




ABBIE TINGSTAD, DAHLIA ANNE GOLDFELD, LANCE MENTHE, ROBERT A. GUFFEY,
ZACHARY HALDEMAN, KRISTA S. LANGELAND, AMADO CORDOVA,
ELIZABETH M. WAINA, BALYS GINTAUTAS

Assessing the Value of Intelligence Collected by U.S. Air Force Airborne Intelligence, Surveillance, and Reconnaissance Platforms



For more information on this publication, visit www.rand.org/t/RR2742.

About RAND

The RAND Corporation 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/principles.

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

Published by the RAND Corporation, Santa Monica, Calif.

© 2021 RAND Corporation

RAND® is a registered trademark.

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

ISBN: 978-1-9774-0693-4

Cover: Rick Penn-Kraus; photo, aapsky/ Getty Images/iStockphoto.

Limited Print and Electronic Distribution Rights

This document 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. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of its research documents for commercial use. For information on reprint and linking permissions, please visit www.rand.org/pubs/permissions.

Preface

Changes in intelligence, surveillance, and reconnaissance (ISR) and processing, exploitation, and dissemination capabilities over the past two decades have led to ever-increasing demand from warfighters. Commanders, planners, and operators across the U.S. Air Force (USAF) ISR enterprise face difficult decisions about how to best meet ISR needs at the strategic, operational, and tactical levels. Yet USAF currently lacks a consistent, quantitative, empirically grounded method of assessing the value that the service's airborne ISR provides—which is essential to good resourcing decisions.

Headquarters, Air Force, Director of Intelligence, Surveillance and Reconnaissance commissioned RAND Project AIR FORCE (PAF) to examine the value of MQ-1 Predator and MQ-9 Reaper platforms to operations in the U.S. Central Command area of responsibility. After initial investigation, the research focus was adjusted to address the need for an assessment methodology to answer such questions across the USAF ISR community.

This report presents an approach to ISR assessments that seeks to quantify the costs and benefits of USAF airborne ISR in specific operational contexts. Though aspects of this may be applicable across different USAF ISR organizations, this report focuses on the Distributed Common Ground System and the operational theaters it does or could support. The assessment methodology is designed to be flexible enough to support ISR resourcing decisions at different echelons, yet consistent enough to foster feedback, standardize data collections, and make use of empirical analysis methodologies. The report details our approach, provides notional examples to illustrate how it can be applied, and recommends steps USAF can take to establish such an assessment capability. It is intended to complement recent or ongoing research efforts by Air Combat Command, the Intelligence Community, the Joint Staff, and others.

RAND Project AIR FORCE

RAND Project AIR FORCE, a division of the RAND Corporation, is the U.S. Air Force's federally funded research and development center for studies and analyses. PAF provides the Air Force 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; Manpower, Personnel, and Training; and Resource Management. The research reported here was prepared under contract FA7014-16-D-1000.

Additional information about PAF is available on our website:
www.rand.org/paf/

This report documents work originally shared with the U.S. Air Force on September 13, 2017. The draft report, issued on September 30, 2017, was reviewed by formal peer reviewers and U.S. Air Force subject-matter experts.

Contents

Preface	iii
Figures	vii
Tables	viii
Summary.....	ix
Acknowledgments	xviii
Abbreviations	xix
1. Understanding the ISR Assessment Challenge	1
Introduction to the USAF Airborne ISR Enterprise.....	1
Why Assess ISR?	3
How Are ISR Assessments Conducted Today?	4
Project Motivation, Objectives, Scope, and Methods	7
Definition of Terms	12
2. Step 1: Identify the ISR Resourcing and Employment Decisions.....	14
Why Context Shapes USAF Airborne ISR Assessments	14
Types of Questions USAF Airborne ISR Assessments Could Help Answer	15
How Do We Measure ISR's Value in Context?.....	20
Conclusion and Suggested USAF Actions to Improve Assessment Step One	23
3. Step 2: Select Cost and Benefit Indicators	25
How Relevant Was the Information to the Operational Decisionmaker?	25
Was the Decisionmaker Able to Make a Decision with the Available Information?	28
What ISR Resources Were Expended in Those Functions?.....	34
Conclusion and Suggested USAF Actions to Improve Assessment Step 2	35
4. Step 3: Collect Data.....	36
ISR Data Challenges	36
Short-Term Mitigation Approaches	38
Long-Term Mitigation Approaches	41
Recommended Next Steps for USAF.....	51
5. Step 4: Analyze the Value of ISR Using Real-World Data.....	53
Reallocating ISR During Dynamic HVI Operations.....	54
Deciding Whether to Reallocate Resources Between Theaters	64
Conclusion and Suggested USAF Actions to Improve Assessment Step 4a	68
6. Step 4b: Analyze the Value of ISR Using Simulations	71
Approach	72
Vignette	74
Results	77
Conclusions	93
7. Step 5: Apply Results	95

Supporting Decentralized Assessments	95
Communicating Assessment Results	99
Conclusion and Suggested USAF Actions to Improve Assessment Step 5	100
8. Improving the USAF’s ISR Assessments Capability: Conclusions and Recommendations ..	101
Conclusions	101
Recommendations for the USAF to Improve ISR Assessments	102
Appendix A. HVI Simulation Analysis Results	105
Appendix B. Mathematical Model of Fixed Point Security	108
Appendix C. Takeaways from Review of Selected Intelligence Community and Commercial	
Bibliography	120

Figures

Figure S.1. Timeline of Example ISR Decisions Before, During, and After an Operation	ix
Figure S.2. Primary Source Material Used to Synthesize Findings About ISR Assessment Steps	xiv
Figure 1.1. Primary Source Material Used to Synthesize Findings About USAF Airborne ISR Assessment Steps.....	11
Figure 2.1. Timeline of Example ISR Decisions Leading Up to and After an Operation.....	16
Figure 5.1. (U) Notional List of Prioritized Targets.....	55
Figure 5.2. PED Super Crew and Analytical Pool	63
Figure 6.1. Modeling Approach	72
Figure 6.2. SEAS/SCOPEM Targeting Logic	74
Figure 6.3. Compound.....	75
Figure 6.4. Cloud Grid.....	76
Figure 6.5. Run Matrix	78
Figure 6.6. Baseline, Clear Weather, One RPA	79
Figure 6.7. Baseline, Clear Weather, Two RPA.....	80
Figure 6.8. Baseline, Clear Weather, Three RPA.....	81
Figure 6.9. Baseline, Clear Weather.....	82
Figure 6.10. Baseline, Medium Clouds, One RPA.....	83
Figure 6.11. Baseline, Medium Clouds	83
Figure 6.12. Civilian Accounting Summary.....	84
Figure 6.13. HVI Location Summary.....	85
Figure 6.14. Strike Outcome Possibilities	86
Figure 6.15. Strike Outcomes, Baseline, Clear Weather.....	87
Figure 6.16. Strike Outcomes, Baseline, Light Clouds	88
Figure 6.17. Strike Outcomes, Baseline Plus Other Intelligence, Light Clouds	89
Figure 6.18. Strike Outcomes, All Cases	90
Figure 6.19. Nonstrike Outcomes, All Cases	91
Figure 6.20. Missed Opportunities	91
Figure 6.21. Time to Strike.....	92
Figure B.1. Search Area	108
Figure B.2. Confidence of Detection.....	109
Figure B.3. Allocation Across Many Points.....	110

Tables

Table 1.1. Definition of Terms	12
Table 3.1. Example USAF Airborne ISR Roles and Benefit Indicators	30
Table 3.2. PAF USAF Airborne ISR Taxonomy First, Second, and Third Order Roles	32
Table 5.1. ISR Status as of Hour 0	57
Table 5.2. Potential “Fix” of HVI 8	58
Table 5.3. Resourcing and Employment Decision: Reallocate to Enable Fixing of HVI 8?	58
Table 5.4. ISR Status as of Hour 4 (After Reallocation)	59
Table 5.5. Resourcing and Employment Decision: Reallocate to Enable Targeting of HVI 8? ...	60
Table 5.6. ISR Status as of Hour 8 (After Reallocation)	60
Table 5.7. Notional Revisit Rates and Sensors Required for Theaters A and B	67
Table A.1. Civilian Accounting Accuracy in the Steady State—Averages	105
Table A.2. Civilian Accounting Accuracy in the Steady State—Percentiles	105
Table A.3. HVI Location Accuracy in the Steady State—Percentiles	106
Table A.4. Strike Outcomes in the Steady State	106
Table A.5. Nonstrike Outcomes in the Steady State	107

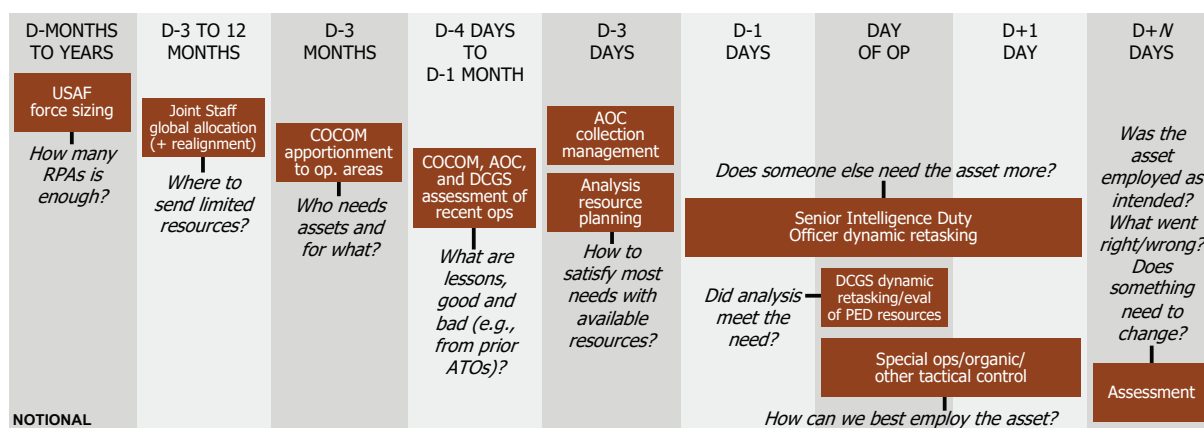
Summary

Assessment Challenge

Growth in the number and capabilities of intelligence, surveillance, and reconnaissance (ISR) sensors, platforms, and processing capabilities over the past two decades has given the United States and its allies a major advantage in recent conflicts. However, it has exacerbated existing challenges to the ability of the U.S. Air Force (USAF) to know, in quantitative terms, just how much value airborne ISR provides. How does a given level of commitment of USAF ISR resources (physical resources and analysts) translate into operational effectiveness (providing appropriate, timely information)?

This issue is especially important today, when a finite supply of USAF airborne ISR resources must meet a large and seemingly insatiable demand for information. Planners, commanders, and operators across the USAF ISR enterprise (and its partners) face an array of resourcing decisions at the strategic, operational, and tactical levels, as illustrated in Figure S.1. In this figure, orange boxes represent resourcing decisions made by different individuals and organizations at different points in time; underneath each box is a question that could be used to assess the resourcing decision. Meeting these resourcing challenges requires a quantitative, empirically grounded method of assessing the value of ISR that is both consistent across echelons and responsive to particular decisionmaking contexts.¹

Figure S.1. Timeline of Example ISR Decisions Before, During, and After an Operation



NOTE: RPA = remotely piloted aircraft; COCOM = combatant command; AOC = Air and Space Operations Center;

¹ Note that, although the study focused largely on USAF airborne ISR, gaining awareness and consistency in assessments with its intelligence and joint partners is important. Broader ISR assessments were not the focus of this study, however.

DCGS = Distributed Common Ground System; PED = processing, exploitation, and dissemination.

Individual teams within the AOCs, DCGS, Air Force Special Operations Command, the National Air and Space Intelligence Center, and other organizations track the effort expended on ISR. In some cases, these teams also note the benefits ISR provides (e.g., new information of interest, actions taken). However, these efforts are not performed consistently across different organizations. Moreover, many USAF and air component organizations are not able to regularly assess airborne ISR; if they do conduct assessments, they focus on quantity (e.g., number of hours of surveillance), rather than on quality or even on mission objectives. This is not because of a lack of motivation or analytic ability on the part of the USAF airborne ISR community. Rather, the USAF (not uniquely) faces two major challenges:

1. The process of planning and conducting USAF airborne ISR operations, along with the databases that support this process, **are not designed for systematic, real-time, or retrospective analysis of how well ISR activities support particular overarching goals.**
2. As a result, there is **no common assessment approach** between (or even within) USAF airborne ISR organizations; **very limited reliable, accurate data** for conducting assessments about USAF airborne ISR; **a lack of common terminology and data standards** for assessments; and, in many cases, **lack of either feedback from end users or access to the contextual information** necessary for USAF airborne ISR analysts to make assessments.

Research Objectives, Scope, and Methods

Awareness of the above challenges prompted USAF leaders to seek better ways to perform assessments across the USAF airborne ISR enterprise. The research discussed in this report was initially motivated by a need within the Office of the Deputy Chief of Staff for Intelligence, Surveillance and Reconnaissance, Headquarters United States Air Force to examine the value of MQ-1 Predator and MQ-9 Reaper platforms to operations in the U.S. Central Command (USCENTCOM) area of responsibility (AOR). However, through a series of initial discussions with experts at Headquarters Air Force, at Headquarters Air Combat Command (ACC), and within the 480th ISR Wing that operates the DCGS, researchers recognized that the airborne ISR assessment challenges faced by USAF are much broader than MQ-1 and MQ-9 operations in USCENTCOM. Thus, the team, with guidance from the sponsoring office, shifted focus to reexamine how USAF can develop more robust foundations for airborne ISR assessment. The analysis is intended to complement other recent or ongoing research efforts by ACC, the Intelligence Community (IC), the Joint Staff, and others.

The research team conducted interviews with stakeholders across the USAF ISR enterprise and at organizations that work with or rely on the enterprise. We reviewed USAF ISR, PED, and assessment procedures and examined prior assessments. To better understand how and why ISR

is employed, we took an approach similar to a strategies-to-tasks framework,² in which broad strategic goals are broken down into increasingly granular activities and tasks required to achieve those goals. We also developed and demonstrated quantitative data analysis methods, using both operational data (i.e., data collected by USAF and other organizations during the course of live operations) and computer simulations. Finally, we explored a limited number of organizations outside USAF that face similar issues with valuing information.

Based on the analysis, we propose a five-step approach to USAF airborne ISR assessment that is flexible enough to be applied to different resourcing decisions, yet consistent enough to foster feedback, standardize data collections, and make use of empirical analysis methodologies. This approach is a framework intended for the USAF airborne ISR community; aspects of it may not apply to other parts of the Joint Community or IC, which may have their own distinct missions and challenges. For example, this approach would require substantial review and filling in with appropriate analytic methodologies before it could be considered for the review of major intelligence failures.

Below, we summarize the five steps and our recommendations for how USAF can improve its capability to carry them out for the purposes of enhancing USAF airborne ISR assessment.

Proposed ISR Assessment Approach

The heart of our proposed approach is a cost-benefit framework that the USAF airborne ISR community can use to begin to quantify the role and value that airborne ISR plays across a range of supported efforts. This framework was formulated on the basis of five analytic assessment steps.

Step 1: Identify the Resourcing Decisions to Be Informed by Assessments

ISR's value cannot be calculated until we answer the question: *value for what?* Thus, the first step in an ISR assessment is to understand the types of resourcing decisions that assessments are meant to inform. As suggested by Figure S.1, the decisionmaking context may vary greatly, from long-term force planning (during which assessments can identify capability gaps and the potential value of more or different assets), to global or regional allocation of existing resources (during which assessments can help deconflict COCOMs' requests or support dynamic reallocation of assets), to tactical management of collection assets and PED resources (during which assessments can help fit the number and type of resources to warfighter information needs and operational objectives). Knowing what resourcing decision will be informed by an ISR assessment guides the development of indicators, data gathering, and interpretation of results.

² Sherrill Lingel, Carl Rhodes, Amado Cordova, Jeff Hagen, Joel Kvitky, and Lance Menthe, *Methodology for Improving the Planning, Execution, and Assessment of Intelligence, Surveillance, and Reconnaissance Operations*, Santa Monica, Calif.: RAND Corporation, TR-459, 2008.

Step 2: Select Cost and Benefit Indicators That Are Relevant to Those Decisions

The assessment framework seeks to quantify and weigh the benefits and costs of ISR operations in a particular operational context. Thus, while the indicators selected for a given assessment should be tailored to the type of decision being made, they will generally fall into *benefit-side indicators* and *cost-side indicators*. These categories may be further broken down into four questions, with the first two pertaining to benefit and the last two to cost:

1. **How relevant was the information to the operational decisionmaker?** Was there a clear link between the sought-for information and the operational decisionmaker's objectives? Did the information help the operational commander make a decision? (Chapter 3 of this report details some indicators to answer these questions.)
2. **Was the decisionmaker able to make a decision with the available information?** The benefits of ISR and PED can be judged by how these capabilities advance good and timely decisionmaking in pursuit of strategic, operational, and tactical goals.³ While it is difficult to characterize what is ultimately a subjective process, we can approximate ISR benefit to decisionmakers by tracking such observable indicators as the number of objects positively identified, the ratio of friendly actions against an adversary to adversary actions against friendlies, and many others (additional details are provided in Chapter 3).
3. **What functions did ISR resources perform?** The cost side of the equation concerns not only *which* ISR resources were used, but *how* they were used. Identifying the functions that ISR performed is crucial to understanding why resources did or did not provide the expected benefits. We propose a taxonomy of ISR roles at five levels of granularity, starting with broad categories (e.g., targeting, support to operations) and ending with specific pieces of information gathered in an individual ISR mission. Such a taxonomy provides insight into what specific ISR functions are intended to achieve and a consistent way to compare different assessments throughout the ISR enterprise.
4. **What ISR resources were expended in those functions?** This category includes information about the number and types of aircraft, sensors, and manpower expended in ISR operations. USAF and others already collect fairly abundant data about resources expended. Although this is sometimes dismissed as “bean counting,” it is very relevant for assessments—as long as it is used with other cost and benefit indicators.

Step 3: Collect Data on These Indicators

Having selected the cost and benefit indicators most relevant to the ISR resourcing decision at hand, the analyst collects the required data. Frequently, however, analysts within the USAF ISR enterprise experience great frustration with accessing and using operational data and contend with limited visibility into mission context and feedback from those requesting information. A consistent, empirically grounded ISR assessment process will be possible only if relevant data are collected and made available to analysts. In the Recommendations later in this section, we

³ We do not mean the ISR resourcing decisions that assessments are meant to inform, but, rather, the operational decisionmaking that commanders and others make on the battlefield, which is enabled by relevant intelligence. Assessments should take into account the possibility for adversary denial and deception.

propose short-term and long-term measures USAF can take to improve collection and dissemination of data needed for ISR assessments.

Step 4: Analyze the Value of ISR Using Historical Data, Simulations, or Both

The next step is to perform the actual analysis. Just as the indicators and datasets will vary depending on the decisionmaking context, analytical methods will vary for each given assessment. While we might think of ISR assessments as occurring after completed operations, assessments based on real-world data are not always possible or desirable because of the nature of intelligence problems, the operational environment, or the technical maturity of new capabilities. Under these circumstances, simulations can help. We may also encounter situations where both real-world and simulated analyses are desired.

In Chapters 5 and 6 of this report, we illustrate just a few of the types of assessments that could be carried out to support tactical to strategic ISR resourcing decisions. For example, we show how ISR analysts supporting a high-value individual (HVI) strike mission can use operational data to inform the dynamic reallocation of ISR and PED resources to take advantage of unexpected strike opportunities while minimizing the cost to parallel HVI missions. We also show how assessment of historical usage data can inform the reallocation of finite ISR resources among COCOMs. In another illustration, we show how computer simulation using the RAND Systems and Concept of Operations (CONOPs) Operational Effectiveness Model can be used to assess the marginal value of adding ISR collection assets to support HVI strike missions under different conditions. While specific analytic methods are different for different types of resourcing decisions, all are grounded in the cost-benefit approach outlined above, and all demonstrate the importance of having consistent, clearly defined indicators and accurate, readily available data.

Step 5: Apply Results Back to the Decisionmaking Process

If the previous steps have been conducted effectively, the assessment product will be relevant and accurate, and it should not be difficult to apply the resulting information to decisionmaking about ISR resourcing. However, ISR assessments do not always address the most relevant (or correct) questions, provide implications for upcoming resourcing and employment decisions, or reach all audiences who would benefit from the information. In Chapter 7, we discuss some ways to address these issues.

The primary source material we used to develop findings and recommendations related to each assessment step is summarized in Figure S.2.

Figure S.2. Primary Source Material Used to Synthesize Findings About ISR Assessment Steps

<i>Assessment step</i>	<i>Primary source material for analysis</i>
1. Identify the decision to be informed	Joint publications and Air Force instructions/doctrine documents; expert interviews within the USAF ISR Enterprise and Joint Community (Joint Staff, CENTCOM, PACOM)
2. Select relevant cost and benefit indicators	Expert interviews with analysts, primarily those working within and supporting missions in CENTCOM and PACOM; review of selected IC, “smart market,” ride sharing, online search algorithm, scientific journal index, and healthcare industry practices
3. Collect data on these indicators	Expert interviews with analysts, database maintainers, and programmers; documented project team lessons from using existing databases; review of selected IC and healthcare industry practices
4. Analyze the value of ISR	New PAF cost-benefit framework; data insights derived from analysis of step 3; review of IC practices; expert interviews with analysts, primarily those working within and supporting missions in CENTCOM and PACOM; computer simulation; previous RAND research
5. Apply results to decisionmaking process	Expert interviews and site visits within the USAF ISR Enterprise and Joint Community (Joint Staff, CENTCOM, PACOM); review of selected IC, smart markets, ride sharing, online search algorithm, scientific journal index, and healthcare industry practices

NOTE: PACOM = U.S. Pacific Command.

Recommendations to Improve Airborne ISR Assessments

Not all of the changes required to improve airborne ISR assessments are within USAF’s power to implement. However, USAF can do much to improve airborne ISR assessments. The recommendations below are steps USAF can take in the areas of oversight; guidance; and tactics, techniques, and procedures to move toward the kind of airborne ISR assessment process described in this report.

Guide a Consistent Approach to ISR Assessment

Guidance should cover what constitutes an assessment, the basic steps for conducting assessments (such as the five-step approach outlined above), and key questions about benefits

and costs that should be answered as part of any assessment.⁴ Guidance should also establish or reaffirm the expectation that organizations will use assessments to support their own decisionmaking about ISR (e.g., within an air tasking order [ATO]) as well as joint or IC decisionmaking under certain circumstances (e.g., an AOC feeding assessments to a COCOM for planning purposes). It might be useful to disseminate guidance in a short format, such as an information graphic, that can be widely distributed and displayed.

Use a Common Lexicon for Assessments, Requirements Articulation, Collection Management, and PED

The lexicon should include a granular taxonomy of ISR roles (such as the one described here) and clearly defined terms such as *benefit*, *cost*, and *indicator*. This lexicon should be coordinated with the joint community, in some cases leveraging or providing a cross-walk to existing terms and definitions.⁵

Plan to Improve Fidelity and Discoverability of ISR Data

- **In the short-term:**
 - Introduce a method for recording ISR roles in requirements databases as well as in the ATO, Unified Collection Operations Reporting Network (UNICORN),⁶ and other databases. This will enable better assessment of why sorties flew and how they were intended to support commanders' objectives.
 - Automatically scrub UNICORN data to alert users to problems (e.g., logical inconsistencies, misspellings).
 - Add useful fields to UNICORN, such as target prioritization and ISR asset tail number.
- **In the long-term:**
 - Implement automatic machine-to-machine transmission of data from aircraft and sensors to data repositories.
 - Tag collection to produce searchable metadata.
 - Store historical data and make it easily retrievable.
 - Create a new, indexed database for ISR assessments that enables data discovery and analysis.
 - Include an ability to track intelligence citations.

⁴ It would be beneficial for this guidance to be informed by and coordinated with assessment organizations within the joint community and IC, including the Joint Staff and COCOM staffs.

⁵ Note that some of the terms; definitions; and tactics, techniques, and procedures needed for ISR assessments are not available in joint or USAF doctrine. Developing these in coordination with the broader joint community and IC will be very important so that more than one lexicon is not developed.

⁶ This database contains ISR sortie, collection, and PED information for missions analyzed within the 480th ISR Wing.

