	Provisional Irish Republican Army
PPE	personal protective equipment
SF	Security Forces
SME	subject-matter expert
sUAS	small unmanned aerial systems
UAS	unmanned aerial systems
USAF	U.S. Air Force
WSA	weapons storage area

References

- Air Force Doctrine Publication 3-10, "Force Protection," U.S. Air Force, November 19, 2019.
- Air Force Instruction 31-118, "Security Forces Standards and Procedures," U.S. Air Force, August 18, 2020.
- Air Force Instruction 36-2646, "Security Forces Training and Standardization Evaluation Programs," U.S. Air Force, December 4, 2020.
- Box, George E. P., Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung, *Time Series Analysis: Forecasting and Control*, John Wiley & Sons, 2015.
- Briar, David P., "Sharpening the Eagle's Talons: Assessing Advances in Air Base Defense Doctrine," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Brooks, J. Paul, David J. Edwards, Toni P. Sorrell, Sudharshana Srinivasan, and Robyn L. Diehl, "Simulating Calls for Service for an Urban Police Department," *Proceedings of the 2011 Winter Simulation Conference*, Institute of Electrical and Electronics Engineers, December 2011.
- Buontempo, Joseph T., and Joseph E. Ringer, "Airbase Defense Falls Between the Cracks," *Joint Force Quarterly*, No. 97, 2nd Quarter, 2020.
- Camm, Frank, Lauren Caston, Alexander C. Hou, Forrest E. Morgan, and Alan J. Vick, *Managing Risk in USAF Force Planning*, RAND Corporation, MG-827-AF, 2009. As of October 16, 2023: <https://www.rand.org/pubs/monographs/MG827.html>
- Carter, David L., and Jeremy G. Carter, "Effective Police Homicide Investigations: Evidence from Seven Cities with High Clearance Rates," *Homicide Studies*, Vol. 20, No. 2, May 2016.
- Caudill, Shannon W., ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Caudill, Shannon W., and Bryan A. Keeling, "Law Enforcement and Base Defense: Improving Interoperability to Benefit the War Fighter," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Caudill, Shannon W., Anthony M. Packard, and Raymund M. Tembreull, "Defending the Joint Force: Lessons Learned from Joint Base Balad," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Caudill, Shannon W., ed., Andrea D. Tullos, and Erik K. Rundquist, *Defending Air Bases in an Age of Insurgency: Integrated Base Defense Principles for Commanders*, Vol. III, Air University Press, 2021.
- Clive, Peter D., Jeffrey A. Johnson, Michael J. Moss, James M. Zeh, Brian M. Birkmire, and Douglas D. Hodson, "Advanced Framework for Simulation, Integration and Modeling (AFSIM)," *Proceedings of the 2015 International Conference on Scientific Computing*, July 2015.
- "Clowns Break into a Nuclear Missile Launch Facility," KLTN 7, July 1, 2006.
- Correll, Diana Stancy, "Kirtland Airman Charged for Unlawful Importation of Firearm After Investigators Find Trove of Weapons at His On-Base Residence," *Air Force Times*, March 20, 2020.