Enhance Airmen Skills and Resources to Perform Assessments

A more analytically rigorous and sophisticated assessment capability will require analysts with comparable skills who can regularly execute a variety of analyses and assessments needed to ascertain ISR's value and drive future decisionmaking. USAF should consider several options to bolster the analytic force, such as:

- emphasizing data science and computer programming in the recruitment and development of ISR airmen
- examining whether the Reserve component can be used to enhance ISR assessments capabilities
- taking advantage of USAF computer coding and data science expertise by encouraging local commanders to select computer- and data-science-savvy airmen for temporary task reassignments
- finding ways to leverage USAF scientists within the total force to work temporarily within the ISR community (e.g., through a program that organizes opportunities for visits and temporary assignments)
- exploring the use of more contractors for short-term outsourcing of coding or database development
- hiring temporary consultants from technology companies to advise on how to best proceed with technological innovation plans.⁷

Coordinate with IC, Joint, and Industry Partners to Identify Opportunities for Gathering Feedback

Two types of feedback to institute include active (e.g., an automated window asking users to select a score) and passive (e.g., statistics on how often products are accessed) approaches.

Further Analysis to Determine the Advisability of Changing PED Force Presentation and the Impacts on Assessments

Shifting PED analyst operations to more a “problem-centric” approach might enable these analysts to better communicate their findings in the context of progress towards solving problems, which would further enable assessments to draw upon the results of PED. This recommendation is similar to that made in a previous RAND report,⁸ but it is also supported by new evidence revealed during the course of this research effort.

⁷ The service should weigh the pros and cons of these (and perhaps other) options as well before determining any specific course of action.

⁸ Lance Menthe, Amado Cordova, Carl Rhodes, Rachel Costello, and Jeffrey Sullivan, *The Future of Air Force Motion Imagery Exploitation: Lessons from the Commercial World*, Santa Monica, Calif.: RAND Corporation, TR-1133, 2012.

Identify When Simulations Are Needed to Support Assessments

Simulations may be especially useful at developing tactics and procedures for informing ISR employment in different contexts, situations where “ground truth” in the real-world is unreliable (either because too little data has been collected or there are many confounding factors), and cases where risk must be understood in mathematical detail, such as exploring options for minimizing collateral damage.

Continually Refine and Update Processes and Guidance for ISR Assessments

USAF must host regular, periodic discussion forums to refine guidance as necessary, update the ISR roles taxonomy as needed, and identify the latest assessments challenges and possible areas of “low-hanging fruit” to target. **This last recommendation is needed to ensure that the previous eight recommendations move forward.** These forums could also be used to communicate and update a list of ISR assessment priorities for the Air Staff related to emerging issues (e.g., ISR in new operating environments, pending investment or divestment decisions) and enduring issues (e.g., ISR benefits in support of long-term national strategic priorities) so that communities within the USAF enterprise can identify existing assessments and datasets within their purview that are relevant and pass them to an appropriate point of contact.

Although some of the recommendations may pose budgetary and organizational challenges or require cultural shifts within USAF, the potential benefits of having better-managed ISR resources are surely worth the effort.

Acknowledgments

The research team would like to thank the sponsor, Brig Gen Aaron Prupas, and the project action officers—Lt Col Brian McCreary, Lt Col Bradley Readnour, and Lt Col John Walsh—in the U.S. Air Force (USAF) Office of the Deputy Chief of Staff for Intelligence, Surveillance, and Reconnaissance. There were also a number of people within the USAF and Joint Staff communities that provided valuable inputs to this project including from organizations such as: Office of the Under Secretary of Defense for Intelligence; U.S. Joint Staff; Headquarters, Air Force; Headquarters, Air Combat Command; Headquarters, U.S. Central Command; Headquarters, U.S. Pacific Command; U.S. Air Forces Central Command; U.S. Pacific Air Forces; Air Force Special Operations Command; 25th Air Force; and 480th Intelligence, Surveillance, and Reconnaissance Wing. Although many individuals contributed content to this report, the following repeatedly offered their time in order to provide information, give feedback, and facilitate site visits: Col Jason Brown, Maj Joel Nieman, Maj Daniel Pagano, Capt Brandon Poje, Capt Ezra Caplan, Capt Alexander Wilkie, JudyAnn Wehking, Dean Milne, and Greg Sanders.

We would also like to thank our RAND colleagues that assisted in reviewing this report: Natalie Crawford, Sherrill Lingel, Daniel Gonzales, and Michael Decker. The research team also benefitted from discussions with Project Air Force Fellow Col Donald Brunk. Pardee RAND Graduate Student Fellows Eugene Han, Jesse Lastunen, David Mannheim, and Hilary Smith assisted with writing up assessment practices in other industries and organizations. We also appreciate the assistance of Karin Suede in document preparation.

Abbreviations

ACC	Air Combat Command
AFI	Air Force Instruction
AFSOC	Air Force Special Operations Command
AOC	Air and Space Operations Center
AOR	area of responsibility
ATO	Air Tasking Order
COCOM	combatant command
COD	Combat Operations Division
DART	DCGS Analysis and Reporting Team
DCGS	Distributed Common Ground System
DoD	U.S. Department of Defense
EEI	essential elements of information
F2T2EA	find, fix, target, track, engage, and assess
FMV	full motion video
FOV	field of view
HQ AF	Headquarters, Air Force
HQ AF/A2	Headquarters, Air Force, Director of Intelligence, Surveillance and Reconnaissance
HUMINT	human intelligence
HVI	high-value individual
IA	imagery analysts
IDEA	ISR Data Enrichment and Aggregation
IC	intelligence community
IPoE	intelligence preparation of the environment
ISR	intelligence, surveillance, and reconnaissance
ISRD	Intelligence, Surveillance and Reconnaissance Division

I&W	indications and warnings
JFACC	Joint Forces Air Component Commander
JFC	Joint Forces Command
JS	United States Joint Staff
NASIC	National Air and Space Intelligence Center
NGA	National Geospatial-Intelligence Agency
NRT	near-real-time
OBP	object-based production
PAF	RAND Project Air Force
PCPAD	planning and direction, collection, processing and exploitation, analysis and production, and dissemination
PED	processing, exploitation, and dissemination
PRISM	Planning Tool for Resource Integration, Synchronization, and Management
ROE	rules of engagement
RPA	remotely piloted aircraft
SCOPEM	Systems and Concepts of Operations Operational Effectiveness Model
SD	Strategy Division
SEAS	System Effectiveness Analysis Simulation
SIGINT	signals intelligence
SME	subject-matter expert
UNICORN	Unified Collection Operations Reporting Network
USAF	U.S. Air Force
USCENTCOM	U.S. Central Command
USPACOM	U.S. Pacific Command

1. Understanding the ISR Assessment Challenge

The U.S. Air Force (USAF) is in the midst of an information renaissance. The proliferation of airborne sensors over the past two decades, driven by increasing numbers of remotely piloted platforms available to carry them, has greatly expanded the volume of information that can be collected to support warfighters. The MQ-1 Predator and MQ-9 Reaper's ability to carry weapons has further integrated these assets into strike operations. USAF airborne still and motion imagery and signals intelligence (SIGINT) have played a major role in recent conflicts, during which the ability to know what adversaries are doing in real time and in fine detail has often meant the difference between mission success and failure—and between life and death for enemies, friendly forces, and civilians.

These developments have raised questions about the ability of the USAF to know, in quantitative terms, how its commitment of airborne intelligence, surveillance, and reconnaissance (ISR) resources (including physical resources and how they are employed) translates into operational effectiveness (providing appropriate and timely information to enable good decisionmaking). Challenges in assessing ISR's value have persisted over decades and throughout different conflicts. How can a finite supply of ISR resources be best employed to meet a large and growing demand for information? "More is better" is not a useful guide for the difficult resourcing and employment decisions that must be made at tactical, operational, and strategic levels. Yet USAF has so far lacked a quantitative, empirically grounded method of assessing airborne ISR value that is consistent across echelons and responsive to the challenges that decisionmakers face about investing in and allocating ISR resources. This problem is expected to continue to grow as the number of USAF sensor-bearing platforms increases, demand grows and diversifies, and adversaries potentially use increasingly sophisticated denial and deception.

This report attempts to fill the USAF airborne ISR assessment gap by proposing a methodology that USAF can use to inform its airborne ISR resourcing and employment decisions at tactical, operational, and strategic levels. It discusses the challenges inherent in the task, reviews current USAF methods of addressing those challenges, and outlines a roadmap for near-term and long-term improvements. First, we briefly describe some relevant aspects of USAF ISR.

Introduction to the USAF Airborne ISR Enterprise

USAF conducts airborne ISR operations in support of a diverse variety of strategic, operational, and tactical objectives and decisionmaking worldwide. Collecting data and conducting analysis helps USAF contribute to such high-level goals as personnel safety,

increasing knowledge of adversary intentions and capabilities, supporting diplomatic activities, maintaining a plausible deterrent, conducting offensive action with precision, and aiding operational efficiency.¹ USAF airborne platforms used to collect ISR include manned assets, such as the U-2 Dragon Lady and the RC-135 Rivet Joint, and remotely piloted aircraft (RPA), such as the MQ-9 Reaper and RQ-4 Global Hawk. Pilots and crews undergo extensive training and perform their missions under a variety of challenging circumstances.

Planning and supporting ISR operations requires ISR specialists within the AOCs and analysts conducting processing, exploitation, and dissemination (PED) within the Distributed Common Ground System (DCGS) located in different theaters of operation around the world. The primary vehicle (a USAF unit and a weapon system) through which the Joint Forces Air Component Commander (JFACC) fulfills their responsibilities is the Air and Space Operations Center (AOC), which “provides the capability to plan, coordinate, allocate, task, execute, monitor, and assess the activities of assigned or attached forces.”²

PED is critical for providing information for decisionmaking in the battlefield. *Processing* refers to the (typically automated) conversion of collections into formats that can be digested and analyzed by USAF intelligence specialists. *Exploitation* consists of analysts’ actions to extract useful information from the collections. *Dissemination* is the process of transmitting the information to the users. The 480th, 70th, and 55th Wings are in charge of performing much of the initial PED of the collections from the USAF ISR assets used in theater. Most of the time, USAF analysts focus on performing near-real-time (NRT) and short-term exploitation; other organizations within USAF—notably the National Air and Space Intelligence Center (NASIC)—and the intelligence agencies are in charge of longer-term (such as forensic) analyses.

The USAF airborne ISR community contains thousands of people and leverages billions of dollars in hardware, software, and facilities. It is connected with the broader intelligence community (IC) and the warfighting combatant commands (COCOMs) and relies upon communications, logistics, aircraft and system maintenance, and a host of other supporting functions. USAF airborne ISR delivers intelligence to a variety of stakeholders, including those supporting strategic government decisionmaking, those conducting warfighting operations, and those planning for potential contingencies. These users are distributed throughout USAF, the IC, COCOMs, and broader U.S. military and government.

The term *ISR* can encompass many elements, including both the results of ISR activities (e.g., information) and the processes of planning and direction, collection, processing and exploitation, analysis and production, and dissemination (PCPAD). In this report, we use *ISR* to

¹ These examples of goals arose out of collective knowledge gained through project interviews and are consistent with recent U.S. strategic visions and USAF doctrine, although they are not directly derived from these references.

² Air Force Instruction (AFI) 13-1AOC, *Operational Procedures—Air Operations Center (AOC)*, Vol. 3, Washington, D.C., May 18, 2012; Joint Publication 3-30, *Command and Control of Joint Air Operations*, Washington, D.C., February 10, 2014, p. xi.

refer to both results and processes, except where we refer to specific elements as such. We primarily refer to ISR in the context of USAF airborne operations.

Why Assess ISR?

The importance, size, and dollar value of the USAF airborne ISR enterprise—which, for the purposes of this report, includes the people, hardware, software, and facilities used to plan and conduct airborne ISR operations, as well as to analyze and deliver intelligence—drives the need to understand the value of USAF airborne ISR employment. Otherwise, managing investment in USAF airborne ISR assets and related system and resource allocation is based on use projections that are not substantiated by analysis of employment benefits. This is a key concern for stakeholders, such as Congress and the U.S. Department of Defense (DoD), which must navigate competition for USAF airborne ISR resource use (among other forms of ISR) and the uncertainty of the future within budgetary constraints. The difficulty inherent in this task of understanding ISR's value was a driving force behind commissioning various research efforts, including our own, to look at assessing the value of ISR. Naturally, this challenge is not generic to the USAF airborne ISR community; we did examine some relevant insights from other organizations, but we did not fully explore their challenges and assessment approaches, nor do we necessarily suggest that any findings or recommendations in this report would be helpful to others. Other parts of the joint community and IC may conduct some missions and face particular challenges distinct from the USAF airborne ISR community.

Another driving force behind this research is what appears to be an ever-growing, insatiable appetite for USAF airborne ISR support (among other forms of ISR, including within the IC and joint communities) among users conducting tactical operations, planning for possible future campaigns, and making strategic and operational decisions. The most recent trend of increased USAF airborne ISR requests appears to have started around outset of Operation Enduring Freedom in Afghanistan (2001–2014) and Operation Iraqi Freedom (2003–2011) and is likely correlated with the growth in the use of USAF RPAs for ISR and the particular characteristics of counterinsurgency and counterterrorism warfare. Regardless of the precise root causes, organizations providing, directing, and analyzing USAF airborne ISR have noticed a tendency in which units receiving intelligence support request more regardless of whether prior missions closed knowledge gaps (i.e., saying “ISR was great, we want more” or “ISR was terrible, we need more.”) A consistently applied ISR assessment process is needed by the USAF airborne ISR community to understand and validate this seemingly endless growth in demand and make decisions about how to best allocate and employ resources across different missions and timescales.³

³ Whether it is useful to employ an ISR assessment process that is also consistent with that used within other organizations outside the USAF airborne ISR community was not addressed by this research. Some joint partners

Some generic examples of important USAF airborne ISR assessment questions include:

- What value did ISR bring to a particular operation?
- Could resources (e.g., aircraft, personnel) be employed more efficiently and more effectively, and how can these goals be achieved?
- What is the return on the USAF investment in analysts and collection platforms?
- What is the effect of rebalancing resources toward emerging problems (and therefore taking resources away from existing obligations)?

To answer these types of questions, the value of USAF airborne ISR in a particular context must be understood. There is no “marketplace” for ISR in which value to users can be determined based on supply and demand dynamics.⁴ Therefore, determining ISR’s value requires a detailed understanding of the benefits and relative costs of using it to support making progress towards various strategic, operational, and tactical goals. Ideally, this would be accomplished using a USAF ISR assessments process that

- provides feedback that helps both affect change within an individual organization or team and justify resource requests and other decisions in discussions with leadership and other, external organizations
- links ISR activities to achievement of overarching strategic, operational, and tactical decisionmaking
- is able to examine whether resources were used as intended and if capabilities and personnel were employed in a beneficial manner
- proves helpful for understanding the value of a single mission or activity with respect to overarching goals, in addition to long-term patterns
- relies on well-understood approaches, a common lexicon, enforced data standards, and feedback
- is achievable by anyone in the USAF airborne ISR community
- uses well-maintained databases that incorporate information needed for assessments to support ISR resourcing and employment decisions across a range of time frames and missions
- aids analysts in advancing the intelligence narrative, even in the presence of adversary denial and deception activities
- helps support decisionmakers makes the right call—even if that is to do nothing.

How Are ISR Assessments Conducted Today?

However, no such ISR assessments process currently exists within USAF, at least not with all of these attributes and in a form that can be distributed across the USAF ISR community. We argue that having some commonalities in assessment and data guidelines is valuable to facilitate

(e.g., at a COCOM) did express in staff interviews that having a common lexicon with the USAF airborne ISR community to facilitate communication about ISR’s value in different contexts could be useful.

⁴ There are priorities that are set for ISR employment, but we argue that these do not necessarily reflect value or enable determination of return on investment.

communication and enable (when necessary) comparison within the USAF ISR community. Establishing commonalities in assessment procedures with IC or joint partners could also be valuable in some contexts (not least to facilitate the exchange of ideas and helpful assessment practices), but we did not conduct an extensive analysis of practices within organizations outside of USAF.

Assessment teams within USAF organizations, such as the AOCs, DCGS, Air Force Special Operations Command (AFSOC), NASIC, and others, track measures that quantify effort expended on ISR. In some cases, they gather information that suggest benefits for informing decisions (e.g., if new information about something of interest is found, when an insight directly enables action). Although these efforts are valuable, many organizations are not able to conduct regular assessments or, if they do, the assessments largely focus on the quantity of ISR effort—i.e., the amount of resources expended. If any “benefits” (e.g., dangers averted, significant discoveries made) are apparent, they are typically documented in PowerPoint presentations. These presentations are succinct and portable, but they are not suited to understanding the complexities of longer-term resource allocation. In addition, these presentations obscure the typically incremental progress toward discovering new or otherwise significant information as well as ISR contributions to monitoring situations, in which a *lack* of activity is an important indicator for supporting critical decisions, and situations in which ISR contributes to a correct “no-go” decision (e.g., with respect to a strike).⁵

As a result, USAF airborne ISR contributions to any particular goal cannot currently be identified across the enterprise or, in many cases, even within a particular organization or team. This is not due to a lack of motivation or analytic ability on the part of the USAF airborne ISR community. Rather, there are two major challenges that the USAF airborne ISR community (perhaps not uniquely) faces:

1. The process of planning and conducting USAF airborne ISR operations, along with the databases that support it, **are not designed for systematic, real-time or retrospective analysis of how USAF airborne ISR activities support particular overarching goals of information users.**
2. As a result, there is **no common assessment approach** between (or even within) USAF airborne ISR organizations; **very limited availability of reliable, accurate data**; a **lack of common terminology and data standards**; and, in many cases, **lack of either feedback from end users or access to contextual information** needed for ISR specialists to make assessments.

Planning and conducting USAF airborne ISR operations around the world (but in large part supporting ongoing missions in and around the U.S. Central Command [USCENTCOM] area of operations [AOR]) requires tremendous focus on the upcoming 72 hours or so, in line with Air Tasking Order (ATO) planning and execution. It has proven difficult for the USAF airborne ISR

⁵ It may also not often be known when a “no-go” decision was truly correct, but this is something that computer simulations can assist in examining, as we describe later in this report.

community to consider needs for retrospective analysis given the heavy burden of orchestrating and conducting initial analysis of the results from real-time and very near-term airborne ISR missions.

Awareness of these challenges and of the importance of enabling longer-term airborne ISR assessments prompted USAF leaders to seek better ways to perform such assessments across the ISR enterprise. Doctrinally, some assessments of ISR are performed at the AOC in response to the planning, production, and execution of the ATO, the airspace control order (ACO), the joint integrated prioritized target list, and the ISR synchronization matrix. These assessments are supposed to be conducted across the AOC's five divisions (the ISR Division [ISRD], the Strategy Division [SD], the Combat Plans Division, the Combat Operations Division [COD], and the Air Mobility Divisions). The importance of data transfer between these divisions is recognized and doctrinally accomplished through embedding personnel across the divisions. The goal of assessments is to determine how well JFACC/Joint Forces Command (JFC) priorities were fulfilled and track the basic employment of ISR.

The ISRD is responsible for "assessing and anticipating adversary activity in the operational environment, managing ISR operations and developing dynamic targeting strategies in order to rapidly, discreetly and efficiently achieve the JFACC/JFC priority effects."⁶ ISR personnel are embedded throughout the rest of the AOC, providing critical links between tactical, operational, and strategic battlefield outcome data with ISR data. ISRD Operations team support personnel who are specially embedded in the SD and are *explicitly* tasked to "conduct ISR operations assessment."⁷ ISRD operations support personnel embedded in the COD are directed "provide feedback on ISR operations accomplished each ATO day in support of ISR operations assessment."⁸ This should, in theory, give the ISR assessment teams accurate battlefield outcome data which is essential to tracking how operational goals (JFACC/JFC priorities down to specific missions) are progressing.

There is also a PED Management Team (some of whom are embedded in the COD) who are tasked with tracking collection requirements throughout the PED process, ensuring that it is responsive to customer timelines. They are directed to coordinate with ISR personnel enterprisewide as well as JFC/component collection managers to assess accomplishments of collection requirements and provide ISR operations assessment data during the JCMB. Lastly, they are also supposed to provide metrics for ISR assessment to the SD.

The Operational Assessment Team (OAT) within the SD is tasked with analyzing completed operations to assess compliance with strategy, and to link intelligence data with operational courses of action. The OAT is supposed to use "a foundation of data and the fusion of that data" to identify and evaluate the effectiveness of friendly operations, highlight opportunities to

⁶ AFI 13-1, 2012, p. 15.

⁷ AFI 13-1, 2012, p. 64.

⁸ AFI 13-1, 2012, p. 69.

influence adversary courses of action, and evaluating the effects achieved on the adversary's strategic and operational centers of gravity.⁹ These analyses are supposed to provide the JFACC with the information needed to make decisions regarding air, space, and cyberspace operations *strategy*. While not ISR-specific (if anything, this is a far more daunting assessment directive), certainly ISR assessments should play a role in the OAT's work.

The AOC is not the only type of USAF organization that performs ISR assessments. Among others, the 25th Air Force, NASIC, and analysts within the DCGS perform assessments. In addition, organizations such as Headquarters, Air Force (HQ AF) and Air Combat Command (ACC) have engaged in ISR assessment activities. These analyses can, to some extent, describe what ISR assets are doing and associate finished intelligence products with ISR sorties. There is less emphasis (because of data access and other challenges) on analyzing the effectiveness of ISR employment or tracking ISR's relative contributions to operations. At the heart of many of these USAF airborne ISR assessment analyses is the UNICORN database, as well as the ISR Assessment Tool and the Mission-Summary Analysis Tool. Although these are valuable repositories of information, we describe the data-related difficulties that current USAF airborne ISR assessment efforts face (as well as the tools that USAF assessors currently have at their disposal) in Chapter 4. Note that other organizations that conduct ISR assessments that may use a broader range of data and tools, and thus may not face the same challenges the USAF does. We also describe enhancements that the USAF can make to data and databases that should help to overhaul ISR assessment capabilities.

Project Motivation, Objectives, Scope, and Methods

The research discussed in this report was initially motivated by a need within the Office of the Deputy Chief of Staff for HQ AF, Director of Intelligence, Surveillance and Reconnaissance (HQ AF/A2) to examine the value of USAF-owned and -operated MQ-1 Predator and MQ-9 Reaper platforms to missions in the USCENTCOM AOR. This was, in part, related to ongoing discussions within the USAF, and among the Joint Staff and Congress members, regarding how many RPAs are needed for operations. At the time the research was being formulated, it was also recognized that this topic was representative of a broader need to assess ISR's value in order to help the USAF make future decisions about airborne ISR force-sizing, resource allocation, and force employment.

This topic has also recently been examined by others in the USAF ISR enterprise, the Joint Community, and the IC. For example, the Joint Staff is working on ways of establishing measures of effectiveness on the basis of relating operational effects, ISR effects, and ISR tasks—a concept that served as one basis for some of the attributes of defining ISR value

⁹ AFI 13-1, 2012, p. 28.

discussed later in this report.¹⁰ The staff within the office of the Director of Intelligence, ACC (ACC A2) is making progress on developing overarching guidance for ISR assessments, that includes steps for assessment, linking a commander's objectives to ISR objectives, and defining needs for assessments data and databases.¹¹ Caplan writes about comparing anticipated and actual operational outcomes and linking results with ISR employment.¹² The research presented in this report complements the ideas presented in other efforts, and in some cases suggests new concepts or provides more granular detail on specific approaches. However, this research is only intended to provide insight with regard to the USAF airborne ISR assessments and does not attempt to rigorously examine or provide recommendations about ISR assessments conducted within the Joint Community or IC. It also does not fully consider the implications of adversary deception and denial activities.

The challenges associated with USAF airborne ISR assessments are not new. For example, Lingel et al. examined ISR assessments in the context of planning for AOC operations in the Pacific theater. This previous study was one of the first efforts to explicitly define a methodology for linking top-level commander's guidance, operational objectives, and operational tasks to particular information needs and ISR tasks using a strategy-to-tasks framework. Lingel et al. observed several issues with ISR assessment that still appear relevant to the USAF airborne ISR community:¹³

- The Planning Tool for Resource Integration, Synchronization, and Management (PRISM), one of the main collection-management tools used in theater, does allow operators to associate collections with Priority Intelligence Requirements (PIRs), but a detailed understanding of the role of the collection in satisfying the PIR is not included in PRISM.
- Few written links (if any) exist between a commander's priorities and individual collections. Collections are only prioritized by their ranking in a list (e.g., the Joint Integrated Prioritized Collection List); thus, there is insufficient information to make informed trade-offs between collections, causing difficulties in making a decision between pursuing a deliberate collection or performing dynamic ISR asset retasking for an emerging target.
- Making intelligent decisions about retasking ISR assets is also somewhat precluded by the difficulty in unraveling what is lost at the strategic level by not satisfying a particular preplanned collection requirement.

¹⁰ Joint Functional Component Command for Intelligence, Surveillance and Reconnaissance Assessments Division, "Measures of Effectiveness," briefing slides, April 20, 2016.

¹¹ Office of the Director of National Intelligence, ACC, "ISR Assessment Framework," white paper, September 27, 2016.

¹² Ezra Caplan, *Recommendations for Mission-Centric Airborne ISR Employment and Assessment*, Maxwell Air Force Base, Ala.: LeMay Center for Doctrine, March 2017.

¹³ Sherrill Lingel, Carl Rhodes, Amado Cordova, Jeff Hagen, Joel Kvitky, and Lance Menthe, *Methodology for Improving the Planning, Execution, and Assessment of Intelligence, Surveillance, and Reconnaissance Operations*, Santa Monica, Calif.: RAND Corporation, TR-459, 2008.

Lingel et al. proposed using a strategy-to-tasks framework—which was already in use in the Strategy and Combat Plans divisions of the AOC—to better understand the contribution of individual ISR collections vis-à-vis the campaign objectives and commander’s guidance and to help guide the prioritization of the collections, as well as any necessary retasking of ISR assets.¹⁴ This framework was expected to make more transparent throughout the chain of command why certain collections were being pursued when others were not and to help determine when to replace a planned collection with an ad-hoc collection for an emerging target. Moreover, the methodology aimed to help with separating the importance of a particular collection requirement from the likelihood of successfully collecting against that requirement. The development of a web-based application containing an up-to-date strategy-to-tasks framework, as well as relevant essential elements of information (EEIs) and outcomes associated with each ISR task was proposed as one way to operationalize this assessments framework. As discussed in subsequent chapters, we adapt a similar approach to measure how USAF airborne ISR activities advance operational objectives.

The research documented in this report follows from the work by Lingel et al. and is intended to offer a complementary perspective to similar recent or ongoing efforts within USAF. It was originally formulated to answer the following assessment questions:

- What is the value of USAF MQ-1 and MQ-9 to operations writ large in the USCENTCOM AOR?
- How many of these RPAs are enough for the USAF to support these operations?

However, through a series of initial discussions with experts at HQ AF, at Headquarters Air Combat Command (HQ ACC), and within the 480th ISR Wing that operates the DCGS, the research team recognized that the airborne ISR assessment challenges faced by the USAF are much broader than MQ-1 and MQ-9 operations in USCENTCOM. Any proposed solutions must be formulated from a broader perspective in order to avoid some of the problems associated with separate assessment processes that are discussed in more detail later in the report.¹⁵

Therefore, the research team—with guidance from the sponsoring office—shifted focus somewhat to re-examine how the USAF can develop more robust foundations for airborne ISR assessment, by addressing the following research questions:

1. Why does USAF need to conduct airborne ISR assessments, and what should this process broadly look like?
2. What is a framework for understanding and tracking USAF airborne ISR’s value?

¹⁴ These divisions used the framework to create target nominations, starting with commander’s guidance and ending with a prioritized target list. Moreover, rather than having two frameworks with different sets of objectives and tasks, the team suggested that the ISR Division and the other AOC divisions coordinate their efforts and use a single and unified framework for ISR allocations in support of overall campaign planning.

¹⁵ Although each organization with a need for ISR assessments might require different specific data sources and approaches, the current system provides so little guidance and standardization that it inhibits effective development, communication, and application of assessment results.

3. How do operational and environmental factors impact the number of distinct USAF airborne full-motion video (FMV) and SIGINT collectors needed to achieve a particular level of confidence in intelligence?
4. What steps should the USAF take to improve the service's ability to conduct airborne ISR assessments?

At the heart of our proposed approach to ISR assessments is a framework that the USAF airborne ISR community can use to begin to quantify the role and value that airborne ISR plays across a range of supported efforts. Importantly, this approach is a framework intended for the USAF airborne ISR community and aspects of it may not apply to other parts of the Joint community or IC. For example, this approach would require substantial review and filling in with appropriate analytic methodologies before it could be potentially considered for the review of major intelligence failures. This framework was formulated on the basis of five basic analytic research steps, which are discussed at length in subsequent chapters, and apply in this research specifically to USAF airborne ISR:

1. **Identify the ISR resourcing and employment decisions** to be informed by ISR assessments (Chapter 2).
2. **Select cost and benefit indicators** that are relevant to those decisions (Chapter 3).
3. **Collect data** on these indicators (Chapter 4).
4. **Analyze the value of ISR** using historical data (Step 4a; Chapter 5) and/or simulations (Step 4b; Chapter 6).
5. **Apply results** back to the decisionmaking process (Chapter 7).

This approach was developed based on several lines of research. We conducted a series of interviews with a sample of stakeholders across the USAF ISR enterprise and at organizations that work with or rely upon the USAF ISR enterprise.¹⁶ The project team supplemented these interviews by obtaining and reviewing documents describing USAF doctrine and procedures, prior ISR assessments, and observations about challenges in conducting assessments and potential improvements that could be made. To better understand how and why airborne ISR is employed, we employed an approach similar to the strategy-to-tasks framework outlined in Lingel et al. in which broad strategic goals are broken down into increasingly granular activities and tasks required to achieve those goals. In the context of this research effort, this method was employed as a step in the process towards tracking USAF airborne ISR's value (discussed in Chapter 3). We developed and demonstrated quantitative data analysis methods using both operational data (i.e., data collected by USAF and other organizations during the course of live operations) and computer simulations, as discussed in Chapters Five and Six. Finally, the project team explored a limited number of industries outside of the USAF that face similar issues with valuing information (e.g., managing collections within the IC, how online search algorithms

¹⁶ Including, but not limited to, the Office of the Under Secretary of Defense for Intelligence, the Joint Staff, HQ AF; HQ ACC; Headquarters, USCENTCOM; Headquarters, U.S. Pacific Command (USPACOM); U.S. Air Forces Central Command; U.S. Pacific Air Forces (PACAF); AFSOC; the 25th Air Force; and the 480th ISR Wing.

determine the order in which information in a search is returned to the user, ride-sharing service dynamic pricing and feedback, the allocation of limited medical resources, quantifying the value of academic research). The primary analytic source material used to inform our discussion and findings regarding each step of the USAF airborne ISR assessment process summarized above is detailed in Figure 1.1 below.

Figure 1.1. Primary Source Material Used to Synthesize Findings About USAF Airborne ISR Assessment Steps

Assessment step	Primary source material for analysis
1. Identify the decision to be informed	Joint publications and Air Force instructions/doctrine documents; expert interviews within the USAF ISR Enterprise and Joint Community (Joint Staff, CENTCOM, PACOM)
2. Select relevant cost and benefit indicators	Expert interviews with analysts, primarily those working within and supporting missions in CENTCOM and PACOM; review of selected IC, “smart market,” ride sharing, online search algorithm, scientific journal index, and healthcare industry practices
3. Collect data on these indicators	Expert interviews with analysts, database maintainers, and programmers; documented project team lessons from using existing databases; review of selected IC and healthcare industry practices
4. Analyze the value of ISR	New PAF cost-benefit framework; data insights derived from analysis of step 3; review of IC practices; expert interviews with analysts, primarily those working within and supporting missions in CENTCOM and PACOM; computer simulation; previous RAND research
5. Apply results to decisionmaking process	Expert interviews and site visits within the USAF ISR Enterprise and Joint Community (Joint Staff, CENTCOM, PACOM); review of selected IC, smart markets, ride sharing, online search algorithm, scientific journal index, and healthcare industry practices

Consistent with the sponsor’s original concerns, we focused the initial interviews and illustrative analyses on USAF MQ-1 and MQ-9 collection of FMV and SIGINT in the USCENTCOM AOR in current and recent (within ten years of project initiation) operations. However, several interviews and some aspects of operational data analysis were conducted at a larger—in some cases, global—geographic scale and included a wider range of USAF airborne collection platforms, missions, and operations to ensure that project findings and recommendations would be applicable across the USAF ISR enterprise (at least for airborne platforms). Further, we conducted this research with a mindfulness that, in many cases, organizations conducting ISR assessments for the USAF or air component are undermanned.

Further, the USAF airborne ISR assessment framework presented here may not apply to challenges experienced elsewhere within the IC and Joint Community. We also do not explicitly address how to explore challenges related to articulating collection requirements, examining impacts of adversary countermeasures (denial and deception), understanding the root causes of intelligence failures, and other complex issues that would require even further development and application of an ISR assessment framework than displayed in some of our examples in Chapters Five and Six.

The remainder of the report is organized according to the assessment steps outlined above, starting with a discussion of different resourcing and employment decisions that the assessment is intended to inform and how to evaluate benefits and costs in different contexts. The final chapter concludes with a summary of our major recommendations to USAF. Some of these recommendations are actionable by USAF alone. Others require collaboration between USAF and partners in the Joint and broader IC communities.

Definition of Terms

There are several terms that will be used throughout the report that we define here. These are listed in Table 1.1 below.

Table 1.1. Definition of Terms

ISR	(intelligence surveillance and reconnaissance) "... synchronizes and integrates the planning and operations of sensors, assets, and processing, exploitation, and dissemination systems in direct support of current and future operations..." ^a
ISR mission or sortie	The dispatching of aircraft to accomplish one particular task. ^b
ISR role	ISR Role Categorization of the functions ISR assets execute. ISR roles are resource-agnostic and instead emphasize ISR capabilities to resolve an intelligence gap. ^c
ISR planning and direction	One of six categories of intelligence operations in the joint intelligence process, planning and direction develops intelligence plans, determines intelligence requirements, and manages their execution. ^d
ISR collection management	One of six categories of intelligence operations in the joint intelligence process, collection management is the process by which intelligence-related information requirements are converted into collection requirements, thereby establishing priorities and tasking and coordinating resources for collection execution. ^d
Assessment	As assessment is at the core of this report, we define it in three similar, yet distinct, ways: "1. A continuous process that measures the overall effectiveness of employing capabilities during military operations. 2. Determination of the progress toward accomplishing a task, creating a condition, or achieving an objective. 3. Analysis of the security, effectiveness, and potential of an existing or planned intelligence activity." ^e
Indicator	Indicators are used to track progress toward mission or operation accomplishment and allow us to evaluate effectiveness in a repeatable, possibly quantitative way. ^e

^a Air Force Doctrine Annex 2-0, *Global Integrated Intelligence, Surveillance, and Reconnaissance Operations, Introduction to Global ISR*, Washington, D.C., January 29, 2015.

^b Joint Publication 3-30, 2014.

^c Office of the Director of National Intelligence, ACC, *ISR Assessment Framework*, white paper, Washington, D.C.: September 27, 2016.

^d Joint Publication 2-0, *Joint Intelligence*, Washington, D.C., October 22, 2013.

^e Joint Publication 3-0, *Joint Operations*, Washington, D.C., January 17, 2017.

In the chapters that follow, we dig deeper into each of the assessment steps, ending each chapter with specific suggestions as to how USAF can start improving its overall ability to conduct analysis of airborne ISR at each step. We begin with a close examination of the first step, which in some ways is the most important because it defines the direction for an entire assessment.

2. Step 1: Identify the ISR Resourcing and Employment Decisions

In Chapter 1, we introduced the concept of weighing costs and benefits as a way of understanding USAF airborne ISR's value. Importantly, this value cannot be illustrated or calculated until we answer the question: *value for what?* Costs and benefits vary depending on the context of the USAF airborne ISR resourcing and employment decisions they are intended to inform. As this chapter discusses in detail, the USAF ISR enterprise involves many organizations with different responsibilities that have different reasons for conducting assessments. Therefore, the first step in an assessment is to understand the types of resourcing and employment decisions that it is meant to inform. In this chapter, we survey a spectrum of decisions concerning USAF airborne ISR that spans echelons, timescales, and parts of the ISR enterprise. We examine different types of decisions based on our expert interviews and examination of doctrine documents.¹ We then discuss how a cost-benefit analysis framework can be adapted to perform consistent USAF airborne ISR assessments across these contexts. This consistency in terms of our general approach is important for the USAF airborne ISR community for two primary reasons: (1) to help ensure that any organization with a need for these assessments could have guidance for how to do them and (2) to enable more effective communication between organizations when assessments can affect more than one party. Although some consistency might also benefit relationships with the Joint Community and IC, organizations outside the USAF airborne ISR community have different functions and needs and may thus require different approaches to assessments.² This framework guides the selection of indicators, data sources, and analytic methods for any given assessment, as described in later chapters.

Why Context Shapes USAF Airborne ISR Assessments

Assessments are fundamentally tools that USAF and air component ISR decisionmakers use to inform their resourcing and employment decisions. The decision(s) being informed by an analysis must be understood before any analysis can be done. Knowing which decisions will be informed guides the development of indicators, gathering of data, and interpretation of results. For example, a decision about redistributing ISR support from one theater of operations to another will require knowing what operational efforts ISR were supporting in each theater, at what level of effort, and to what effect. With this information in hand, a member of the Joint Staff community or other qualified decisionmaker can weigh the costs and benefits of using ISR capabilities in each theater and compare these with national strategic priorities to help inform

¹ For example, Air Force Doctrine Annex 2-0, 2012; Joint Publication 2-0, 2013; AFI 13-1, 2012.

² That is, we do not intend to suggest that “one size fits all.”

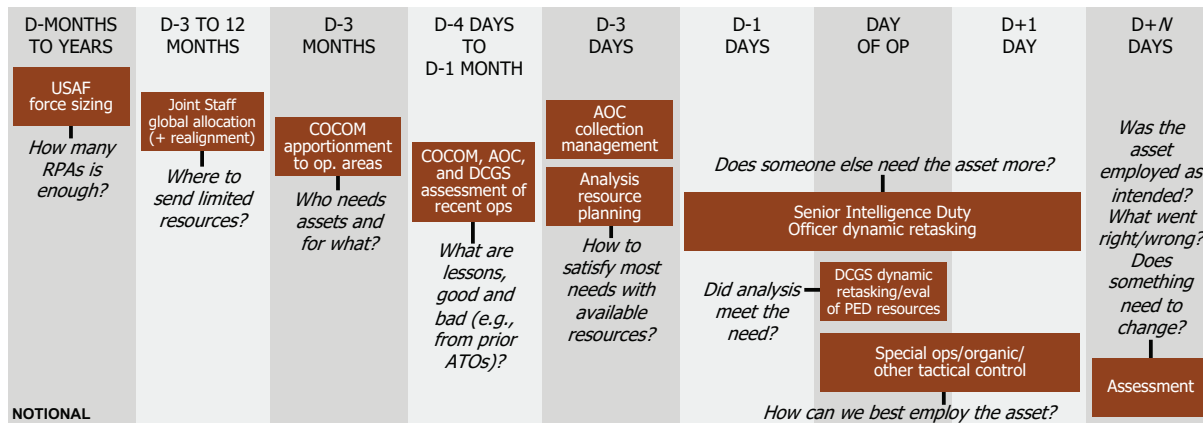
which capabilities should go where. In contrast, a theater collection manager frustrated by a lack of requirements satisfaction by a particular sensor will want to focus on the particulars of how that sensor was employed and its data exploited in different contexts to determine whether weather, limitations on sensor capabilities, maintenance, PED availability, or other constraints are contributing to the problem.

However, it is too often the case that USAF ISR assessments are designed to accommodate analysts' limited time and data, rather than to inform particular decisions about ISR. In some cases, this leads to only partially answering the most relevant assessment questions, or even answering irrelevant or incorrect assessment questions. Answering the wrong assessment question can be as unhelpful as not conducting an assessment at all or can even be more damaging if the analysis suggests a course of action that exacerbates problems. For example, an assessment that emphasizes the large number of ISR sorties being flown in a particular theater could conceal the fact that these sorties are not yielding useful information because of sensor malfunction or mismatch. A resulting decision could be that the operational theater in question is overlooked for additional, potentially more helpful ISR resources ("they're already getting a lot"), and other theaters that could use the extra ISR assets do not get them.

Types of Questions USAF Airborne ISR Assessments Could Help Answer

Management of the USAF airborne ISR enterprise requires a wide range of resourcing and employment decisions at different echelons and over different timescales. Decisions about ISR in general can range from long-term investments in platforms and capabilities (years) to allocation of resources to specific theaters (months) to employment of resources in operations (hours, days, and weeks). Figure 2.1 illustrates a notional timeline of events leading up to and following a day of operation supported by ISR capabilities. In this figure, the orange boxes represent types of resourcing decisions that are made by different individuals and organizations at different points in time. The question associated with each box is one an USAF (or other) ISR assessment could ask to help support the resourcing decision. The figure provides a generic sampling of the types of ISR assessment questions that arise based on the resourcing and employment decisions that must be made at different points in time, as revealed during our expert interviews and review of doctrine documents. This timeline includes types of decisions and questions that are not strictly considered within the USAF airborne ISR community, but they are included here in order to present a more complete picture. As this figure illustrates, these decisions about ISR cover a broad temporal and topic spectrum but are also related to each other. For example, investment and global allocation decisions made years and months, respectively, prior to the operation will ultimately impact what ISR capabilities can be drawn upon to support the operation. In turn, information about how ISR supported a particular operation could theoretically be included in a longer-term analysis that may help future investment and allocation decisions.

Figure 2.1. Timeline of Example ISR Decisions Leading Up to and After an Operation



Answering each of the questions in Figure 2.1 requires different sets of indicators that rely on the evaluation of various quantitative, qualitative, or both kinds of data sources derived from real-world observations and, in some cases, via simulations. Some assessments may rely on common data that is analyzed over different timescales, on using different definitions of what constitutes ISR “success,” or on both. We illustrate how some of these differences shape analysis approaches in later chapters. Despite these differences, however, the connections illustrated in the figure between resourcing and employment decisions made in different contexts suggest that using a common assessment approach, data standards, lexicon, and other guidelines will help ensure continuity between decisions made in different parts of the USAF airborne ISR community. Having some similarities (e.g., in data formats) with assessments conducted elsewhere in the USAF, Joint Community, IC, or some combination thereof may also have benefits, but we did not directly assess these. Furthermore, too much similarity between assessment approaches between very different organizations could inhibit the flexibility needed to solve very different kinds of problems.

In what follows, this chapter summarizes the different USAF organizations that are involved with assessments. It then looks more closely at some potential roles for assessment in each part of the timeline in Figure 2.1, beginning with long-term ISR asset portfolio planning. This discussion will further illustrate why identifying the ISR resourcing and employment decisions central to a particular assessment is so important.

Now, we turn to some specific examples of assessment questions that might be asked within some of the organizations described here at different points along the timeline in Figure 2.1. Again, some of these questions are not strictly within the decisionmaking purview of the USAF airborne ISR community; however, these questions all impact that community in one way or another and are thus worth discussing for the purposes of awareness. Further, there are instances in which the USAF airborne ISR community will need to be aware of the results of ISR assessments conducted elsewhere in order to better inform its own decisionmaking.

ISR Assessment to Support Long-Term Planning

Long-term planning is used by USAF to acquire assets and develop an optimal ISR portfolio. We define planning as separate from tasking because here it is referencing long-term force sizing and acquisition decisions. Planning assessments are needed to inform acquisition decisions and serve to identify capability gaps and the potential for added value from additional or more-capable assets. To this end, assessment in this category requires the analysis of which assets are being used the most often and for what purpose. This analysis should also identify when capability gaps exist and whether and which additional assets could fill those capability gaps. This is one area in which understanding how assets outside of the USAF airborne ISR community can answer key intelligence questions is useful, so that future investment plans consider overarching Joint Community and IC needs and gaps.

ISR Assessment for Global and Regional Allocation of ISR Resources

In each fiscal year, the USAF has a certain amount of ISR assets and associated PED capabilities in its portfolio that are available for use worldwide within the context of the Joint Community and IC. Use of these capabilities requires forward planning to allocate them to the appropriate theaters of operations. This is done to help ensure that when ISR is requested to support information needs, these capabilities are theoretically available in roughly the quantity needed at any given time. For example, an RC-135 Rivet Joint located halfway around the world from an AOR that needs it is unlikely to be terribly responsive in helping address the associated commander's intelligence priorities. Forward planning is needed to determine how many of these Rivet Joints can be sent to the AOR in question (e.g., based on its priority and needs), where these aircraft can be based to be close enough to be useful, who is available to analyze the information collected, and other matters that help ensure ISR operations are as smooth as possible.

The process of dividing up the portfolio occurs approximately every 12 months with a Joint Staff-moderated process in which each COCOM requests use of certain capabilities. This planning process is a complex negotiation that is based on both historic need patterns and projections of future requirements. There is some flexibility in asset allocation subsequent to these annual decisions, which are reviewed periodically throughout the year to allow for adequate response to a large unforeseen demand (for example, in 2014, USCENTCOM anticipated an Iraq drawdown, which did not materialize). However, COCOMs are unlikely to gain access to additional ISR resources during the year, other than in circumstances where demand for ISR changes for a high-priority operation. This can sometimes lead COCOMs to request as much ISR as possible, in order to ensure that each will have enough at any point during the year. As a result, there is an increasing push to justify ISR demand projections and to put processes in place that might enable a rapid temporary reassignment process to become more responsive to meeting demand.

Data-driven assessment could play an important role in this process. No one can predict the future, but having the ability to explain historic use and shortfalls and forecasted demand will help the Joint Staff to make informed decisions and the COCOMs to know they are making requests that can reasonably meet their demand. Examples of the types of analytic results that could drive better allocation would be a COCOM showing data from dozens (or more) ISR missions that demonstrate a sensor mismatch with important targets.

Within a COCOM, ISR portfolios are allocated and reviewed on a somewhat more frequent basis during the year, to ensure that resources are appropriately allocated across the different AORs for which the COCOM is responsible. Resources can be transferred during the year to different AORs should the need arise. However, because of fears of losing an asset should it be loaned out, even temporary reassignments are associated with vigorous debate. Any insight into what contributions assets are having toward positive strategic, operational, or tactical outcomes would help clarify the case for either keeping or moving those assets. Furthermore, explanations for information gaps linked to insufficient ISR resources would aid additional apportionment. Hopefully, quantitative analysis underlying these apportionment decisions would also make operational commanders less anxious about “losing” an asset, because such analysis could also justify bringing it back.

ISR Assessment for Collection Management and PED

The ISR PCPAD cycle broadly summarizes all the stages associated with a single ISR mission.³ PCPAD is typically thought of as a cycle, because the end result of one mission theoretically feeds into planning for the next mission. The primary USAF entities involved in the PCPAD process include the AOC for planning and direction, pilots and sensor operators for collection, and the analysts within the DCGS and intelligence partners within and external to USAF for processing, exploitation, analysis, production, and dissemination.

We can examine each of the PCPAD stages in turn to highlight the unique characteristics of the decisions made in each stage of the cycle. These decisions are directly tied to the types of outcomes that need to be used to do the assessment analysis. This section details how the appropriate approach to ISR assessment depends on the decision context dictated by the PCPAD cycle and how the types of assessment questions to be asked and answered at each point in the ISR cycle are in many cases characteristic of a particular phase within in this cycle. Within each of the parts of the PCPAD cycle, an understanding of how ISR assets were used and the outcome of this utilization should inform both prior and future decisionmaking.