- Del Valle, Anabel, "Robot Dogs Arrive at Tyndall AFB," 325th Fighter Wing Public Affairs, March 29, 2021.
DoD—See U.S. Department of Defense.
- Douhet, Giulio, *The Command of the Air*, Office of Air Force History, 1983 (originally published in 1921).
- Hale, Kimberly N., *Expanding the Use of Time/Frequency Difference of Arrival Geolocation in the Department of Defense*, RAND Corporation, RGSD-308, 2012. As of October 17, 2023:
https://www.rand.org/pubs/rgs_dissertations/RGSD308.html
- Holmes, Robert H., Bradley D. Spacy, John M. Busch, and Gregory J. Reese, "The Air Force's New Ground War: Ensuring Projection of Air and Space Power Through Expeditionary Security Operations," in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Jackson, Brian A., John C. Baker, Peter Chalk, Kim Cragin, John V. Parachini, and Horacio R. Trujillo, *Aptitude for Destruction: Vol. 2, Case Studies of Organizational Learning in Five Terrorist Groups*, RAND Corporation, MG-332-NIJ, 2005. As of October 17, 2023:
<https://www.rand.org/pubs/monographs/MG332.html>
- Jackson, Brian A., David R. Frelinger, Michael J. Lostumbo, and Robert W. Button, *Evaluating Novel Threats to the Homeland: Unmanned Aerial Vehicles and Cruise Missiles*, RAND Corporation, MG-626-DTRA, 2008. As of October 17, 2023:
<https://www.rand.org/pubs/monographs/MG626.html>
- Johnson, David, "The Triangular Distribution as a Proxy for the Beta Distribution in Risk Analysis," *Journal of the Royal Statistical Society: Series D (The Statistician)*, Vol. 46, No. 3, 1997.
- Jones, Seth G., "Violent Domestic Extremist Groups and the Recruitment of Veterans," statement before the House Committee on Veterans' Affairs, October 13, 2021.
- Knox, Kenneth, "Automated Risk Analysis Tool to Support Transformation of US Air Force Security Forces (Briefing Slides)," June 2009.
- Kotz, Samuel, and J. René Van Dorp, "A Novel Method for Fitting Unimodal Continuous Distributions on a Bounded Domain Utilizing Expert Judgment Estimates," *IIE Transactions*, Vol. 38, No. 5, 2006.
- LaForge, John, "First US Citizen Convicted for Protests at Nuclear Weapons Base in Germany," Counterpunch, May 26, 2020.
- Law, Averill M., and W. David Kelton, *Simulation Modeling and Analysis*, McGraw-Hill, 1982.
- Leemis, Lawrence, "Input Modeling Techniques for Discrete-Event Simulations," in Brett A. Peters, Jeffrey S. Smith, D. J. Medeiros, and Matt W. Rohrer, eds., *Proceedings of the 2001 Winter Simulation Conference*, Institute of Electrical and Electronics Engineers, December 2001.
- Losey, Stephen, "Security Forces Airman Accused of Trying to Burn Cop Car During Protests While Wearing Air Force Gas Mask," *Air Force Times*, August 20, 2020a.
- Losey, Stephen, "Demoted Master Sergeant with White Nationalist Ties Is Now Out of the Air Force," *Air Force Times*, September 9, 2020b.
- McLaughlin, Nancy, "Nuns Proudly Recount Protest Against Missiles," *Greensboro News & Record*, November 11, 2008.
- Milner, Joseph A., *Integrated Defense: Lessons Learned from Joint Base Balad*, Research Paper 2012-3, Air Force Research Institute, 2013.