ISR Planning and Direction is defined as: “the determination of intelligence requirements, development of appropriate intelligence architecture, preparation of a collection plan, and

³ Air Force Doctrine Annex 2-0, 2012.

issuance of orders and requests to information collection agencies.”⁴ Planning and Direction identifies intelligence needs for answering the commander’s Critical Information Requirements then uses PIRs to develop detailed EEIs. For each operation type, an assessment process should allow decisionmakers to make the operation more effective and efficient. For planning and direction, this assessment should provide better visibility into which intelligence requirements best support mission objectives and how well. Assessment of planning and direction should therefore answer questions like: *How well do EEIs match mission objectives? Does the intended collection match the mission prioritization?*

Once these collection requirements are established, these requirements are aligned with available collection capabilities to determine how ISR assets will be employed.⁵ Tasking includes the work of collection management and mission planning in order to prioritize and coordinate ISR assets intelligently. When done successfully, collection management and mission planning will task “appropriate assets or resources to *acquire the data and information required.*”⁶

The tasking stage of the ISR cycle includes the identification of information required to support objectives (PIRs and their corresponding EEIs). These are used to generate an integrated prioritized collection list, after which subject-matter experts (SMEs) must decide how to satisfy the EEIs and then ISR assets are tasked. Once these have been identified and prioritized, collection managers decide how to satisfy requests in order of priority with existing limited assets.⁷ Assessment useful for collection management will help them answer questions like: *is this collection sufficient to meet a prioritized requirement? Are assets and collections being used as tasked or requested? Is there a more optimized way to use this asset?* Tasking assessments are needed to inform future allocation decisions to most effectively and efficiently meet ISR demand with limited supply. Analysis in this area should help collection managers more beneficially allocate ISR assets by helping to determine the most appropriate ISR asset for a given collection and the relative value of each collection. This will enable each asset use to be optimized for the maximum ISR value across the entire prioritized collection list.

It is important to consider that, within an AOR, there may be ISR assets owned or employed by organizations other than USAF. Assessments are also valuable for considering, as an example, how to make best use of assets across different military services within a Joint operation. As with Joint Staff decisionmaking, having a consistent assessment approach and

⁴ Joint Publications 1-02, *Department of Defense Dictionary of Military and Associated Terms*. Washington, D.C.: Joint Chiefs of Staff, November 8, 2010.

⁵ Air Force Doctrine Annex 2-0, 2012.

⁶ Air Force Doctrine Annex 2-0, 2012.

⁷ Carl Rhodes, Jeff Hagen, and Mark Westergren, *A Strategies-to-Tasks Framework for Planning and Executing Intelligence, Surveillance, and Reconnaissance (ISR) Operations*, Santa Monica, Calif.: RAND Corporation, TR-434-AF, 2007.

lexicon is helpful for understanding the various ISR contributions different services and other organizations can make.

In contrast to planning and direction, collection is the act of gathering data. Here, the assessment focus is on the appropriateness of asset-target pairing and the method of asset employment in order to gather the required information. Therefore, assessment of collection should answer questions like: *How are current assets being used according to mission type, and what is the outcome? How are sensors being tasked and how likely are they to satisfy the EEIs?*

Assessing the execution of a collection mission will help identify when and why collection requests are not satisfied, and in turn inform future mission planning. The evaluation of the appropriateness of collection techniques for answering specific types of intelligence questions in a particular AOR is a very important function of ISR assessment. If a collection mission fails (e.g., is unable to satisfy the EEIs), then an examination of the reason(s) for this failure could result in, for example, a decision to use another asset for similar missions in the future, a decision to layer more assets to achieve a higher success rate, or the identification of EEIs that are not appropriate to the task. Therefore, the questions for assessment are *whether the asset performed as tasked, and if not, why not?*

Processing and exploitation turns collected ISR data into intelligence with the objective of advancing the overall mission supported by this ISR. Assessment in this category requires analysis that determines if the collection contributed to mission success and if the collection was sufficient to meet the intelligence requirements. Since processing and exploitation priorities need to be synced with the PIRs, assessment here should address: *Is there sufficient PED to support high-priority missions? Were tasked sensors sufficient for the collection goals? Were PED analysts appropriately apprised of the context of collection?*

How Do We Measure ISR's Value in Context?

The above discussion shows that the purpose of assessment can vary substantially, depending on the resourcing and employment decision(s) that are being supported. Once this context has been identified, the assessment analyst can begin to identify the indicators and types of data that will be relevant to the assessment of USAF airborne ISR. A useful starting point is to form a hypothesis that can be tested. Examples of notional hypotheses relevant for USAF airborne ISR (and perhaps other) assessment include: *ISR has not helped provide warning of a potential adversary's maritime movements this past week because of severe weather issues. A month-long effort to collect imagery on potential new missile sites helped inform national diplomatic strategy. Analysts focused on identifying road-side explosives successfully diverted most troops this year. Yesterday's collection did not advance the narrative for understanding our highest-priority target's movements.* Making an “educated guess” (or even an uninformed one) provides two benefits for ISR assessment. First, it links the decision(s) informed by the assessment to the benefits the assessment is trying to measure. Second, it generates a focused goal for the

assessment. Using questions to structure assessments (e.g., *Why were we surprised by the adversary's maritime movements? Did ISR help troops avoid road-side explosives?*) can also be helpful, but such a format can easily drive analysis that is too broad or otherwise undefined. A hypothesis (or hypotheses, for some assessments) forces the specificity needed to generate a useful ISR assessment.

Although the specifics of the approach will vary between different assessments questions and hypotheses, it will be important for the USAF ISR enterprise to define a common understanding of how this approach is formulated in any given context. Again, this will help ensure consistency across the enterprise and enable links to be made between assessments and resourcing and employment decisions made at different echelons and timescales, in addition to avoiding having personnel learn completely new assessment practices when they move between organizations. As discussed in Chapter 1, this type of overarching assessments guideline is not currently available to members of the enterprise, though the office of the ACC A2 and others are working towards developing one.

In this report, we propose a basic assessment framework for USAF airborne ISR that seeks to quantify and weigh the benefits and costs of ISR operations in a given operational context. Although cost-benefit analysis is not a new approach, we integrated findings from an examination of practices within the Joint Community and IC (to include the USAF ISR enterprise) as well as examples external to these communities described in Appendix C—to distill an approach that is tailored to needs and challenges for USAF airborne ISR assessment discussed in Chapter 1. This proposed approach serves as a starting point for USAF airborne ISR assessment that will necessarily be shaped by each individual use case. This is a starting framework; the final assessment approach will be shaped individually by the different users. Wider IC and Joint Community ISR challenges and needs may be quite different from those of the USAF airborne ISR community; thus, the assessment approach discussed here may not widely apply. The assessment framework described below is organized around four main questions developed by the research team, and constructed based on interviews with USAF analysts and the examination of practices within the IC, Joint Community, and elsewhere:

1. How relevant was the information to the operational decisionmaker?
 - a. Was there a clear understanding of the specific information needed to support the decisionmaker?
 - b. What was the additive value of the information collected?
2. Was the decisionmaker able to make a decision with the available information?
3. What functions did ISR resources perform?
4. What ISR resources were expended in those functions?

Questions 1 and 2 examine benefits in different ways, both of which are important. Interviews with intelligence analysts and an examination of assessment practices within one particular IC organization highlighted the importance of answering both questions to examine

benefits from complementary perspectives.⁸ The first asks whether ISR was successful in being relevant. Put in terms of what we learned from our examination of IC intelligence assessment practices and in interviews with analysts working in the Pacific theater, this answers whether ISR *advanced an important intelligence narrative*. Did we make incremental progress towards understanding an important unknown? The second question examines the overall effect of ISR on some broader operational goal. The latter might include understanding potential adversaries' military capabilities (e.g., airfield lengths and construction), providing strategic warning of diplomatically significant events (e.g., breaking of treaty terms), or force protection (e.g., reducing casualties). This, too, was revealed as important in the examination of IC intelligence-assessment practices where (as illustrated in Chapter 5) incremental intelligence goals were linked to broader operational objectives. What was the overarching impact of our efforts? Interviews with analysts working in the Middle East theater suggested some ways in which ISR could be linked to operational effects, something we discuss in more detail in Chapter 2.

Answering the first question requires subject-matter expertise from analysts at the DCGS, AOC, IC, or a combination of those in order to confirm why and how a particular collection or set of collections is relevant (or not) for filling key information gaps. Addressing the second question requires the analyst to identify one or more indicators (and associated data) that point to the supported commander's ability to make a decision based on sufficient information. Importantly, these indicators must consider situations where ISR is able to support a good judgment not to take action, even if that action is ultimately desired (e.g., not striking a high-value individual's [HVI's] safehouse when noncombatants are present).

Both approaches to understanding benefit are needed for several reasons, which include the multiple confounding factors that can influence strategic, operational, and tactical outcomes (i.e., ISR are one of many factors at play) and the additive nature of intelligence information (i.e., important insights are often gained through extended periods of collection and analysis and sometimes the ultimate value of ISR activities is only understood later on). An analyst may not need to answer both questions for a given assessment and could rely periodically on only one or the other. Over the long term, however, a team or organization will need to answer both questions to provide a complete picture of ISR benefit.

Questions 3 and 4 characterize different aspects of cost, which we argue must be measured both in terms of the numerical amount of resources expended (e.g., sorties, sensors, manpower) and in terms of how these resources were used. It is perhaps easiest to count how many sorties were flown in support of a given mission, but this information is not relevant to a cost-benefit assessment unless we also know what those sorties were doing and how those tasks related to the larger ISR and operational objectives. Interviews with experts involved in various aspects of collection management revealed the continued importance of documenting the number of

⁸ Due to potential sensitivities, we do not reveal the organization's name, but an illustration is discussed in detail in Chapter 5.

aircraft, types of sensors, number of people, and other measures that might be dismissed as “bean-counting,” in order to track what resources are being used. However, these numbers should be supplemented with better understanding of *why* the resources are being expended, according to interviews with those managing assets and collections at different echelons, to enable better precision and efficiency.

The cost-benefit approach proposed above provides a common framework for thinking about and performing USAF airborne ISR assessments at many echelons and for many resource-decisionmaking contexts. Armed with this framework and an understanding of the larger context, and with a working hypothesis of the potential factors that may have influenced outcomes, the analyst is ready to select the indicators of cost and benefit that will be most relevant to the assessment. We discuss these indicators in the next chapter.

Conclusion and Suggested USAF Actions to Improve Assessment

Step One

Every USAF airborne ISR mission is itself tied to several decisions. For example, which aircraft carrying what sensors should be used? How many aircraft must be sent, how often, and for how long? Over time, ISR employment is tied to bigger decisions, ranging from overall force sizing to how many aircraft to apportion to regions within an area of responsibility and the distribution of those assets amongst different types of use (e.g., tracking high value assets or missile movements) within different regions. ISR assessments are needed to support decisions about ISR investments, allocation, and employment across timescales ranging from years (e.g., in preparation to purchase new aircraft) to minutes (e.g., for calculating risk in reassigning resources to an emerging problem).

Addressing each of these resourcing and employment decisions—and others like them—consistently and meaningfully requires systematic and widespread use of common approaches for understanding ISR’s value (or costs and benefits) as well as standard lexicons and databases in order to ensure consistency and translatability of findings between different assessments. As described in Chapter 1 and reiterated in this chapter in more detail, the USAF airborne ISR community presently finds itself challenged in both of these aspects of generating enterprise-wide ISR assessment activity: Although several organizations (and even teams within those organizations) perform assessments, there is little use of common approaches, assessments lexicons, and means for understanding value in the context of a particular decision. This may be the result of lacking overarching guidance on assessments—something HQ AF/A2 and ACC and others are attempting to rectify—and organizational processes and outputs that are siloed, and hence encourage individualized assessments. Even the AOCs and DCGS, whose missions are neatly related through the PCPAD cycle and who function on similar timelines have in recent years spent a large amount of effort and relied upon their own initiatives to develop a more common way of thinking about, conducting, and sharing assessment information.

This chapter has highlighted the importance of understanding the ISR resourcing and employment decisions that assessments can inform and has explored some of the contexts in which assessments can or could be helpful. Although the precise assessment design and indicators used will vary depending upon this context, we have emphasized that it is important to employ a consistent overarching approach to understanding ISR value for the purposes of assessment. We have suggested that using a cost-benefit analysis framework, organized around four generic questions, could provide needed direction for assessments across the USAF ISR enterprise. In combination with a widely-used lexicon, this generic framework can help ensure that assessments across the USAF ISR enterprise speak a common language that is useful for enabling decisionmakers to make resourcing and employment decisions.

The final chapter of this report presents a comprehensive list of recommendations to help USAF develop or improve its ability to carry out ISR assessments. However, we will conclude each intervening chapter by highlighting actions USAF can take to improve the particular assessment step under consideration. With respect to Step 1, USAF can make the following improvements, based on our compiled findings from interviews, site visit observations, and examination of doctrine documents:

- Ensure that the reasons for requested assessments related to USAF airborne ISR are clearly indicated and that the importance of understanding this context is acknowledged in doctrine and standard operating procedures, such as those governing AOC and DCGS operations and those communicated by AF/A2 and ACC/A2 to the USAF ISR enterprise.
- Establish a USAF-wide ISR assessment directive that will enable a common understanding of how to use a cost-benefit framework to value ISR activities in different contexts. Such a framework could emphasize the importance of formulating assessment hypotheses and addressing the four questions outlined in this chapter.
- Choose an existing or new Air Staff-level role to help manage ISR enterprise- or USAF-wide assessments activities. Developing and encouraging the use of common assessment practices will probably require a position devoted to providing umbrella guidance and advocating for consistency in assessment approaches and data across the ISR enterprise, and perhaps across USAF as a whole. Otherwise, it is possible that good progress in USAF data science could be lost if this area is just a “pet project” that appeals to one leader at a point in time, but not necessarily to those that follow.

In the next chapter, we turn to the selection of indicators for cost and benefit using the questions posed above.

3. Step 2: Select Cost and Benefit Indicators

As discussed in the previous chapter, precisely defining analysis questions and hypotheses is important for guiding the overall USAF airborne ISR assessment effort. The next step, then, is to determine which suite of quantitative and structured qualitative indicators will appropriately illustrate the benefits and costs of ISR over the time frame relevant to the assessment. Importantly, not only must these benefits and costs be defined, they must also be measurable. The previous chapter introduced a series of four questions developed by the research team which broadly outline the types of information we argue are needed to explore benefits and costs (or, in this case, value) of employing ISR in a particular context. Here, we step through each of those questions individually to suggest approaches for determining cost and benefit indicators. Our discussion draws heavily upon the synthesis of findings from expert interviews, especially with those analysts within and supporting operations in the USCENTCOM and USPACOM regions. We also drew upon ideas from the IC, smart markets, ride sharing, online search algorithms, scientific journal indexes, and healthcare industry assessment practices, which we reference within our discussions below and are available in Appendix C. That said, the findings we derive are not necessarily applicable beyond the USAF airborne ISR community. Some examples of indicators are given here for USAF and others to consider, but these are less important than the overall approaches and reasoning behind the general methods, which are the primary focus of this chapter.

How Relevant Was the Information to the Operational Decisionmaker?

Information should not be collected and analyzed simply for its own sake. Ideally, there should be some important reason attached to every collection and PED activity. Importance should be sufficiently loosely defined to account for previously unknown or emerging problem sets, the relevance of which may not yet be known. Adding relevant information to existing data sets and defining new problem sets can be referred to as “advancing the narrative.” ISR that advances an existing or potential new narrative provides a tangible benefit to information users (e.g., the Secretary of Defense, a commander, AOC). Specifically, it helps build a repository of knowledge that could ultimately lead to an enhanced ability to make a good decision about strategic, operational, or tactical actions at the appropriate time (irrespective of whether that decision is to move forward, pause and wait for more information, or decline action).

Illustrating benefit by answering this question requires having and understanding the information needed to support the decisionmaker and interpreting whether the result of the collection and analysis has yielded additive information, including through the creation of

products.¹ Achieving these two things is not simple and requires foresight, planning, and good communication among key players.

Imagine relocating to a completely new place, something USAF members and other military service members do routinely. There are several topics that are immediately relevant (“known unknowns”): moving logistics, housing, vehicles, location of stores and other conveniences, new friends and neighbors, and schools, among other things. Surprises (“unknown unknowns”) might yield other matters of importance: local customs and events, unexpected language or cultural barriers, and the like. Navigating these areas of concern requires gathering knowledge, and not all of it might be “additive.” If the location of a pharmacy is already known, it is not helpful to have a kindly neighbor point out the way. However, it is helpful to meet neighbors. Thus, this neighborhood exchange may not have benefit for understanding where the pharmacy is, but it might be useful for identifying who lives down the street.

The ISR environment is not so different. These resources are applied to gather information about topics of strategic, operational, and tactical interest. Generally speaking, these broad topics are readily identifiable in priority lists and other documents that are periodically updated. What is challenging is understanding all the intermediate pieces of knowledge that must be gathered in order to fill a larger knowledge gap or other area of concern. As a notional example, an operational priority could be to understand a foreign nation’s actions during a military exercise. Achieving this could notionally require tracking aircraft, ships, vehicles, and other vessels, identifying which military units are doing what, and surmising the intended purpose of the exercise. Developing a granular understanding of information needs, which not only include what must be known, but also how and how often different aspects need to be monitored, requires an intensive understanding of the problem set. In some cases, it is even more difficult to track what is already known, what is unknown, and whether new bits of information are emerging.

Once these more-granular information needs have been laid out, it is theoretically possible to track the progress of ISR activities towards answering relevant questions in an additive fashion. This requires establishing a framework to keep track of this information. In the simplest case, a checklist (or spreadsheet) could be used to indicate what granular information needs have been fulfilled within a relevant time period. A more nuanced approach that has been discussed and demonstrated in various formulations in the USAF ISR enterprise involves ascribing fixed scales that describe the value to a particular ISR activity with respect to a granular objective or task. For example, an analyst could use the following four categories to describe the benefit of information gathered:

- none = not gathered or already known and confirmed
- low = already known but not previously confirmed

¹ This includes the number and type(s) of products. Production is a necessary step towards providing relevant information to a decisionmaker.

- medium = not previously known but lacking in one or more critical details (e.g., specific location)
- high = not previously known and containing all critical details.

The challenges associated with developing scales to track information relevance are multifold. First, organizations with common missions or roles must develop a common scale or risk not being able to communicate assessment results. Information must also be stored in an accurate, retrievable format that can be associated with other cost and benefit indicators.

Second, feedback from information users (other than directly reporting or associated analysts) can typically not be relied upon for understanding information relevance. Many information users have numerous demands on their time and cannot always take time to provide thoughtful feedback or work closely with collection managers to tailor ISR. Employing simple, automated feedback processes could help to some extent and is discussed at greater length in Chapter 7.

Finally, subject-matter expertise is needed to both identify what important knowledge advances are as well as evaluate each individual collection and analysis activity. This is particularly important because of the previous point made regarding the lack of feedback that can be generally expected. However, analysts at the AOCs, DCGS, and elsewhere are generally quite busy simply responding to urgent requests. Still, if USAF needs to conduct assessments where ISR benefits are considered, these analysts will certainly have to play a large role in feeding those assessments, perhaps as part of the postmission procedures they already undertake. Further, ambiguity in the results from intelligence collection is also something important to consider. Deriving benefit from USAF airborne ISR collection to the operational decisionmaker should account for ambiguity in the results and interpretation of collection. Tracking ambiguity and effects on decisionmaking will be an important function of ISR assessment in some contexts.

Three advances may help these analysts consistently provide feedback. The first is to enable information rating scales in common databases containing post-mission analyses and other information. These should not require analysts to devise free-form text, but rather involve check-boxes or drop-down menus with limited choices and an optional comment box. The second is to organize analysts into groups that work on specific problem sets, helping to ensure sufficient contextual understanding. To some extent, this shift towards a “problem-centric” force presentation, rather than a “platform-centric” construct is already present in parts of USAF. The third is to continue to support efforts for prototyping and then fully developing applications and tools to reduce other aspects of analyst workload (e.g., transfer of data between different formats) that are routinely performed by computers in other industries. In time, this will place greater analyst focus on tasks that require higher-order analytic skills. Benefits from these three advances are contingent upon having sufficient numbers of analysts with the appropriate training and career depth to conduct assessments, which is not the case for most—if any—USAF ISR organizations in 2017.

In the long term, understanding the relevance of ISR collections will be aided by concepts such as object-based production (OBP), if the enterprise shifts analytic efforts in this direction. Put simply, OBP is an approach for creating a structured data environment in which pieces of information can be related to each other through different attributes, such as location and topic. There are ongoing initiatives within the IC and USAF to facilitate expanded use of OBP as a means of structuring information and as a tool for aiding different forms of intelligence tradecraft. For the purposes of ISR assessment, OBP might provide a means of quickly understanding the relevance of recently collected data because it could show the user whether the data leads to conclusions that represent known information or something new has emerged.

Was the Decisionmaker Able to Make a Decision with the Available Information?

The benefits of USAF airborne ISR can also be judged by how these capabilities advance good and timely decisionmaking in pursuit of strategic, operational, and tactical goals. In this sense, we do not mean the ISR resourcing and employment decisions that assessments are meant to inform, as described in previous chapters, but rather the operational decisionmaking that commanders and others make on the battlefield, which is enabled by relevant intelligence.² For example, a theater collection manager might wish to know whether ISR has contributed to decisions that led to finding, avoiding, and neutralizing explosives on particular roads. A strategic decisionmaker may wish to know whether ISR collections can increase her confidence in proposing additional sanctions on a country accused of a malicious act.

It is difficult to quantify, or in some cases even observe, what is ultimately a subjective decisionmaking process. Operational decisionmaking is based on many factors in addition to intelligence. But we can approximate the benefit to decisionmakers by tracking measurable indicators associated with different operational goals. These indicators must be understood in the context of time, place, and other factors that impact their values. For example, maintaining the safety of personnel might be crudely tracked by monitoring the number of people wounded and killed in action. These numbers will be greatly impacted by the length of time considered, location, and whether there is an active conflict going on. Determining whether the employment of ISR contributed to good operational and tactical decisionmaking that helped to keep personnel safe requires deep understanding of how ISR was employed and if safety outcomes varied depending on whether or not ISR was employed. Confounding factors, such as operational tempo, the nature of tactics, or weapons used, must also be held constant. This is true when attempting to link ISR usage to any particular favorable (or unfavorable) set of outcomes.

² Importantly, assessments should maintain awareness of the possibility for adversary denial and deception, which will impact the ability of a decisionmaker to make good judgements using intelligence.

Furthermore, information-gathering is often an iterative, long-term process. A relatively small subset of ISR missions might be clearly deemed beneficial if, for example, a key building is identified as a potential target or the transport of a possibly dangerous weapon is tracked. In both of these examples, however, the mission would likely not have been successful without a large number of other ISR activities that had taken place hours, days, months, or even years beforehand. No one prior sortie yielded dramatic insight; rather, a body of collections set the stage for the key piece of information to be found. Are these numerous prior missions any less successful or examples of poor employment, given that they were required for the big breakthrough?

As a result, it is important to understand and compensate for the fact that many indicators can only approximate (in some cases, weakly) the quality of decisionmaking, and that many other factors need to be considered and eliminated from the analysis before ISR's benefit can be surmised. We found that ISR benefits to decisionmaking can best be approximated by using indicators that—if considered through an analytic lens that filters out the appropriate confounding factors—relate specific ISR activities to their potential impact on good decisionmaking. Whereas Figure 3.1 illustrates some potential indicators for specific, overarching goals that require good and timely decisionmaking, it might be beneficial to take a more ISR-focused approach when it comes to assessments conducted for ISR itself.

One way to think about this is to consider the broad roles that ISR plays in different contexts. The concept of roles has been explored by USCENTCOM, ACC, and the 480th ISR Wing, among others.³ In the example illustrated in Table 3.1, we use four broad roles: intelligence preparation of the environment (IPoE), targeting, warning intelligence, and support to operations. (As we demonstrate in the next section, these roles can be decomposed into increasingly granular functions that track how ISR resources are being used on the “cost” side of the equation.)

³ For example: USCENTCOM, “ISR Roles,” White Paper, undated; Headquarters, ACC, “ISR Assessments Deliverable Review #2,” briefing slides, August 2–9, 2016.

Table 3.1. Example USAF Airborne ISR Roles and Benefit Indicators

RAND PAF ISR Role	Example Indicators of Benefit
IPoE	# of significant indicators of potential change identified for an area of target set, ratio of revisit rate to preferred revisit rate
Targeting	# of objects positively identified (or confirmed not to be) valid targets, average time required to regain track to high value target
Warning intelligence	Percent of significant indicators for which status is confirmed within the time required, whether significant changes in specific indicators have been missed over the last relevant time period (e.g., days or months—depends on level of activity)
Support to operations	Ratio of friendly actions taken against an adversary to adversary actions taken against friendlies, # of times confidence for time-sensitive targets has risen (e.g., from probable to confirmed)

NOTE: PAF = Project AIR FORCE.

A good, timely decision supported by USAF airborne ISR in one or more of the roles in Table 3.1 might result in enabling valuable offensive action or in ensuring robust defenses. For example, beneficial IPoE could highlight promising regions for operations as well as keep-out zones, helping to avoid any undesired outcomes in the event of a conflict. Strong ISR support to operations could help troops navigate around obstacles and dangers, which could save lives. Identifying increasingly granular ISR roles in a particular decisionmaking context will ensure that the benefit indicators identified are relevant both to how ISR is being used and to the strategic, operational, and tactical objectives the decision is supporting.

There are several challenges with the approach described here, though they can be fairly efficiently overcome through the use of good data-gathering strategies and databases; employment of strategies-to-tasks or logic model approaches for linking overarching goals with ISR activities; and engagement with SMEs, who can help identify not only which indicators are relevant but which ones are measurable. To develop benefit indicators, the analyst must understand the operational decisions being made on the basis of information collected (in addition to the ISR resourcing and employment decisions that the assessment will inform). Second, many of these indicators require the analyst to identify events or outcomes that may not be reliably knowable. It is difficult to declare that an explosive device was avoided if no one present knew it was there or it did not explode. Are the number of mobile missile movements truly knowable by an external party? This problem can never be fully avoided, and requires one to compensate by using multiple indicators, examining indicators over a period of time, and (in some cases) employing detailed examples that describe hypotheses about a broader pattern of success or failure. Third, data must be gathered and stored in databases that can be frequently updated to accommodate the new types of information that are gathered for assessment purposes.

It is also important to remember that these indicators themselves do not measure the “goodness” of decisionmaking, which is subject to many other considerations apart from ISR, including strategic or operating conditions and other enablers (e.g., communications, plans). These additional factors should be accounted for in any assessment. One good way to help tease out whether ISR had a beneficial impact on decisionmaking is to compare indicators for similar operations that were conducted with and without ISR support. As we demonstrate in Chapter 6, simulations provide an excellent tool with which to compare operational outcomes with varying levels and types of ISR support while keeping other influential factors (e.g., the mission) constant.

What Functions Did ISR Perform?

We turn next to the cost side of the assessment and examine two types of indicators: the functions that USAF airborne ISR performed and the resources invested in performing those functions.

Determining what USAF airborne ISR operators are doing is important for categorizing costs. Simply counting the resources employed is not helpful, in itself, for understanding the value that ISR provided. To gain a true understanding of ISR’s value, it is necessary to know what functions or roles were carried out and how those roles relate to overarching operational objectives.⁴ The importance of doing this has been recognized by airmen for years. However, it is just now emerging as a more common practice within the USAF ISR enterprise, as USAF examines the infrastructure needed (e.g., databases, a common lexicon, tools) to facilitate regular, consistent recording of the function(s) ISR performs during the course of any given sortie.

The concept of ISR roles, which we briefly introduced in Table 3.1, provides a starting point for associating each ISR or PED mission with a formal, standardized concept of purpose to enable more responsive and responsible accounting, allocation, collection management, and PED. We first examined the ISR roles (and subroles, when available) previously developed by USCENTCOM, ACC, and the 480th ISR Wing. Though these provided a useful starting point, we found their lack of granularity and consistency (between different schemes for ISR roles) limiting. Using EEIs, target names, and other information in reports summarizing individual ISR sorties was both too time-consuming to be practical for routine assessment and, in some cases, impossible, since these data are not sufficiently tracked (as discussed further in the next chapter).

In response to these problems, we developed a “taxonomy” of ISR roles that relates high-level roles to increasingly granular subroles (Table 3.2). The hierarchy leverages schema and feedback from USCENTCOM, US PACOM, PACAF, ACC, the 480th ISR Wing and others, as

⁴ Within the context of this report, we equate the terms *function* and *role*.

well as doctrine and operational data analysis, and was inspired by a strategies-to-tasks approach.⁵

Table 3.2. PAF USAF Airborne ISR Taxonomy First, Second, and Third Order Roles

RAND PAF First Order Role	Second-Order Roles	Third-Order Roles
IPoE	Evaluate the environment	Characterize the operating environment Identify planning considerations with respect to potential threads
	Collect intelligence on the adversary	Identify and characterize potential targets Evaluate adversary capabilities (Blue's perception of Red)
Targeting	Non-HVI target development	Support basic, intermediate, and advanced target development
	HVI targeting	Characterize, identify, and target HVI-associated combatants Identify and target the HVI Determine the disposition of noncombatants
	Postengagement support	Assess strike Support plans for follow-on action (if necessary)
	Indications and warnings (I&W)	Identify emergent instability and violence, military acquisition events, and mobilization
Warning Intelligence	Fixed-point security	Conduct perimeter surveillance Monitor known and suspected threats
	Overwatch of maneuver forces	Search for adversary positions and activity ahead of the maneuvering force Monitor identified or suspected threats
Support to Operations	Dynamic targeting (time-sensitive engagements)	Real-time collection on target and environment Support engagement in real time

This taxonomy, which includes five levels of granularity (the fourth- and fifth-order roles are available in a classified appendix to this report), was constructed using a structured, iterative approach. We started with the most recent US CENTCOM ISR roles and sub-roles available at the time the research was being performed. After reviewing hundreds of reports detailing ISR sorties and conducting interviews with analysis, we worked toward defining component roles at five levels of granularity:

- **First-order** roles are broad categories of functions ISR can perform (e.g., targeting, support to operations).

⁵ David E. Thaler, *Strategies to Tasks: A Framework for Linking Means and Ends*, Santa Monica, Calif.: RAND Corporation, MR-300-AF, 1993; Rhodes, Hagen, and Westergren, 2007.

- **Second-order** roles are important subsets of first-order roles. These generally match the USCENTCOM ISR taxonomy, unless no sub-roles were specified by the COCOM. In the latter case, we propose second-order roles in line with doctrine or understanding developed from examination of recent operations in USCENTCOM (e.g., HVI targeting, overwatch of maneuver forces).
- **Third-order** roles represent the finer level of detail present in doctrine or standard operating procedures (e.g., determine the disposition of noncombatants, monitor identified or suspected threats).
- **Fourth-order** roles are aligned to general categories of information that could be associated with ISR missions. These are unique types of ISR employment that assets do worldwide (e.g., monitor flow of materiel, identify maritime domain characteristics, conduct functional characterization, identify obstacles along route of travel, determine target presence and location).
- **Fifth-order** roles represent generic types of information that could form the basis for part of an individual ISR or PED mission (e.g., determine type and speed of incoming threat, locate and characterize anomalies, follow personnel who engaged in operational activity, evaluate signs that the hostile entity is an adversary military or terrorist force, broadly characterize communication systems).

While the decision to stop this functional decomposition at the fifth-order level was, to a certain degree, a pragmatic one, we also found that once we arrived at this level of detail, it was possible to answer the question, “was an asset able to fulfill fifth-order role X?” with a binary “yes” or “no.” This is not true of higher-order roles.

As a final step, we evaluated the taxonomy for inconsistencies, logical incongruities, and similarities between different branches linking lower-order roles to higher-order ones. The research team then began working backwards from fifth-order roles back to first-order roles to ensure that each role was unique within the hierarchy and represented a range of activities that ISR might conduct in any circumstance around the world. Part of this process included additional analyst interviews and further review of ISR sortie data to deconstruct what types of activities or functions belonged together on the basis of similarities in analytic purpose and approach. As a result of this step, we decided to condense the first-order ISR roles into IPoE, targeting (including HVI), warning intelligence, and support to operations.⁶

Here, we discuss the taxonomy further, bearing in mind that this work has a distinct USAF airborne ISR focus. IPoE establishes information baselines and changes to those baselines. In our taxonomy, IPoE second-order roles focus on the environment itself, as well as the potential adversary. Understanding the operating environment requires both documenting its characteristics as well as specifically focusing on threats. Basic intelligence on an adversary means both seeking out potential targets as well as determining an adversary’s capabilities (what locations or entities need to be defeated and why).

⁶ The original set of ISR roles that we began working from, based on a USCENTCOM approach, consisted of support to deliberate targeting, force protection and fixed-point security, HVI targeting, intelligence, and operations.

For targeting, there is a need to separate ISR support to standard target development versus HVI, since the latter has special considerations and processes. Poststrike engagement is also included within targeting, because it involves assessing the effect of any strike and providing any information needed if a re-strike is required. Dynamic targeting is not included under the targeting first-order role for two reasons: (1) It represents ISR support that occurs during an active operation and is linked with support to forces, and (2) its time-sensitive nature makes the processes potentially somewhat different from other targeting. It could be argued, however, that dynamic targeting should fall under targeting, that HVI should fall under ISR support to operations, or that both of these are true. With regard to HVI, it is not necessarily entirely a dynamic process, and can represent its own type of operation. Here, and for other roles, USAF and others must evaluate what will make the most sense in practice.

Warning intelligence is divided between the broader, strategic concerns reflected in I&W and the smaller-scale problem set of fixed-point security, which is focused on protection of permanent or semipermanent infrastructure. I&W could be decomposed in many different ways, but we start with signs of instability, military acquisition, and mobilization as three dominant signals that are closely associated with I&W. Fixed-point security focuses on surveillance of a perimeter and threat monitoring.

Finally, support to operations can be thought of as force overwatch potentially leading into dynamic targeting when necessary. Force overwatch means searching for threats and then monitoring potential threats once identified. Dynamic targeting supports then assists making decisions about eliminating confirmed threats.

In later chapters we illustrate how this taxonomy of ISR roles can be used to inform assessments and resourcing and employment decisions. The key point for Step 2 of the assessment process is for the analyst to clearly identify the roles that were conducted in the operation(s) under consideration, to the finest level of granularity.

What ISR Resources Were Expended in Those Functions?

The other key cost indicator has to do with how many USAF airborne ISR resources are expended and what has been produced as a result. This is especially important for answering efficiency questions, but it is also used to correctly place resources. For example, there might be two sensors with different ISR capabilities. Using one may require 100 hours to reach an answer. Using the other may require half that time to reach the same answer and perhaps provide higher confidence. A detailed and readily accessible database of such information can greatly inform resourcing and employment decisions, as well as an understanding of why some operations provided greater benefit than others.

There are several types of indicators that might be useful to include. Broadly, these fall under four categories of data: aircraft data, sensor data, manpower data, and product data. Aircraft data indicate what the platform was, where it took off from, how far it flew, how long it flew, how

long it may have dwelled in a particular area, and other aircraft-related statistics. Sensor data concerns which sensors were used, how often, for which targets, and other sensor-specific information. Manpower data indicates which personnel were used in what quantities, how long they worked, and other information about personnel. Product data summarizes what was produced as a result of an ISR mission and in what quantities, ranging from verbal or electronic interactions to formal reports.

USAF and others already collect fairly abundant data about resources expended. Although this is sometimes dismissed as “bean counting,” it is certainly very relevant for assessments—so long as it is used together with other cost and benefit indicators.

Conclusion and Suggested USAF Actions to Improve Assessment Step 2

This chapter introduced a set of indicators that can be used to capture the costs and benefits of USAF airborne ISR operations. While the selection of specific indicators will depend on the resourcing and employment decisions that are to be informed by ISR assessments, the cost-benefit framework and approach to specifying indicators such as ISR tasks can provide a common language for assessments at many levels of the ISR enterprise. However, these indicators are only useful insofar as there is data to support them. In the next chapter, we discuss the challenges USAF faces in gathering and disseminating consistent data for ISR assessments.

USAF can take the following actions to improve the tracking of cost and benefit indicators as part of Step 2:

- Refine the taxonomy of ISR roles so it applies consistently across the USAF ISR enterprise.
- Develop methods and doctrine that spell out how to consistently track additive collection activities.
- Examine which indicators could be added to specific databases to help track ISR support to decisionmaking and educate analysts and operators to use those indicators (e.g., data sampling strategies, which confounding factors to remove from the analysis).
- Enable analysts to consistently provide feedback on whether the content of new collections are additive; some areas to consider include changes to databases, force presentation, and availability of applications to reduce other aspects of analyst workload.

4. Step 3: Collect Data

At this stage of an assessment, an analyst will need to collect data that indicate the outcome of operations as well as the intent and extent of USAF airborne ISR employment. However, analysts within the USAF ISR enterprise (and, perhaps, others as well) have extreme difficulty with accessing and using operational data. Additionally, they often contend with limited visibility into ISR mission context and get limited feedback from those requesting information (we encountered frustrating problems with the availability and quality of data during our own research). We describe the origins and extent of these problems below and provide short- and long-term remediation approaches that could improve access to quality data for ISR assessments. The insights and conclusions presented in this chapter are synthesized from a number of sources, including expert interviews with analysts, database maintainers, and programmers, as well as documented project team lessons from existing databases, documents,¹ and IC and healthcare industry assessment practices (see Appendix C).

ISR Data Challenges

The current USAF ISR enterprise data environment does not lend itself well to assessments. Data varies in quality and fidelity, and what data is recorded is scattered in different databases, Excel spreadsheets, PowerPoints, and other documents. These problems manifest as a major opportunity cost² to USAF, as decisions are being made without the solid footing data analysis can provide. Conducting even the most basic types of ISR assessment—such as knowing the ground truth on what ISR sorties flew, what targets were collected or historic trends about those targets—is extremely time consuming and impractical as a result of having data and information in varying locations, formats, quality, and fidelity. For example, we learned from one organization that it took them 300 man-hours to document information about a month of ISR sorties (this did not include Army or Special Operations Forces assets), and after all of this work, the outcome of a significant fraction of those sorties remained unknown because of inconsistencies in reporting. This and similar anecdotes illustrate the near-impossibility of ISR assessment in the current data environment. If USAF wants to be able to assess ISR in any meaningful sense, it is going to have to update its data environment, which includes both information technology (IT) solutions and data quality improvements.

¹ Lingel et al., 2008; Office of the Director of National Intelligence, ACC, 2016; Air Combat Command, ISR Assessment Framework (Draft), September 20, 2016; Joint Functional Component Command for Intelligence, Surveillance and Reconnaissance Assessments Division, 2016; Caplan, 2017.

² Opportunity cost is defined here as the loss of potential gain from other alternatives when one alternative is chosen.

Further, USAF ISR data is typically manually recorded, as opposed to machine-generated, which is one of multiple reasons for the data quality and fidelity problems discussed above. For example, the ATO building process has many steps that requires sortie planning that transfers data from spreadsheets, to PowerPoint slides, to Microsoft Word ATO requests. Mission information is also input by hand and then ported to a requirements database.

Redundant, manual transcription of data can be ripe with errors. Generic problems with any type of manually entered data can include typing errors, misunderstandings of when an activity started or ended or if it happened at all, and the insertion of logical fallacies resulting in data mismatches (e.g., geographic coordinates not correlating to a location's name, or takeoff times reported after landing times). Mitigating detected problems is also problematic, especially if the data has made it into an existing centralized database of information. At that point, chasing down the ground truth about a sortie could involve accessing chat logs from months ago, old PowerPoint presentations, or calling someone. Attempting to fill in data about missing sorties is equally challenging, and the process is arduous to repeat.

Another data issue USAF contends with is the many different types of formats that can contain useful information, including briefings or presentations, documents, spreadsheets, and media files. Even if an USAF analyst has a strong programming and data science background (many do not), quickly synthesizing information to find correlations between ISR employment and operational outcomes is very tricky. Consequently, people are forced to read through postmission reports and copy and paste (or manually transcribe again) information into their own spreadsheets. They must also have access to people who know what information needs are and whether ISR is helping meet this need, as well as to people or databases that can provide information on the outcomes of operations or other events that could indicate the quality of decisionmaking by proxy. After accessing this information, assessment analysts must be able to link (or at least infer links between) all of the relevant pieces of information, as described in earlier chapters. This type of data linking is intrinsically impeded by stove-piped data sources and important information not being tagged to collection.

The most comprehensive database of USAF airborne ISR sortie and related PED data is UNICORN, maintained by the 480th ISR wing, but even this does not contain all the information that might be needed to conduct ISR assessments. For example, it does not record information about planned sorties that did not fly, nor does it capture all necessary information regarding operations supported by other ISR wings, such as the 55th and 70th. To meaningfully assess ISR, USAF will have to work with partners to improve the ISR data environment. This will include both IT solutions and data quality improvements, some of which are described in this chapter. Solutions to these problems are multifaceted and ultimately require a reimagining of the way ISR collection is collated and recorded across the ISR enterprise.³ In the sections that follow we outline a phased approach to improving to the data environment. The first phase consists of very

³ What exactly we mean by this reimagining is the substance of much of this chapter.

short-term fixes that can be implemented on the order of months with a focus on UNICORN, which is already providing information valuable to assessments. The second phase is significantly longer term, on the order of years, but requires immediate action to set the groundwork for major improvements. Many of the recommendations that can be done quickly apply to the latter vision as well. The timeline for the longer-term vision could very well be shortened if USAF works to mature and re-engineer its IT platforms quickly. Developments, such as introducing the new position of chief data scientist, are promising signs for steering the service towards a more advanced approach to data management.

Short-Term Mitigation Approaches

The research team examined both short-term and long-term mitigation approaches to USAF ISR data environment problems identified above. This is because, in some cases, discussions with analysts, programmers, and database maintainers revealed some possibilities for “quick fixes” that might improve ISR assessment capability, to some extent, within a few months or years. These types of limited solutions are worth considering, given their potential to incrementally reduce the presently large data-related barriers to ISR assessment. However, analyst interviews also revealed a longer-term need for faster, seamless, more unfettered access to ISR data. This will require an “overhaul” in the approach to ISR data management. This is something that has been acknowledged within ISR community for some years and was also affirmed by the PAF project team in site visit observations of the systems, databases, and practices currently in use within the USAF ISR enterprise. This will be a longer-term process, as it involves changes to legacy architecture and approaches, which cannot necessarily be changed overnight (indeed, this could take years or a few decades to complete).

We begin our discussion with some suggestions for what can be done to achieve limited improvements in the coming months or a few years. Based on our interviews and observations, the easiest way for USAF to improve the quality and accessibility of data for ISR assessment (and perhaps for other purposes as well) *right now* is to expand and improve the UNICORN database. As we have discussed, this database is already serving as a useful repository of ISR information and is owned, operated, and maintained by USAF. Although there might be other options for USAF to consider, UNICORN is the primary enterprise-wide database for a large portion (but not all) USAF ISR platform sorties and thus could serve as one useful starting point in developing near-term solutions for improving ISR data accessibility and accuracy.⁴

UNICORN was designed to help manage PED missions across the DCGS.⁵ It includes PED products and information about the accompanying sorties that collected on the target(s), and

⁴ Motion and still imagery collection in support of conventional operations is most consistently reported in detail on UNICORN.

⁵ Ashley Hodges, “UNICORN 101,” 480th Intelligence Wing, briefing slides, April 29, 2008; interview with Benjamin Dunlap, 480th ISR Wing, June 7, 2017.

therefore is a natural place for assessment analysts to turn to (though it was not designed entirely for the purposes of ISR assessment). Airmen manually enter into UNICORN information about ISR sorties and the resulting collection and exploitation. UNICORN draws information from several other databases, including those which contain information about requirements and operational outcomes (notional examples of “outcomes” include whether troops were surprised by enemy forces and the amount of illegal substances captured during an ISR-supported operation). Although this connectivity to other information sources allows UNICORN to provide better context for ISR sorties and PED activities, it also generates additional data quality problems because these external databases also rely heavily on manual entry.

The way most of the data is tracked in UNICORN is often described as “bean counting,” since it primarily involves recording facts that can be explicitly measured while incorporating only a limited amount of information related to the relevance of ISR collections and proxy indicators for decisionmaking quality. Basic knowable facts about what assets flew and what happened during the collection are mostly accounted for (notable exceptions are flight paths and aircraft tail identification, which are not in UNICORN). These facts help to paint a partial picture of asset employment (missing flight information is still a problem), but do less to elucidate the utility of collection: were intelligence gaps closed, were troops on the ground kept safer, or were operations able to run more smoothly? In other words, whereas UNICORN data could help assessments analysts make substantial headway to detail *costs*, there is less information about *benefits*. In addition, systematically compiling information using UNICORN to even do an assessment of costs can be challenging due to lack of data about some ISR sortie characteristics and the manual labor sometimes needed to construct spreadsheets tailored to specific assessments needs (i.e., information contained within UNICORN cannot always be easily manipulated and exported—for example, to compare different types of sorties or similar sorties on different days).

There are some indicators available within UNICORN that could be used to examine benefits or the lack thereof (e.g., killed-in-action or wounded-in-action statistics or information about ordinances dropped), but more could be added to support analysis of decisionmaking quality, as discussed earlier in this report. In addition, post-mission summaries, mission reports, and mission outcome statistics provide a window into what happened during collection. Did an asset spot targets, warn planners about the safety of a geographical area for a future mission, alert troops to retreat from a dangerous situation, or spot intruders trying to penetrate a base? Some of this type of information is stored in the writeups about the ISR sorties and PED activities. Unfortunately, these records vary in specificity and quality and are not easily scrapable.⁶ Furthermore, when UNICORN data is exported, it does not embed the post-mission summaries, images, or histories of targets that have been previously the subject of collection. Thus, in the current data

⁶ *Data scraping* is a technique by which a computer program extracts data from human-readable output coming from another program (in this case, the output is coming from programs like Microsoft Word or Adobe).

environment, it is extremely difficult to track how pieces of information fit together and why some operations go better than others.

Even if data within UNICORN were used to correlate ISR sorties with outcomes, it is impossible to judge if a positive correlation means ISR assets were being smartly allocated and employed or if it was mostly because of luck (or somewhere in between) without knowing the intent and reason that ISR assets were deployed in the first place. Before January 2017, the best way to ascertain this intent was by reading through mission-associated EEIs and pairing them with the associated commander's intent at the time. Like many other human-composed text files, EEIs have huge variations in quality and are not easy to search through (nor are they necessarily always the best reflection of why an ISR sortie was conducted). Commander's intent and target prioritization are generally identifiable, but this is not neatly tracked within UNICORN either, nor are these details always sufficiently descriptive to help deduce the purpose of an ISR sortie.

As of January 2017, however, the 480th ISR Wing made a first pass at codifying the intent, or purpose, of an ISR sortie within UNICORN, by having individual PED analysts attach an ISR role to each record in the database. These ISR roles are similar, but distinct from, the four first-order roles developed by PAF.⁷ This is an important step forward, but analyst interviews, DCGS site observations, and experimental UNICORN data analysis performed by our team revealed two limitations of this advance within UNICORN. First, analysts may interpret the role associated with an ISR sortie differently (in addition, a sortie may have more than one role, but this is not tracked in UNICORN). Second, a more granular understanding of ISR collection purpose (e.g., PAF third-, fourth-, and fifth-order roles) is necessary for conducting a range of assessments. For example, an AOC or Task Force interested in seeing how ISR assets are being used in the context of an operation might not learn enough to support allocation decisions if they find that all ISR is "supporting operations" (equivalent to a PAF first-order ISR role). Receiving an assessment that strike-capable ISR platforms are accounting for civilians (a PAF third-order ISR role) whereas reconnaissance-only aircraft are searching for adversaries ahead of a moving force (a PAF third-order role) might, however, spur helpful discussions about which aircraft should be allocated towards what purpose.

An alternative approach for including ISR roles in UNICORN would be for ISR requestors to select one or more third- or higher-order ISR role as part of the ISR requirements generation process. This could be made simple and consistent, for example through the use of an automated querying system that inputs a user's initial keystrokes to return suggested roles (rather like online search engine suggestions). Another way of doing this might involve providing collection requirements managers with automated tools to tag ISR sorties before they are put onto the ATO.

Regardless of the exact mechanism used to identify an ISR sortie's intended role(s) at a fairly granular level before the aircraft takes off, the result could help prevent PED analysts from having to continue second-guessing the purpose(s) of a particular collection. This could lead to

⁷ These were among those that the research team reviewed in constructing the PAF ISR taxonomy.

potentially improved cataloging of ISR sorties as well as enabling PED analysts to focus more on other tasks they have to do. Of course, PED analysts should be able to amend the roles input into UNICORN if the use of an asset changes post-takeoff.

Through interviews with UNICORN users and maintainers, we learned that doing basic manipulations to the database—including such things as creating drop-down menus and autocorrecting spelling—would take relatively little time and effort and could improve data usability for assessments. Scrubbing the data inside of UNICORN could also, at the very least, remove recorded data from sorties that contain logical fallacies, leaving a cleaner dataset for analysis. Even with such changes, there will still be the problems of missing sortie data and poor-quality data imported from other sources.