- National Research Council, *Review of the Department of Homeland Security's Approach to Risk Analysis*, National Academies Press, 2010.
- Nelson, Barry L., *Foundations and Methods of Stochastic Simulation: A First Course*, Springer, 2013.
- Ondaatje, Elizabeth Heneghan, *Trends in Anti-Nuclear Protests in the United States, 1984–1987*, RAND Corporation, R-3696, 1989. As of October 17, 2023:
<https://www.rand.org/pubs/reports/R3696.html>
- Parsons, Doug, John Surdu, and Ben Jordan, “OneSAF: A Next Generation Simulation Modeling the Contemporary Operating Environment,” *Proceedings of the Euro-Simulation Interoperability Workshop*, June 2005.
- Robbert, Albert A., Lisa M. Harrington, Louis T. Mariano, Susan A. Resetar, David Schulker, John S. Crown, Paul Emslie, Sean Mann, and Gary Massey, *Air Force Manpower Determinants: Options for More-Responsive Processes*, RAND Corporation, RR-4420-AF, 2020. As of October 17, 2023:
https://www.rand.org/pubs/research_reports/RR4420.html
- Sagraves, Robert D., “Air Support for Base Defense: Lessons for the Noncontiguous Battlefield,” in Shannon W. Caudill, ed., *Defending Air Bases in an Age of Insurgency*, Air University Press, 2014.
- Sharp, David, “Damaged Sub Miami Leaves Shipyard, to Be Scrapped,” *Navy Times*, June 13, 2015.
- Shlapak, David A., and Alan J. Vick, “Check Six Begins on the Ground”: *Responding to the Evolving Ground Threat to U.S. Air Force Bases*, RAND Corporation, MR-606-AF, 1995. As of October 17, 2023:
https://www.rand.org/pubs/monograph_reports/MR606.html
- Smirnov, N., “Table for Estimating the Goodness of Fit of Empirical Distributions,” *Annals of Mathematical Statistics*, Vol. 19, No. 2, June 1948.
- Judge Advocate General's School, *The Military Commander and the Law*, 15th ed., Air University Press, 2019.
- Under Secretary of Defense for Intelligence, *Internal Review of the Washington Navy Yard Shooting: A Report to the Secretary of Defense*, U.S. Department of Defense, November 20, 2013.
- Under Secretary of Defense for Intelligence and Security, *Physical Security Program*, DoD 5200.08-R, U.S. Department of Defense, April 9, 2007, incorporating change 2, October 19, 2020.
- United States Military Academy West Point, “Combat Simulation Lab,” webpage, undated. As of January 16, 2024:
<https://www.westpoint.edu/academics/academic-departments/systems-engineering/combat-simulation-lab>
- U.S. Department of Defense, *Protecting the Force: Lessons from Fort Hood*, January 2010.
- U.S. Department of Defense Manual S-5210.41, “Nuclear Weapon Security Manual,” U.S. Department of Defense, May 4, 2022.
- Vick, Alan J., *Snakes in the Eagle's Nest: A History of Ground Attacks on Air Bases*, RAND Corporation, MR-553-AF, 1995. As of October 17, 2023:
https://www.rand.org/pubs/monograph_reports/MR553.html
- Vick, Alan J., Sean M. Zeigler, Julia Brackup, and John Speed Meyers, *Air Base Defense: Rethinking Army and Air Force Roles and Functions*, RAND Corporation, RR-4368-AF, 2020. As of October 17, 2023:
https://www.rand.org/pubs/research_reports/RR4368.html

Wilson, Jeremy M., and Alexander Weiss, *A Performance-Based Approach to Police Staffing and Allocation*, Office of Community Oriented Policing Services, 2014.



The U.S. Air Force (USAF) Security Forces (SF) are responsible for detecting, deterring, denying, and defeating security threats in both garrison and expeditionary environments. The SF must address everyday threats that are very similar to those a civilian law enforcement organization responds to on a day-to-day basis. Simultaneously, the SF must address threats that are distinctly military in nature: protecting bases from potentially highly capable attackers. Although existing security planning supports detailed bottom-up, asset-based security planning, such processes do not fully explore the risk trade-offs that are associated with different security strategies, nor do they identify opportunities for SF strategies to manage multiple risks simultaneously.

This report summarizes research on how top-down risk analysis models could help inform USAF decisions regarding SF staffing, systems, and strategies. The goal was not to build a fully fledged risk tool that could be immediately applied to USAF SF personnel and other planning but to take a substantial step in building the foundation for such a tool. Five different scenarios that involved risks to USAF (1) personnel, systems, and facilities and (2) nuclear assets and sites were modeled and explored.

The USAF integrated defense approach has moved from a compliance-based to an effects-based framing for security planning and execution. This research demonstrates how outcome-based risk analysis tools could contribute to that effort, better informing planning and risk tolerance decisions made by commanders at USAF bases at home and abroad.

\$49.95

www.rand.org

ISBN-10 1-9774-1293-9
ISBN-13 978-1-9774-1293-5