Actually improving the quality of original data is a much more difficult task. The Office of the Under Secretary of Defense for Intelligence’s Project ISR Data Enrichment and Aggregation (IDEA) sought to work on exactly this problem. Working with USPACOM, Project IDEA developed algorithms to mine data from various post-mission products and other data files as well as databases used by the Joint Community, including UNICORN. The goal was to standardize, deconflict, and enrich data. In the case of conflicting data (e.g., more than one takeoff time for a particular ISR sortie), the most authoritative data source (closest to the original source) was used. Using Project IDEA algorithms, the scrubbed data are combined into one master source and uploaded to a dashboard for assessment purposes.

Initial employment of Project IDEA appears successful, saving man-hours, detecting challenges in the USPACOM ATO process, and allowing for more accurate and repeatable assessment of basic questions. Analysts report that they can immediately assess what assets flew and determine the reason if an ISR sortie did not occur as planned. Analysts also suggest that they have more awareness of resulting intelligence products, which targets were observed, and which sorties experienced foreign intercepts. Through interviews, we found that there seems to be interest in extending Project IDEA to other COCOMs as well.

Long-Term Mitigation Approaches

Whereas the short-term mitigation approaches discussed in the previous section will better facilitate the answering of basic assessment questions (e.g., did the sorties fly? what products resulted?), long-term improvements can enable USAF ISR assessment analysts to address more advanced questions that require linking of ISR benefits and costs (e.g., have ISR investments benefitted counterterrorism operations? is the ISR asset needed more in one COCOM or another?) by taking more full advantage of modern data transmission and manipulation technologies. Our interviews and observations suggest that a new USAF ISR database could be needed in order to most efficiently and effectively integrate the upgrades outlined below that we believe are needed to ultimately enable more depth and breadth in ISR assessments. However, evolving UNICORN should also be weighed as an option, bearing in mind that there may be

fundamental limitations in its design because it was not originally built to house and manipulate data for assessments. It might also be possible to simply create a new search interface that interacts with numerous existing databases, though this would likely not be able to fully improve data quality and some other issues that are currently hindering ISR assessments. The following long-term improvements, summarized in greater detail below, should also be coordinated with Joint and IC efforts to maximize utility:

- Increase machine-to-machine data transfer to reduce incidence of human error when inputting data.
- Tag data collection to enable more effective data search and analysis.
- Track how well collection helps answer key intelligence questions.
- Enhance the amount and diversity of feedback given from users to analysts.
- Enable fast, secure search for analysts to easily identify and manipulate quality data for assessments.

Machine-to-Machine Communication and Data Transfer

Many of the data fields found in UNICORN and other ISR-related databases and information sources record objective, knowable facts. Enabling systems to directly transmit such information—or, when necessary, automatically transmit manually input data after it has been recorded *the first time*—to an ISR database will help minimize problems associated with manual transcription and retranscription. Further, it will help eliminate the issue of figuring out who is responsible for “clicking the buttons” or actually recording the data.

Systems onboard aircraft, like the Geographic Positioning System (GPS) or Internal Navigation System, record data about a sortie’s geographic location, flight path, important coordinates (e.g., an aircraft entering its “box”), time, sensors used, weapons detonated, ISR asset type, and other facts that could be useful to ISR assessment analysts (and others). These data could be automatically transmitted via telemetry or analogs to a central ISR database that people who need access to the data can query.⁸ This would prevent the information from becoming degraded from manual recording, which is often the current process for inputting such facts (if at all) into databases the USAF ISR enterprise employs.

Other factual data cannot be recorded by onboard systems, because they (currently) require a very basic level of human judgement. These include information such as the target number, target type, and mission number. Currently, many of these data values are assigned more than once in different databases by different analysts, which can result in unnecessary data mismatches and logical errors (for example, if one analyst types in the wrong mission number). Though UNICORN does link to other databases and information sources, it still relies on a manual upload of exported information while it automatically ports data back into some other systems. Until USAF reaches a point that systems can talk to each other and transmit information

⁸ These data are typically being transmitted, just not to a database easily accessible to ISR assessment analysts.

without human intervention, data that is intrinsically human-generated will lose fidelity as it is transferred between systems.

Machine-to-machine communication for the purposes of data recording is not a novel concept, nor does it require great technological development. For example, Apple's iPhone has a GPS, accelerometer, gyroscope, and magnetometer working together to automatically track and save data about a user's daily movement. iPhone users can thus record their running routes or the number of steps taken and transmit that information to an application (app) for human consumption. The information stored in apps can also be uploaded to a cloud account (iCloud) so that locally stored data can be restored were something bad to happen to the phone. Conveniently, multiple cloud accounts are not necessary to save other sensor-generated data, such as photos, or such human-generated data as emails. Data contained in all Apple apps plug into a single cloud account, and there is no manual labor involved with data synchronization (syncing). It is entirely done by machine-to-machine communication.

Note that any transition from manual data entry to machine generated data will likely need to be accompanied by additional efforts to standardize data formats, including enhancing compatibility with Joint and IC partners. In itself, machine-to-machine transmission of data should make it easier to enable this standardization.

Tagging Collection to Produce Useful Metadata

Returning to the iPhone example, user photos are geo-tagged and can be tagged by more advanced features like people's names via facial image recognition technology. Instagram allows users to tag photos with whatever hashtag they want. The metadata associated with iCloud or Instagram photo repositories is desirable to users, as it allows thousands of photos to be rapidly searched and clustered. A user who takes a yearly trip to a vacation destination can trivially see how the environment, activities, or people might have changed over time. It is not a stretch to posit that these types of capabilities—enabled by effective use of data tagging—would be useful to USAF for ISR assessments.

Generally speaking, tagging data allows for relationships between objects, networks, or places that might superficially appear dissimilar to be detected. Indeed, there are many efforts within USAF today aimed at tagging data usefully. There are also nascent projects to automatically analyze images and tag them with key findings. However, coordinating these efforts within the service, and with IC and joint partners will be critical for ensuring that information can be shared and compared.

The USAF ISR enterprise could consider continuing efforts to follow the IC's lead on data tagging, which would reap two benefits. First, this would help ensure consistency with the IC. Second, the IC, in the course of working through its own data management challenges, has come up with some useful insights about tagging data and how to usefully analyze it once tagged.

Much like iCloud or Instagram, the IC stores some types of information for on-demand use. IC data standards are continually evolving, but always emphasize enabling future exploitation.

Telemetry and other machine-to-machine communication tends to capture the physical characteristics of collection, while data content captures observations made by and through these physical characteristics. Efforts to generate content using consistent nomenclature or taxonomy enables automation, parsing, and exploitation by successor systems and techniques. All of this enables real-time comparisons of current and historical activity by metadata tags (e.g., geographical location). Baseline profiles of targets, areas of the world, and associated activity can be easily found on-demand. Databases are used to organize data and facilitate easy retrieval of old observations. Information garnered from collection is linked to specific objects and places of interest and the ISR missions that collected on them. This allows for relatively facile characterization of what ISR activity and employment lead to what outcomes.

USAF could adopt data tagging processes like those of the IC for all collection. Having historical video, imagery, and other forms of collection available on demand and searchable would greatly facilitate assessments and analyses in general. For this to happen, mission summaries and other ISR products must be decipherable enough to be used for subsequent analysis and planning. Ideally, tagging would be accomplished using

- consistent structure and terminology
- clear standards for characterizing observations
- descriptions of context, intent and collection objectives
- associations with networks entities or other known targets.

Importantly, the ISR database itself, or systems that it pulls from, are crucial for processing and tagging data for metadata production.

Tracking the Closing of Intelligence Gaps

Enhancing data quality and tagging will do much to enable analysis of the *costs* portion of the cost-benefit framework described earlier (along with aspects of benefits related to products, or outcomes such as strikes or observations tagged to a particular target). However, we have also argued that understanding *benefits* at the macro level of intelligence gap closure more clearly is also important for assessing ISR's value.

Correlating intelligence successes and failures with ISR employment fundamentally requires tracking how particular initiatives are progressing, if intelligence gaps are closing, and if intelligence problems, no matter how minute, are being solved. Currently, some basic metrics of success are being recorded, but they are extremely difficult to link to the broader progress of an intelligence operation. Individuals working on an intelligence operation keep their own notes and contextualize the output of different ISR missions. Centralizing these activities would enable analysts across an enterprise (ideally the USAF ISR community or IC) to observe progress across different efforts and better contextualize their own efforts. This would also enable coordination to maximize information flow and minimize duplication of efforts, both of which are key challenges for the USAF ISR enterprise today.

Other organizations within the IC delay their mission summaries for a short period of time (hours to days) to allow for coordination and information flow between related but separate missions. Entry into IC centralized databases is delayed somewhat longer (days) to gauge the usefulness and additive effects of intelligence analysis. This allows the database to track the utility of different collections and forms the basis of ongoing intelligence narratives. A more thorough discussion of exactly how these tracked outputs should be used to answer real assessment questions is found in Chapter 7, so we leave this section brief. Acknowledging that there already exist a large number of spreadsheets, wikis, one-off “databases,” and other valuable information for assessments, the most important step forward is for USAF to *centrally* store, timestamp, and organize intelligence analysis so that people can retrospectively understand the level of progress made at any given time, know the results of other potentially related collections, and be able to correlate the outcomes with ISR employment.

The National Geospatial-Intelligence Agency (NGA) employs a process called structured observation management to help analysts make connections between different observations.⁹ This is facilitated by an associated process that ingests, structures, and organizes new data entries to enable consistent formatting and visibility of data. This is a model which USAF is already aware of and working to emulate, and could ultimately help track progress towards closing intelligence gaps and facilitating ISR assessments.

Enhancing Feedback

Although there are exceptions, it can be very challenging to determine whether particular collection and PED activities met the requirements and other needs of one or more end users. Once collection has happened, PED and other analysts will exploit and interpret results, producing information they believe is relevant to the need. However, it is difficult to get consistent feedback from users as to whether the results were useful. This creates clear challenges for assessing the results of particular ISR activities, especially with respect to assessment question 1.

Ideally, long-term improvements in the data environment for assessment will also be able to support additional new ways of getting feedback from end users using the following mechanisms:

- tracking number of views of PED products
- establishing a simple, consistent rating scale for products or ISR missions
- citing information used to build fusion reports.

The first mechanism would be a passive way to see how many views something has gotten, similar to content on YouTube, for example. Not every viewer will find what they needed, but

⁹ See Mark Munsell, “A Focus on Data Management; Transforming the Way We Do Business,” *CIO Review*, May 30, 2017.

this should provide a broad ability to see what has been popular for users and perhaps infer from this what is useful. This would provide the additional benefit of flagging products that are being viewed very little or not at all, which can prompt questions to the intended users about whether they still have a need for this type of information or whether they know where to find the information.

The second mechanism would involve an automated survey and require a limited amount of activity from users to either “like” or “dislike” a product or mission, or rate something from one to five. This could be a requirement for leaving a particular product page. Although some users might initially dislike having an additional step in viewing products, this is probably something most can get used to fairly quickly. These are successful feedback strategies employed by Facebook, Uber, and other online entities and would represent a mere fraction of the time required to access a product in the first place.

Finally, the idea of using more detailed citations for reports and products that build upon information from one or more original sources appears to be gaining traction in some parts of the enterprise. Providing guidance about what information to include, as well as where and when to include it, will help to standardize approaches to citations. Automating citations will not only help with this standardization, it will also help save airmen’s time. There will be numerous benefits once citations are implemented more widely in a standardized fashion. For assessments in particular, there will be two main improvements: the ability to track citation metrics and an additional mechanism whereby products with similar themes can be tracked. This latter idea would help support initiatives to build relational databases and OBP, which in turn can help facilitate better understanding of the “additiveness” of new collections, as detailed in Chapter 3.

Database Construction

We have already discussed four types of initiatives above that will help improve ISR mission data capture and contextualization. Now, we turn to the information database or repository itself. How USAF should ultimately shape a future database (or other effort to support data management and analysis) is a big decision, all aspects of which cannot be examined and weighed in this report more broadly focused on ISR assessments. However, any effort should capture and enable linkages between data that can help describe both costs and benefits associated with ISR missions. Any effort must also be capable of integrating with other data repositories and databases throughout the USAF, IC, and Joint Community, interfacing as seamlessly as possible to maximize information flow. Examples of necessary USAF interfaces include the ATO planning cycle, target prioritization, and aircraft maintenance schedules. There are many important features to consider when creating and implementing any new information repository, database, or other effort to support data management and analysis. While we do not aspire to be comprehensive in our discussion, we discuss some key features in general terms below.

“Open Architecture”

Favoring openness in database construction and linkages, as opposed to having information siloed in different databases and systems that cannot interact, will help facilitate better ISR assessments because data will become more easily accessible. Any long-term solution for enhancing USAF ISR mission data collection and storage should allow different user groups to customize and extend capabilities to meet their needs and requirements. It should also simplify connecting existing systems currently siloed in different units and organizations. It should automatically pull and format information from other relevant databases. For example, flight paths that are already transmitted to existing systems, but this information should then either be uploaded to the central database or linked to so that a user can pull the information easily.¹⁰ Any long-term solution must allow for querying, creating, modifying, and deleting data.

Storage Requirements

USAF intelligence operations produce large amounts of data. Any future database solution will need to take this into account, though the cost and diversity of options for storage—commercially, at least—are becoming more diverse and are relatively low in cost (as might be expected when there are diverse options). A rough notional estimate of storage requirements for 60 Reaper orbits serviced 24 hours a day requiring 80 gigabyte of storage per hour of uncompressed high definition video is about 115 terabytes per day or 42,000 terabytes per year. This is a large amount of data, but the cost of storing it does not necessarily have to be high. For example, Amazon storage costs \$0.022 per gigabyte . For the notional Reaper example above, storing a year’s worth of video would be less than a million dollars per year. Of course, there will be other data that needs to be stored, but in an absolute sense, storage is quite inexpensive.

Therefore, imposing any upfront maximum on data storage capacity for any future database might not be wise. Data should be managed so that they are easily accessible from storage until they reach an age threshold or they are discovered to be flawed in some way once an intelligence analysis is complete; this data can then be compressed and migrated to storage that is infrequently accessed (reducing costs per gigabyte by nearly half). Eventually, aged data should be purged following the authority under which the data was collected. Data related to enduring intelligence challenges, or some fraction of most important intelligence related to various operations, could be stored indefinitely or as needed.

There is also the question of where data should be stored. Should USAF maintain its own servers? Should resources be shared between the services? Should the services share with the IC? Perhaps storage solutions should be outsourced entirely to Silicon Valley. A thorough analysis of these alternatives is out of our research scope, but we mention it as it is an important question to resolve early in the design of a future ISR mission database concept.

¹⁰ Today, users interested in retrieving flight paths have to be aware of exactly where to look, how to access the information, and then how use it.

Data Formatting

The future ISR mission data solution should be very flexible in terms of data formatting. Based on PAF observations at different USAF ISR enterprise sites, there should be capability to store objects ranging from documents to text files to high-definition video, regardless of the format. It should also be possible to embed multi-intelligence collections in order to facilitate cross-checking information across intelligence types. Further, the possibility of new file formats emerging over time should be taken into consideration.

The extent to which USAF (or possibly DoD) should enforce data standards on information stored in databases will depend on how quickly and how well different DoD and intelligence agencies conform to any standards. Currently, even when standards do exist—for example, NGA has created geospatial intelligence (GEOINT) standards consisting of technical specifications and other precise criteria to ensure that GEOINT products are suitable for analysis and visual representation of Earth—USAF may not always be in sync with them.¹¹ Adhering to standards becomes increasingly important for interoperability as intelligence analysis becomes better coordinated. Indeed, agreeing upon and ensuring adherence to data standards may help facilitate better coordination of intelligence operations and analyses, in addition to more effective assessments.

Data Tagging

As discussed previously, data tagging is of great importance. The future ISR mission data solution should have the capability to tag and attach information. This includes both automatically generated mission related metadata (such as weather conditions) or user-defined tags (e.g., a hashtag) to data and files. Tools should be available to parse specific target-related data in the database, whether it is gathered from image recognition technology, cyber targets (e.g., an Internet Protocol address or a port), or human-defined information. Two additional particularly important tags are timestamps and classification levels.

Data Querying

It is important to determine how any new or updated database will be queried. An example of what could be used is the structured query language (SQL). SQL is an international standard language for creating, processing, and querying databases to retrieve, format, report, insert, delete, and modify data. Regardless of which query language is chosen, the future database must enable users to quickly and easily search and analyze the totality of data over time for longitudinal studies, and NRT data to inform dynamic military situations relying on ISR.

We also conducted an examination of the benefits and problems of search algorithms utilized by PageRank and other search engines that scan the internet to automatically populate search results (see Appendix C). A major criticism of search engines is that they enable savvy web

¹¹ NGA, “Geoint Standards,” website, undated.

developers to “game the system” by identifying and exploiting the metrics being employed. Artificially inflating a website’s rank reportedly has become more difficult as the algorithms and processes have been refined over time, yet the vulnerability of search engine algorithms can undermine confidence in the system or create a poor user experience. The important take-home message here is that future USAF intelligence search algorithms need to be carefully designed to preclude people from “gaming the system.” For example, perhaps in the future an analyst will be able to query the database to look for all information related to a blue house on a certain street in Afghanistan. USAF will want to avoid the situation where human intelligence (HUMINT) data always comes up first, while RPA data comes up last, because a search algorithm was tweaked to bias HUMINT reports (and thus potentially make HUMINT look more important).

Apps and Algorithms for Analysis

An open architecture will enable a wide variety of applications and algorithms to be designed for accessing and manipulating data. Importantly, not all future user needs can necessarily be anticipated in the near term, so flexibility is important. Some examples of useful functions applications could be designed to provide include curating and aggregating data to discover hidden relationships between targets, exploiting NRT information to benefit operations, and analyzing what types of ISR operations were more or less successful over time.

Here, we examine battle damage assessment for a notional example of how applications accessing future ISR mission data might be used to support real-time operations. Consider a Reaper involved in targeting a terrorist compound. To minimize ordinance dropped (perhaps to reduce the risk of unintended consequences and to save munitions), knowledge of the success (or failure) of each strike must be generated through analysis of NRT intelligence. An analyst watching streaming video of a strike can, in this theoretical example, develop a baseline of what the compound looked like in the past by querying relevant archived intelligence products using an application designed for this purpose. At the same time, a second analyst could examine SIGINT to examine activity originating at the compound since the strike took place. In this example, the SIGINT analyst might notice the sudden use of a communication device near the compound that could be associated with another HVI. Details about that HVI could theoretically be queried through the use of an application to provide context for the current intelligence hits, helping the analysts decide what information to pass along to those involved in making strike decisions, and with what level of urgency.

Over time, computers can be programmed to automatically execute routine analyses and assessments. In the more distant future, one can imagine deep-learning algorithms that would allow computers to find correlations that grant significant insight into ISR’s value (and into answering key intelligence questions more generally) that humans might have missed without high visibility into operations and deep experience. With a large fraction of the USAF ISR

analytic workforce changing station every few years, this may be an area in which an algorithm can aid analysts with less experience looking at a particular intelligence problem set.¹²

User Access and Security

Finally, the issue of user access must be considered. Different types of access restrictions must be enforced—for classification reasons, to protect unclassified but sensitive information, and to minimize risks posed by those that might exhibit malicious intent, to name a few. All organizations, including USAF, face the challenge of granting enough users access to do their jobs well, while not permitting so much transparency so as to incur undue risk to missions. USAF will have to continue to weigh this problem for the future ISR mission data environment, which in some ways may pose greater security challenges than those experienced today if information is less siloed across different databases and other sources than it is today. The full spectrum of classified information needs to be accommodated, appropriately barring people without the proper security clearance and need-to-know from retrieving some data. Acquiring the most robust cyber and information security solutions for the database is crucial, and it is also important that security be regularly reviewed and maintained to allow for new solutions to be developed as potential threats evolve.

Other Acquisition and Use Considerations

If USAF decides to acquire a new database or other technology, a rapid acquisition approach should be considered in order to ensure that the system(s) do not lag behind new developments before they can even be fully fielded. Funding is always a major consideration, but the technical aspects of a new or updated approach can be drawn from a strong baseline of existing efforts within USAF, broader IC and Joint Communities, and private industry. Part of this will also include agreeing on data and data file format standards, as discussed earlier. It is also imperative that systems integration considerations allow for (1) seamless interfacing between the database and existing systems and data sources across USAF and (2) telemetry and other forms of machine-to-machine communication.

Finally, training additional airmen in data science sooner rather than later will help to build a critical mass of knowledge of the field. Technological advances will not be helpful for ISR assessments or anything else if there aren't sufficient airmen that are familiar with manipulating, analyzing, and interpreting data. Like any skill, data science requires some training, and our analyst interviews suggest that this is not necessarily something that can just be learned "on the job."

¹² There are a number of initiatives to develop these types of algorithms. Key to their success, among other factors, are training personnel broadly enough to be useful and ensuring that analysts are only using these algorithms to the level at which they are comfortable trusting a computer program.

Recommended Next Steps for USAF

This chapter has discussed a number of challenges related to databases and other information sources used for USAF airborne ISR assessments. We also broadly outlined both short-term and long-term mitigation approaches which were directly informed by project team time spent using operational databases, interviews with analysts and database maintainers, and the examples of assessment practices used in other industries and organizations described in Appendix C. We summarize these mitigation approaches here.

In the short term, it makes sense for USAF to pursue some incremental improvements to the UNICORN database. Although this database was not designed for assessment purposes, nor does it contain all of USAF's ISR missions (let alone potentially relevant non-USAF ISR activities), it does currently enable limited-scale ISR assessments. There are some ways to grow UNICORN's capabilities with relatively little scale and scope of effort. First, implementing granular ISR roles (equivalent in detail to PAF's third-, fourth-, or fifth-order roles) will help better define and track how ISR is being used. Ideally, this information would be provided by ISR requestors and then edited by collection managers or PED analysts if and when missions evolve. This would prevent PED analysts from having to guess the purpose of a mission from (many times) vaguely articulated EEIs.

Secondly, creating and implementing more basic algorithms to identify data quality problems (even if these are not resolved, flagging them can still save analysts time) and enhance USAF ISR analyst ability to manipulate what data are in UNICORN. More algorithms can also be designed to pull in and format other available data that could help analysts better document ISR costs and benefits. Project IDEA is one example of such an initiative that is helping COCOMs (USPACOM and potentially expanding to others) conduct better ISR assessments. Generally speaking, USAF should continue keeping an eye on work being done in COCOMs and the IC to improve ISR assessments.

In the long term, USAF should decide on an enduring, effective solution that will enable its ISR assessment analysts and others within the ISR community to rapidly access and manipulate large varieties and quantities of data to inform decisionmaking within the enterprise. This will mean ensuring that data standards, for example, are consistent as necessary with partners in the IC and Joint Communities. The long-term solution may take the form of a new database, hefty changes to UNICORN, the construction of a new search algorithm and user interface that accesses dozens of other databases, or some other option. Regardless of form, we have discussed a number of types of improvements required as part of this long-term solution, which include:

- increasing machine-to-machine data transfer to reduce incidence of human error when inputting data
- tagging data collection to enable more effective data search and analysis
- tracking how well collection helps answer key intelligence questions
- enhancing the amount and diversity of feedback given from users to analysts

- including several important considerations in future data management solutions:
 - favoring openness in database construction and linkages to enable integration of diverse data sources and applications
 - planning for a large amount of storage with the ability to review and eliminate data when no longer needed; this will be critical for enabling analysts to examine historical data and build on existing intelligence threads
 - creating and enforcing data standards to enable data merging and analysis
 - enabling data tagging and querying to allow analysts to retrieve information needed
 - creating applications and algorithms to support analysis, especially in situations where this is needed quickly
 - ensuring user access and security so that those that are appropriately cleared and have a need-to-know can access the right information
 - providing data science training when necessary to help airmen make full use of data for assessments
 - using rapid acquisition processes to ensure that new materiel and processes are not obsolete before they can be fielded.

5. Step 4: Analyze the Value of ISR Using Real-World Data

The previous two chapters have focused on what fundamentally amounts to different aspects of cataloging and describing information for the purposes of USAF airborne ISR assessments. Chapter 3 detailed ways of measuring benefits and costs of USAF airborne ISR. Chapter 4 provided insights on specific data and database issues that require resolution in order to measure benefits and costs related to USAF airborne as this report describes them.

The next step is to perform the actual analysis to assess USAF airborne ISR. Just as the indicators and datasets will vary depending on the decisionmaking context, so too there is a choice of analytical methods that may be employed, singly or together, in a given USAF ISR assessment. We do not specify which methods should be used because these will necessarily need to vary depending on the type of ISR-related question being asked and how the outcomes will be employed. Instead, we provide two notional examples relevant to USAF airborne ISR that each illustrate a type of analysis that an air component might be interested in undertaking if real-world (i.e., not simulated or inferred) data were readily available to support it. We discuss analyses based on simulation in the next chapter. The method used here is the cost-benefit framework introduced in Chapters 2 and 3, and we assume that accurate real-world data are available via database improvements such as those described in Chapter 4. Note that the lack of such data availability for USAF airborne ISR under many circumstances would make either of the assessment examples we are about to describe extremely challenging to conduct within the present data environment. The context for these examples is based on an examination of IC intelligence assessment practices (see Appendix C) and expert interviews with analysts working within or supporting missions in the USCENTCOM and USPACOM theaters.

Here again, the analysis will proceed differently depending on the type of decision it is meant to inform—though always grounded in the cost-benefit approach outlined in preceding chapters. This chapter presents two largely qualitative examples using notional operational data to illustrate how different assessments might be used to inform different ISR resourcing and employment decisions. The first examines a tactical-level resourcing and employment decision that a senior intelligence duty officer might face: the dynamic retasking of ISR resources to respond to changes in operational conditions. This decision requires an assessment of the opportunity costs associated with different resourcing options. The second example highlights a strategic-level resourcing and employment decision that an air component might seek to inform by communicating with their COCOMs that might, in turn, communicate with Joint Staff members: a mid-year shift of an airborne ISR collection asset from one COCOM to another because of an emerging need. Both examples highlight time-sensitive situations in which ad-hoc ISR assessments must be performed. In many cases, these situations may be even more

challenging than routine assessments because of generally shorter lead times for developing effective arguments, which is why we selected them.

In both examples, we assume that there is a well-defined set of indicators to answer the assessment questions discussed in Chapter 3 and a comprehensive, accurate, and retrievable body of data from relevant historical ISR operations that analysts can quickly draw upon for their analyses, perhaps facilitated by automation.

Reallocating ISR During Dynamic HVI Operations

The first example shows how the ISR assessment approach can be used to support a senior intelligence duty officer in an AOC (air component) who is faced with the decision of whether and how to reallocate limited ISR resources among multiple priorities during HVI operations spanning the course of several hours.

In this notional scenario, we assume an operation in which Blue forces seek to find, fix, track, target, and ultimately engage several HVIs active in a given locale. Figure 5.1 describes the three HVIs and their relative importance as determined by the commander. This prioritized list provides the kind of “status and intent” that should be available to planners, collection managers, and mission teams. The collection strategy is a “network-centric” deliberate targeting of HVIs via linked entities in their network. We assume the organization prosecuting these targets can sustain five airborne ISR orbits (24-7 aircraft and PED presence) in the area of operations and can readily move aircraft between targets within the theater of operations. The example begins mid-mission, early in the evening (“Hour 0” for purposes of this example) and spans eight hours, so the decisionmaking that ensues focuses on only the dynamic reallocation of resources.¹

¹ The classified Appendix B includes a more comprehensive scenario with both phases of ISR asset allocation: the planned allocation of resources and the dynamic reallocation of resources.

Figure 5.1. Notional List of Prioritized Targets

Priority 1: HVI #2 is a terrorist planner, recruiter, and media personality.

Status: The HVI has been identified living at his residence with his family. We have been monitoring them for six months, but due to the HVI's tradecraft, we have observed him only three times. Since the last sighting two days ago, we have maintained custody of the HVI and his family at the residence.

Intent: Track movements at the residence to enable a strike against the HVI once he has separated from his four children. The HVI's wife has been determined to be a combatant, so she can be included in the target picture. Do not lose custody of the HVI since that is a precondition for striking.

Priority 2: OBJ Last Resort is a terrorist training camp with two high-value trainers.

Status: The camp was located one month ago. The trainers have been identified, and there are 6-12 trainees at any given time. Training occurs almost every day, but the camp occasionally hosts visitors who do not participate in training. A servant or "tea boy" typically is present during the daytime. A vehicle also is present, which departed along with half of the trainees last week.

Intent: Track activities at the camp to enable a strike against the trainees and both trainers. The "tea boy" has been determined to be a non-combatant, so do not strike the camp while he is on target. Do not strike the camp if anyone on target has not participated in training, such as occasional visitors. Be poised to respond if the trainers or trainees load up the vehicle for departure.

Priority 3: This courier is a developmental lead to HVI #8.

Status: We have identified the residence of a known courier for HVI #8. Intelligence indicates he meets in person with the HVI and other network associates; however, we have not observed any such meetings over two weeks of occasional coverage, during which the courier frequently departed his residence in his personal vehicle to run errands at public places.

Intent: Find leads to HVI #8 by following the courier to meetings or hand-offs.

The overall operational objective is to strike these targets, subject to the rules of engagement. To do this, the operational commander (or others involved with the operation) must make a series of decisions that move along the find, fix, target, track, engage, and assess (F2T2EA) phases of the dynamic targeting cycle. The overall ISR objective is to provide enough information to enable the commander to decide whether and when to strike the targets. Additionally, ISR supports the commander's objectives and decisionmaking by filling information needs. For example, it is necessary to confirm that a person of interest is the sought-after HVI before an operation can move from the "fix" to the "track" phase. At any given F2T2EA phase, certain pieces of information must be filled in before the commander or others can make a decision whether or not to move to the next phase, and hence come closer to meeting the operational objective. From the perspective of ISR value, a decision either way (i.e., to move to the next phase or not) indicates an ISR benefit, in that the commander was able to make that decision. In this sense, it is possible to indicate the extent to which ISR has enabled the

commander to make a decision by looking at how much of the required information have been provided in the current F2T2EA phase. Likewise, we can assess how much ISR has furthered the pursuit of operational objectives by looking at how the operation has progressed along the F2T2EA process.

The rate at which those information needs are filled is not entirely within Blue's control. For example, the HVI either is or is not accompanied by his children, and so the information that he is alone (and therefore potentially eligible to strike under the rules of engagement) may not come for some time. However, the time required to fill information needs can sometimes be estimated based on an understanding of previous HVI activities (e.g., through IPoE), both in the current operation thus far and, perhaps, in historical operations. For example, the analyst may know that the HVI does not generally leave his family at night, and so the information that he is alone is unlikely to come for at least several hours. This type of insight can be valuable in an environment where ISR resources are limited, as we discuss further below.

The ISR collection manager is responsible for tasking ISR resources to support the commander's progress toward operational objectives. The collection manager will need to determine what information is needed to enable the next relevant decision (e.g., to move from one targeting stage to another), what ISR functions are needed to collect that information, and what ISR resources (e.g., sorties, sensors, manpower) are needed and for how long to carry out those functions. An assessment of past operations can directly inform these decisions; assuming historical data is available in a consistent form, such as proposed in preceding chapters, the analyst can understand the benefits likely to be achieved (i.e., timely support to enable the commander's decisionmaking) for a given cost (i.e., the type and level of ISR effort)—in other words, ISR value.

Suppose, for the purpose of this illustration, that the collection manager has made the initial allocation decisions, using input from the analyst's assessment of historical data and perhaps from his own experience. Moreover, a certain number of objectives have been prioritized for ISR collection within the theater for some days, hours, or months. At the moment we open the scenario ("Hour 0"), we see that enough information has been provided to enable the commander or other decisionmaker (perhaps the collection manager, under some circumstances) to move operations to the "track" phase for HVI 2 (see Table 5.1). In this notional example, the commander needs to confirm that the HVI is not with his family in order to move to the "target" phase. This requires ISR to keep apprised of where the HVI and his family members are at all times. The collection manager, again drawing upon the analyst's assessment of similar past operations as well as rules of engagement, standard operating procedures, and intelligence tradecraft, determines that two ISR functions must be carried out to meet this information need: (1) maintain custody of the primary target, and (2) account for all personnel previously observed with the target. Moreover, the collection manager determines that these tasks require two RPA and two PED analysts and that it is likely to take about x hours to gather the information needed

to move to the “target” phase.² This last point is informed by the analyst’s observation that in the operation up until now, the HVI has not left his family at night, as discussed above. The collection manager makes similar determinations for the other targets, as shown in the table.

Table 5.1. ISR Status as of Hour 0

Priority	Target	Targeting Phase	Information Needed to Enable Next Decision	ISR Tasks Required to Fill Information Needs	ISR Resources and Time Required
1	HVI 2	Track	<ul style="list-style-type: none"> Location of HVI and family members 	<ul style="list-style-type: none"> Maintain custody of primary target Account for all personnel previously observed 	<ul style="list-style-type: none"> 2 RPAs 2 PED analysts X hours
2	Training camp	Track	<ul style="list-style-type: none"> Presence of trainers Departure of noncombatant servant 	<ul style="list-style-type: none"> Confirm that targets perform trainer’s activities Account for noncombatant servant 	<ul style="list-style-type: none"> 2 RPAs 2 PED analysts X hours
3	HVI 8	Find	<ul style="list-style-type: none"> Courier transit to potential HVI meeting 	<ul style="list-style-type: none"> Confirm that courier activities are operational, not innocuous or incidental 	<ul style="list-style-type: none"> 1 RPA 1 PED analyst X hours

Now suppose that the situation on the ground has changed by Hour 4: ISR confirms that a courier known to be associated with HVI 8 has joined a group of men in what appears to be an operational meeting. The commander has the opportunity to move the operation to the “fix” phase for this target, which would seek to confirm that the meeting is indeed operational and not innocuous. However, it only makes sense to do so if the necessary ISR resources are available to support the “fix” phase.

Using the same historical assessment techniques, tradecraft, and past experience as described above, the analyst informs the collection manager of the information needs associated with fixing this target, the ISR tasks, and the projected amount of resources and time needed, as summarized in Table 5.2. Moving to the “fix” phase for HVI 8 would require one additional RPA and one additional PED analyst than are currently allocated to this target. The collection manager now faces a resource allocation decision: Should he reassign the necessary resources from one of the other HVIs, and if so, which one?

² We do not specify times in this example, which can vary as a result of a number of factors, including mission, target, and environment.

Table 5.2. Potential “Fix” of HVI 8

Priority	Target	Targeting Phase	Information Needed to Enable Next Decision	ISR Tasks Required to Fill Information Needs	ISR Resources and Time Required
3	HVI 8	Fix	<ul style="list-style-type: none"> Identification of HVI among meeting participants 	<ul style="list-style-type: none"> Confirm that target matches HVI’s visual signature Establish how associates are connected to HVI network 	<ul style="list-style-type: none"> 2 RPAs 2 PED analysts X hours <p>Shortfall:</p> <ul style="list-style-type: none"> 1 RPA 1 PED analyst

NOTE: Yellow = projected.

The analyst can support this resourcing and employment decision by assessing the costs and benefits of each option. The assessment method used in this example measures and compares the opportunity costs for different ISR resourcing options with respect to their ability to answer intelligence problems and achieve operational objectives. The thought process is illustrated in Table 5.3. In this case, the benefits of a given option are indicated by how much that option advances the commander’s overall operational objectives by enabling potential progress along the F2T2EA targeting cycle. These are relatively straightforward in this case: Reassigning ISR resources from other targets to HVI 8 would potentially result in HVI 8 being fixed, moving one step close to achieving the ultimate objective of striking him. The benefit of not reassigning resources is that Blue could continue to track HVI 2 and the training camp—both higher priorities as defined by the commander.

Table 5.3. Resourcing and Employment Decision: Reallocate to Enable Fixing of HVI 8?

Option	Reallocation	Benefits	Costs
A	Transfer 1 RPA and 1 PED analyst from HVI 2	Opportunity to fix HVI 8	Revert to fix phase (costs 2 RPAs and 2 PED analysts for 4 hours to restore progress)
B	Transfer 1 RPA and 1 PED analyst from training camp	Opportunity to fix HVI 8	Revert to fix phase (costs 2 RPAs and 2 PED analysts for 18 hours to restore progress)
C	None	Opportunity to continue tracking and eventually target HVI 2 and training camp	Give up opportunity to fix HVI 8 (costs 1 RPA and 1 PED analyst for possible weeks until next chance to fix HVI 8)

The key cost question, then, is the degree to which removing resources from either HVI 2 or the training camp would hurt the progress already made toward enabling the commander to move to the targeting stage. The analyst can estimate this by determining (1) how much progress along the information-gathering process (for a given F2T2EA stage) or the F2T2EA process itself would be sacrificed by withholding or removing ISR resources and (2) how many ISR resources and time would be required to restore that lost progress? The first question is necessary because of the nature of ISR operations for time-sensitive targets: If the track on the target is lost, then it

may be necessary to return to the beginning of the F2T2EA cycle and find the target anew—a setback that could take days or weeks to rectify.

In the real world, the time to make up lost progress would have to be calculated based on past data about each HVI. The past can never truly predict the future, but in this case past data would inform what the decisionmaker might expect the time to make up lost progress to be. Analysts and decisionmakers in the real world would also potentially be aided in this example by applications and tools that can retrieve and fuse data and information. Having a convenient interface that does not require manual labor to input information would also aid convenience.

In the example at hand, the analyst determines that although removing one RPA and one PED analyst from HVI 2 would sacrifice the progress made in the “track” phase, the remaining RPA and PED analyst could maintain the “fix” on this target. Consequently, it would require only the return of the reassigned resources and four hours of collection time to restore lost progress. This is because past experience suggests that there is a low risk that HVI 2 will leave the house before dawn, as noted above. The cost would be higher for the training target because the tracking task must account for both combatants and noncombatants who follow less predictable patterns. Consequently, even though HVI 2 is a higher-priority target than either the training camp or HVI 8, the cost-benefit assessment supports the decision to reassign resources from this target, resulting in the new status shown in Table 5.4.

Table 5.4. ISR Status as of Hour 4 (After Reallocation)

Priority	Target	Targeting Phase	Information Needed to Enable Next Decision	ISR Tasks Required to Fill Information Needs	ISR Resources and Time Required
1	HVI 2	Fix	<ul style="list-style-type: none"> Continued location of HVI 	<ul style="list-style-type: none"> Maintain custody of primary target 	<ul style="list-style-type: none"> 1 RPA 1 PED analyst X hours
2	Training camp	Track	<ul style="list-style-type: none"> Presence of trainers Departure of noncombatant servant 	<ul style="list-style-type: none"> Confirm that targets perform trainer's activities Account for noncombatant servant 	<ul style="list-style-type: none"> 2 RPAs 2 PED analysts X hours
3	HVI 8	Fix	<ul style="list-style-type: none"> Identification of HVI among meeting participants 	<ul style="list-style-type: none"> Confirm that target matches HVI's signature Establish how associates are connected to HVI's network 	<ul style="list-style-type: none"> 2 RPAs 2 PED analysts X hours

NOTE: Yellow = new status.

Next, suppose that by Hour 8, a candidate for HVI 8 has been identified and is being tracked when an another (unanticipated) decision point arrives: The target has been dropped off and is walking alone in a potentially targetable area. The opportunity to progress to the targeting phase

is valuable but must be weighed against competing priorities as well as its own chance of succeeding. To successfully target for engagement in this area, collection managers estimate (using a similar analysis process as depicted in Table 5.2) that only the two aircraft already assigned to HVI 8 are needed, but more analytical support is desired to confirm what has been observed and to assess the target environment for a potential strike. To determine how a third analyst could be reassigned to HVI 8, the opportunity costs of several options are considered, as depicted in Table 5.5.

Table 5.5. Resourcing and Employment Decision: Reallocate to Enable Targeting of HVI 8?

Option	Reallocation	Benefits	Costs
A	Transfer 1 PED analyst from HVI 2	Opportunity to target HVI 8	Revert to “find” phase (costs 1 RPA and 1 PED analyst days or weeks to restore progress)
B	Transfer 1 PED analyst from training camp	Opportunity to target HVI 8	None (only 1 PED analyst required to fill remaining tasks)
C	None	Opportunity to continue tracking and eventually target training camp	Potential collateral damage effects if targeting proceeds without sufficient resources

Removing the only remaining analyst from HVI 2 could be catastrophic, since no one would be monitoring for the target’s departure. If HVI 2 moves locations, it could take weeks and sometimes months to re-establish custody. In contrast, removing one of the two analysts presently monitoring the training camp would have no opportunity costs because, eight hours into the operation, the presence of trainers has been confirmed, and only one analyst and two aircraft are needed for the remaining task of monitoring for the noncombatant’s departure. Therefore, one PED analyst could be reassigned from the training camp to support the targeting of HVI 8, as reflected in Table 5.6.

Table 5.6. ISR Status as of Hour 8 (After Reallocation)

Priority	Target	Targeting Phase	Information Needed to Enable Next Decision	ISR Tasks Required to Fill Information Needs	ISR Resources and Time Required
A	HVI 2	Fix	<ul style="list-style-type: none"> Continued location of HVI 	<ul style="list-style-type: none"> Maintain custody of primary target 	<ul style="list-style-type: none"> 1 RPA 1 PED analyst X hours
B	Training camp	Track	<ul style="list-style-type: none"> Departure of noncombatant servant 	<ul style="list-style-type: none"> Account for noncombatant servant 	<ul style="list-style-type: none"> 2 RPAs 1 PED analyst X hours
C	HVI 8	Target	<ul style="list-style-type: none"> Room occupancy at HVI compound 	<ul style="list-style-type: none"> Confirm there are no unobserved personnel in targeted infrastructure or environment Identify zones of engagement 	<ul style="list-style-type: none"> 2 RPAs 3 PED analysts X hours

In this example, we have demonstrated how structured assessments can enable decisionmakers to

- plan to allocate significant resources to lower-priority targets due to the *proximity of decision points* that would accelerate *progress toward objectives*
- dynamically reallocate resources from the top priority to an emergent target by *comparing the opportunity costs* of either pursuing or temporarily suspending different objectives among these targets
- assign PED analysts based on changes in analytic needs rather than strictly based on a fixed ration of analysts to ISR collection platforms
- deny additional resources to a potentially rewarding target because it would come at the cost of losing progress toward objectives on other targets, where *restoring lost progress* would require many resources and much effort.

A key point is that this type of analysis is only possible because the costs and benefits have been quantified using relevant indicators (e.g., information needed to enable decisionmaking, resources required to fill those needs) and available data. It is conceivable that a commander or collection manager could make the same decisions illustrated here through intuition alone, but in the absence of objective, quantitative analysis, he or she could just as easily have insisted on keeping maximum resources allocated to the two highest-priority targets, thereby allowing a low-cost, high-benefit opportunity to advance an operational objective slip away. At the same time, we emphasize that the structured assessment methodology outlined here is not intended to replace good human judgment but to aid and support good decisions.

Aligning People with Problems

The USAF ISR community is fundamentally devoted to attacking intelligence problems. However, many members of this community are not organized to focus directly on these problems, as was highlighted in the example just given where analysts were following active intelligence threads. Instead, many analysts within the USAF ISR community help execute planning or PED for particular ISR sorties or missions or are connected with specific collection platforms. This is not true in all parts of the community. Parts of the ISR enterprise that are organized according to particular types of problems appear to have success in not only developing a good understanding of how relevant and valuable ISR is in a particular context, but communicating this to users of both ISR products and ISR assessments.

One weapon system that is very important for ISR assessments is the DCGS. Although the analysts within the DCGS are focused on PED, they typically provide an initial assessment on what happened during an individual ISR mission for collection platforms under their purview. This information is entered into UNICORN, which is used widely in the community for assessment purposes. In this way, DCGS analysts communicate assessments for ISR missions to

the broader community, and this information is used for other ISR assessments by the AOC, HQ USAF, the Joint Community, and others.

However, PED teams that are organized around collection platforms might not be optimally organized for working deeply within the context of their missions, which is needed for fully understanding what information is being sought and whether the collection is providing that information (i.e., providing an accurate assessment). Here, we consider how shifting PED analyst operations to more a “problem-centric” approach might enable these analysts to better communicate their findings in the context of progress towards solving problems, which is a more assessment-focused approach. Although we chose to focus on the DCGS for this example, the general concept may apply more broadly to the communication of assessments in other parts of the enterprise. For this discussion, we rely on interviews conducted for this research with analysts at multiple sites within the DCGS enterprise and AOC analysts working with the DCGS, supported by findings in previously published RAND work examining this topic.³

An example of an alternative construct for PED organization is to exploit information from all ISR assets operating within a particular geographic area and focused on inter-related missions by a single PED Super Crew (Figure 7.1). Note that this is for assessment discussion purposes only; USAF would have to consider a range of options (and whether changes are worthwhile) before implementing any changes to PED crew construction. In Figure 7.1, the common thread for all the missions being supported at one given time by this PED Super Crew is that all of them have to belong to the same operational line of effort within which several intelligence problems have been defined. Moreover, we assume that a set of ISR assets has been assigned to these intelligence problems and, therefore, real-time exploitation of the collections from these assets is performed by the PED Super Crew. All members of the Super Crew are assumed to know the operational line of effort well and to have sufficient familiarity with the intelligence problems. Therefore, if an analyst has to be reassigned from one intelligence problem to another during mission execution, such reassignment (including bringing the analyst up to speed on the latest developments) can take place in a short amount of time.

The PED Super Crew composition is illustrated in left-hand side of Figure 7.1, and consists of the following:

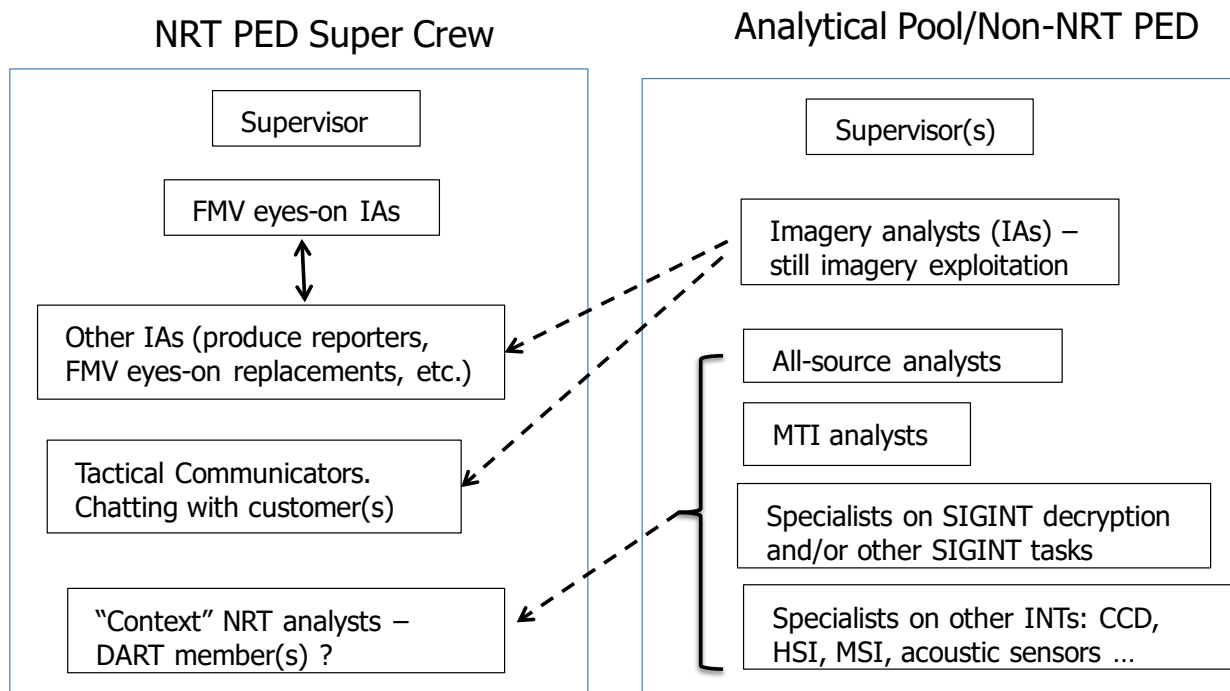
- the Super Crew supervisor
- the Eyes-on Imagery Analysts (IAs), whose number depends on how many ISR assets are providing collections that need to be exploited in NRT by the Super Crew
- other IAs, including product reporters, product editors, and eyes-on IA replacements
- the Tactical Communicators in charge of providing information in NRT to the customer(s) via chat

³ See Menthe et al., 2012.

- the “Context” analysts who are in charge of keeping track of the “big picture;” this small group comprises primarily all-source analysts, but may borrow specialists from the analytical pool, as needed (see analytical pool specialties below).

The analytical pool (right-hand side of Figure 5.8) provides the surge capability, e.g. additional IAs may be needed by the Super Crew at specific times. It also provides intelligence analysts with specialties potentially relevant to a variety of information problems, such as Moving Target Indicator, SIGINT, Coherent Change Detection, Hyperspectral Imaging, Multispectral Imaging and others. Any of these specialists may be borrowed by the Super Crew and temporarily be part of the Super Crew’s “Context” analyst group.

Figure 5.2. PED Super Crew and Analytical Pool



NOTE: CCD = charge-coupled device; DART = DCGS Analysis and Reporting Team; HSI = hyperspectral imagery; MSI = multispectral imagery; MTI = moving target indicator.

Compared to the platform-centric approach, these analysts are broadly focused on the overall Super Crew missions (namely, the different intelligence problems that the Super Crew is working on) instead of only the collections from a single ISR platform. This enables assessment of ISR activities in a holistic manner and links it directly to the intelligence problems.⁴ For

⁴ This construct also has some other benefits, which are only peripherally related (if at all) to assessments. The FMV eye-on analysts are allowed to observe more than one FMV feed, depending on operational conditions (we will illustrate this when discussing a vignette). The construct also allows for prompt reassignment of analysts from one intelligence problem to another when bad weather or aircraft malfunction interrupt or prevent the ISR aircraft sortie.

example, the analysts in the Super Crew can identify how many of the intelligence problems have either been resolved or have had information gathered that can advance the narrative toward solving the problem. This information can then be directly fed back to the information users and AOC, which can drive future collections and perhaps even overall operational strategy. Further, this type of crew composition might also streamline how information related to each problem set is entered into databases such as UNICORN, which is especially important for free-form fields (i.e., those that do not have set responses to choose from).

Deciding Whether to Reallocate Resources Between Theaters

We now turn to a different example, in which air components must assist their respective COCOMs in conducting assessments to support a DoD decision about potentially reallocating ISR resources in response to emerging threats. Like the previous example, this case is notional but based on real-world choices that warfighters, planners, and USAF ISR managers often face. Each year, the DoD allocates ISR resources to different theaters based on projected needs for the year. However, the plan can be reexamined if necessary should theater needs change, as they do in this example. Here, an airborne ISR platform originally assigned to one theater might be needed by another. The air components in both theaters must assist their COCOMs by demonstrating why the platform is needed in their respective areas of responsibility before needs are adjudicated at a higher level (i.e., within the DoD). In the real world, the challenges associated with assessment procedures and data availability hinder the timely delivery of cost-benefit analyses to inform such resourcing and employment decisions. In this example, we assume a structured assessment process that relies on clear theater goals, ISR roles, and platform-usage information that is available to both theaters.

The context of this example is as follows: Theater A is focused on supporting ground troops in a desert environment combating an insurgency linked to global terrorism. It employs two ISR platforms that can take still imagery pictures, among others. One imagery sensor is high-definition but does not operate in all-weather conditions. The other imagery sensor produces images of poorer quality but can operate in all-weather conditions. Generally speaking, the still imagery is supporting the development of targets for possible prosecution in support of ground-troop forward movement to retake areas overrun by insurgents. Theater A also uses still imagery to monitor changes in the size and location of insurgent camps as a means of providing warning ahead of attacks on Theater A's ground troops. Some of the still imagery targets could theoretically also be monitored using wide-area imagery sensing, which is currently being exclusively used to monitor for threats around the main Theater A operating base.

Theater B is a part of the world that has enjoyed relative regional stability for about two years. Consequently, only about one airborne imagery flight per week is allocated to this theater. At the time of this example, however, Theater B alerts its chain of command about what appears to be a hastily planned naval exercise by a potential adversary. After two weeks, military

maritime activity appears to be ramping up, including possible construction of new facilities that are well-positioned to potentially menace a close ally. Theater B currently has no airborne imagery sensors, and satellite imagery from various sources is not reliably available. It is also the wet season, suggesting a need for an all-weather sensor. Updates are required every two days to support decisionmaking about whether to recommend additional diplomatic, economic, or military actions to halt the potential adversary's ramp-up in naval military activity.

Theater B requests an all-weather imagery capability to ensure that its platform and sensor can fly and produce needed intelligence every other day (or more frequently if necessary). A DoD decisionmaker assigned to examine this request decides that Theater A is the only area capable of potentially giving up an ISR collection asset because of the large volume of ISR capabilities in its area and because it does not strictly require an all-weather capability. Theater A, however, fears that losing access to this asset will impact its mission because of the sheer volume of imagery needed to support its operation. The DoD decisionmaker asks each theater to answer the following assessment questions, which are delegated to the respective air components:

1. What operational decisions in the theater require imagery support?
2. What are the imagery revisit times required to support timely operational decisionmaking?
3. How will operational goals be impacted if an all-weather imagery collection asset is not available?

In this notional example, we assume that in each theater, assessment analysts within the AOC would perform the following steps to support the DoD decisionmaker's resourcing and employment decision:

1. Link the ISR roles associated with imagery collection to theater objectives.
2. Identify the decisions that need to be made in each theater on the basis of information derived from imagery collection and calculate or infer the appropriate interval between imagery collection needed to support this decisionmaking.
3. Specify metrics for operational goals and calculate or infer what the impact of not having an imagery collection asset would be.

Theater A analysts find that several ISR roles could potentially meet their theater's information needs. These include support to operations (third-order roles: search for adversary positions and activity ahead of the maneuvering force, monitor identified or suspected threats), and, perhaps also to a lesser degree, IPoE (third-order role: identify and characterize potential threats) and targeting (third-order role: support basic, intermediate, and advanced target development). Theater B analysts find identify IPoE (third-order roles: characterize the operating environment, identify planning considerations with respect to potential threats, evaluate adversary capabilities) and warning intelligence (third-order role: identify signs of mobilization) as important ISR roles for their theater.

Revisit rates would be ideally calculated by examining past similar situations and comparing outcomes with ISR employment. For the sake of this illustrative example, we can assume that the revisit rates featured in Table 5.7 are reasonable based upon an assessment of this nature—facilitated, of course, by databases with accurate information made readily retrievable using algorithms and user interfaces that enable assessment analysts to pull together and understand the information.

This quantitative assessment can be reinforced by examining in detail the types of additive information that would support each theater's objectives, which can be logically examined to confirm whether the quantitative analysis makes sense. In Theater A, types of additive information might notionally include the location of sites confirmed as camps, evidence that camps have moved or may be about to move, the size of camps, identification of weapons at the camps, movement of suspicious people or vehicles towards the main base, or other signs that an attack could be imminent. The types of additive information that would support Theater B's goals include observations of freshly broken ground or new foundations being built, detection of changes in the size or geometry of known military construction sites, presence of cement trucks or other equipment associated with construction, and the number of military ships in particular ports.

Table 5.7. Notional Revisit Rates and Sensors Required for Theaters A and B

Theater	Goal	Third-order ISR roles that support goal	Imagery sensor	Revisit rate required
A	Compel adversary to retreat	Search for adversary positions and activity ahead of the maneuvering force	Any imagery	Daily
		Monitor identified or suspected threats	Any imagery	Daily
		Support basic target development	Any imagery	Weekly
		Support intermediate target development	Any imagery	Weekly
		Support advanced target development	Any imagery	Weekly
	Protect forces	Identify and characterize potential threats	Wide-area	Daily
		Search for adversary positions and activity ahead of the maneuvering force	Any imagery	Daily
		Support advanced target development	Any imagery	As needed based on threat emergence
B	Identify new areas of military construction	Characterize the operating environment	Any imagery	Weekly
		Identify planning considerations with respect to potential threats	Any imagery	Weekly
		Evaluate adversary capabilities	Any imagery	Weekly
	Monitor areas of existing construction	Identify planning considerations with respect to potential threats	All-weather	Every two days
		Evaluate adversary capabilities	All-weather	Every two days
	Provide awareness about any changes in military ships in port	Identify signs of mobilization	All-weather	Every two days

Finally, assessment analysts would identify overarching goals for each theater and ways to measure progress towards them as a first step in understanding how the lack of an imagery collection asset might impact these goals. Theater A's goals include compelling the adversary to retreat and protecting Blue forces. Overarching progress towards these goals (not specifically ISR's role in supporting them) could be measured by tracking the distance retreated by the adversary as indicated by camp positions, the number of surprise attacks on the main base, and the number of casualties among Theater A ground troops during the course of operations. Theater B's goals include identifying new areas of military construction, monitoring areas of existing construction, and providing awareness every two days as to whether the status of military ships in port has changed. Overarching progress towards these goals could be measured by the number of new areas of construction identified, whether changes in existing construction have occurred within 48 hours, and whether the status of military ships in port can be affirmatively given every other day. The measures for Theater B include specific times because we assume that a high-level decisionmaker is requesting updates at this interval to ensure timely reaction to developments.

Based on the information in Table 5.7, Theater B has a compelling argument for securing the all-weather sensor. The all-weather sensor is the only imagery asset that can perform the ISR roles required for two out of three goals. Without the all-weather asset, Theater B cannot hope to accomplish the majority of its goals. Theater A cannot make the same claim, since the all-weather sensor is not strictly required for ISR roles in support of any of its goals. However, Theater A could claim that losing the all-weather imagery sensor will stretch its other imagery resources, potentially leading to negative consequences for one or more operational goal. Such an assessment might be compelling if Theater A could provide, for example, a statistical demonstration of increased attacks on the main base at a time when imagery was in shorter supply than it is now.

Although Theater A is supporting an important operation with lives at risk, Theater B may need to inform a tricky strategic decision. A high-level decisionmaker would ultimately have to which goals are more important and would suffer most from the lack of an all-weather sensor.

In this second illustrative example, we have demonstrated how concepts such as ISR roles, benefit indicators linked to operational objectives, and a structured assessment approach can enable air components to support joint decisionmaking that requires

- weighing the needs of different theaters using an “apples-to-apples” assessments approach
- logically working out the benefits of ISR in each theater, as well as the potential consequences for overarching objectives
- quantitatively characterizing indicators for mission success or laying the foundations for calculating such indicators
- consistently explaining the reasons for making a particular decision.

As in the previous example, this type of assessment would not be used in place of a decisionmaker’s judgement. Rather, it would provide a useful tool to aid or justify that judgement.

Conclusion and Suggested USAF Actions to Improve Assessment Step 4a

This chapter has offered two notional examples of air component ISR assessments that might be conducted using real-world data. The first involved a very short-duration dynamic reallocation decision in the context of an HVI operation. The second highlighted a broader-scale, longer-term reallocation decision in which the air components in two respective theaters had to help justify COCOM reasons for having access to a particular imagery sensor.

The examples presented in this chapter have highlighted some important points with respect to USAF ISR.⁵ First, ranked lists of targets alone do not capture the value of each target at any given time or in every situation. ISR often uncovers activities and relationships that defy existing

⁵ These insights may also apply to other ISR communities, but we did not explicitly study these.

priorities and can potentially influence future priorities. Our notional HVI scenario provided just such an example, and the results of this use case show how *proximity to decision points* and *progress toward objectives* can be help measure when a high-priority target may temporarily have less value than a lower-priority target, for instance. Although the theater allocation scenario would have had to take into account national priorities for a final decision to be made, it did highlight that the sensor in question was potentially more useful in one theater than in another.

Second, linking overarching goals to ISR roles and the ISR assets that could perform these roles helps to parse exactly which resources are needed and why. Linking benefit indicators with these goals and ISR roles provides a basis for tracking progress towards these goals and ISR's real or potential role in making progress (note that a similar analysis could be performed to examine specialty PED resource allocation—for example, how to use a notional team with a specialty in a particular problem set, such as improvised explosives). The theater reallocation example was able to show what would be lost and what would be gained by reallocating the all-weather sensor in terms that permit a decisionmaker to lower the number of dimensions considered in the decision. With sufficient data available, assessments can also compare potential future situations with similar past ones (e.g., risk to the main base in the theater reallocation example) to show what might happen if ISR of a certain type is no longer available, thereby providing a decisionmaker with a sense of the opportunity costs.

Third, assessments in both of these examples would benefit from reliable databases, tools, and analysis applications to help with data fusion and other analytic activities. The databases used to conduct these assessments would contain information entered by PED analysts performing initial assessments of what happened during ISR sorties; organizing their efforts by the problems they are trying to help solve could help to better link their analyses to problems that are central for understanding ISR's value in assessments. As we have stressed elsewhere in this report, even if real-world data are theoretically available, they must be accurate, retrievable, and understandable by an USAF ISR assessments analyst.

There are steps USAF can take to further enable these kinds of assessments using real-world data, assuming that the steps detailed in previous chapters have already been implemented. These steps are supported by the following information collected during DCGS site observations, analyst interviews, and previous RAND research:⁶

- Train USAF ISR assessment analysts to lay out key questions and steps using notional examples such as those provided here and real-world examples of limited-scope until databases are sufficiently ready to conduct larger-scale cost-benefit assessments.
- Prepare USAF and air component assessments to inform resource allocation and reallocation decisions at every service and joint echelon on the basis of specific goals and intelligence problems and how these are linked to ISR roles. When possible, analysis should include a detailed look at what is lost or gained by different ISR resourcing and

⁶ See Menthe et al., 2012.

employment decisions. Some aspects of this will also apply to making investment decisions, where long-term costs and benefits must be weighed based on projected or assumed strategic, operational, or tactical needs.

- Consider new USAF PED force presentation models (alluded to in Chapter 3), to include analysts “on the line” (i.e., conducting real-time exploitation) that could have benefits for assessment.
- Encourage airmen with coding skills—or the ability to learn them—to help their units solve emerging assessments problems by developing prototype applications or tools.⁷

As mentioned at the beginning of this chapter, there may be situations where it is impossible or impractical to conduct USAF and air component ISR assessments using real-world data alone. The next chapter discusses how computer simulations can be used to support assessments in these cases.

⁷ If necessary, these can ultimately be aided by contractor or other support to develop fully functional tools.

6. Step 4b: Analyze the Value of ISR Using Simulations

Real-world operations and computer simulations can be complementary sources of information for USAF airborne ISR assessment. Operational data sets are, of course, the gold standard. But as described throughout this report, such data sets are often unavailable to USAF ISR assessment analysts, and even when they are available, the data are often incomplete and ill-conditioned. Simulations can help fill the gap when operational data sets are insufficient to answer important ISR assessment questions.

Simulations are particularly valuable when attempting to make assessments involving the future and the contrafactual such as projecting the potential benefits of new sensor technologies that have yet to be fielded or answering “what if” questions about courses of action not taken. Simulations can also be used to augment operational data sets where outcomes are obscured, such as assessing the effectiveness of an air campaign where there are no eyes on the ground to verify battle damage assessment indicators.

However, it is important to bear in mind that simulations are, by definition, merely shadows of the phenomena under study.¹ *Garbage in, garbage out* is a common caution when working with computer simulations—bad input data or flawed assumptions can render the output meaningless. The same scientific steps used for ISR assessments based on operational data must be employed when using computer simulations: The assessment question must be clearly defined, the indicators must be selected in advance, and the vignettes must be designed to measure them. Operational data should be used to validate aspects of the simulation wherever possible, but this can usually be done only in a piecewise manner by comparing different indicators with different data sets. If a complete operational data set were available for comparison, we would simply use that data for the assessment itself, and there would be no need to conduct simulations.

In this chapter, we illustrate how computer simulations can be used effectively to inform ISR assessments by simulating the surveillance operations in the HVI example described in the previous chapter. Note that our discussion here applies to the USAF airborne ISR community; the focus and methods of analysis may change if conducted to examine an ISR assessment question within the broader IC or Joint Communities. We show here how simulation can be used to identify critical points for decisionmaking and can help select appropriate indicators for assessment. Though the specific details of the notional HVI example are not necessarily realistic, the purpose of this is to provide an illustration which can be done using a stylized situation.

¹ We are mindful of Box’s famous dictum: “Essentially, all models are wrong, but some are useful” (George E. P. Box and Norman R. Draper, *Empirical Model Building and Response Surfaces*, New York: John Wiley & Sons, 1987).

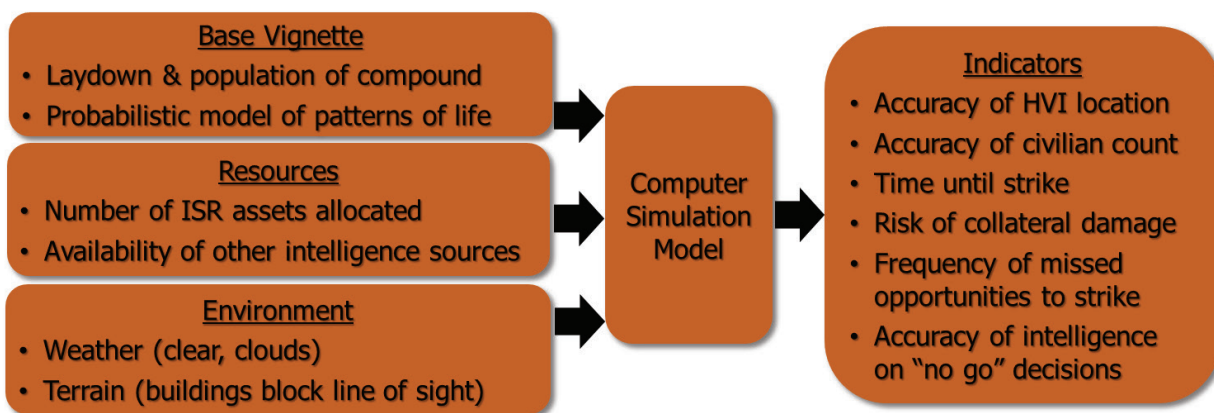
Selected results are presented graphically in the following sections, while the full results are tabulated in Appendix A. We also describe a simple mathematical model for fixed-point security in Appendix B, to illustrate how less-complicated modeling efforts can be useful in some cases.

Approach

Methods

Our overall approach to simulation for ISR assessments for the HVI vignette is one of theme and variations: we define a base case and then vary the resources and environment to see how each affects the indicators. Whenever we asked SMEs what values the various indicators should be for this type of operation, the response was always the same: “It depends.” Our goal here, therefore, is not to determine what the “right” values are—for there are none—but rather to reveal the dependencies. Figure 6.1. illustrates this approach.

Figure 6.1. Modeling Approach



For the HVI vignette, the base case does not attempt to replicate a specific historical operation, but rather is an average case that we believe to be sufficiently representative of this type of operation. To analyze the HVI vignette, we use a stochastic, agent-based model. For our purposes, we define an agent as “an autonomous entity situated in an environment that it can perceive, and in which it can act.”² Agent-based modeling is a good choice for many ISR assessments because it allows one to model perceptions and ground truth separately and to observe the decisionmaking. This enables us to measure both the quality of the intelligence and its effect on the final outcomes. There are many different agent-based models and the types of agents in these models may vary greatly in complexity and number. While the authors favor

² Clint Heinze, *Modelling Intention Recognition for Intelligent Agent Systems*, Canberra, Australia: Defence Science and Technology Organisation Systems Science Laboratory, DSTO-RR-0286, 2004.

models with fewer but more complex agents, this choice should be tailored for the ISR assessment question.³

Stochastic modeling is important because each individual iteration of the model will then be different. Weather, sensor snapshots, and enemy activity all vary probabilistically. This enables us to understand both typical and atypical values as well. While the average often represents the most probable outcome, the exact average outcome is usually still quite unlikely. (For example, for 100 fair coin tosses, the chance of getting a perfect 50/50 split is less than 8 percent.)⁴ Stochastic simulations are important to ISR assessment because they allow us to gauge the likelihood of unlikely outcomes, i.e., to estimate *risk*.

Simulation

The specific computer simulation model used for this analysis was RAND's Systems and Concepts of Operations Operational Effectiveness Model (SCOPEM), a suite of functions, subroutines, and libraries written for the System Effectiveness Analysis Simulation (SEAS) modeling environment. SEAS is a time-stepped, stochastic, agent-based model maintained and developed by *ExoAnalytic Solutions* for the USAF Space Command, Space and Missile Systems Center, Advanced Systems & Development Directorate; it is part of the USAF Standard Analysis Toolkit and the USAF Space Command Modeling and Simulation Toolkit. SEAS version 3.11.1 was used for this analysis.

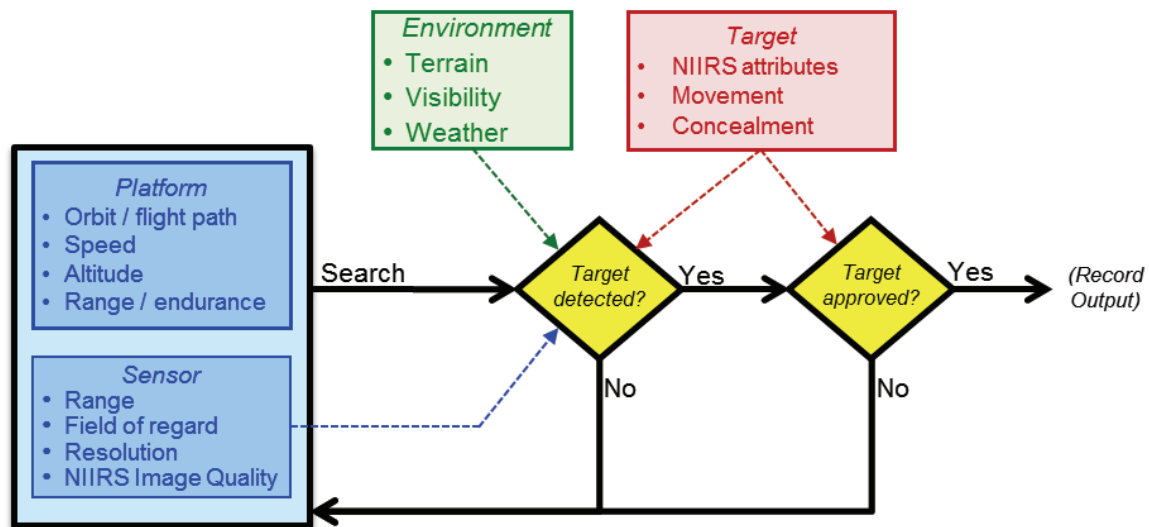
SCOPEM has been developed over the past 12 years by PAF for a wide range of ISR-related analysis, from RQ-4 Global Hawk maritime surveillance, to F-22 Raptor sensor performance, to the use of RPA in the hunter-killer mission. More information on this model can be found in previously published reports.⁵ An overview of the targeting logic in SCOPEM is shown in Figure 6.2. Target detection and strike approval result from the interaction between aircraft (platform and sensor) capabilities, target characteristics, and environmental conditions. In this model, the aircraft, HVI, insurgents, and civilians are all represented as independent agents.

³ Previous work has varied from a handful of fairly sophisticated agents to tens of thousands of very simple ones. The analyst should guide the simulation, not the other way around.

⁴ James L. Hein, *Discrete Mathematics*, 2nd ed., Sudbury, Mass.: Jones & Bartlett Publishers, 2003.

⁵ Lance Menthe, Myron Hura, and Carl Rhodes, *The Effectiveness of Remotely Piloted Aircraft in a Permissive Hunter-Killer Scenario*, Santa Monica, Calif.: RAND Corporation, RR-276-AF, 2014; Sherrill Lingel, Lance Menthe, Brien Alkire, Scott Grossman, Robert A. Guffey, Keith Henry, Lindsay D. Millard, Christopher A. Mouton, George Nacouzi, Edward Wu, and John Gibson, *Methodologies for Analyzing Remotely Piloted Aircraft Effectiveness in Future Roles and Missions*, Santa Monica, Calif.: RAND Corporation, DB-637-AF, 2012; and Lance Menthe and Jeffrey Sullivan, *A RAND Analysis Tool for Intelligence, Surveillance and Reconnaissance: The Collections Operations Model*, Santa Monica, Calif.: RAND Corporation, TR-557-AF, 2008.

Figure 6.2. SEAS/SCOPEM Targeting Logic



NOTE: NIIRS = National Imagery Interpretability Rating Scale.

The model reports every sighting of every agent by another agent. If an agent would have seen another agent but the line of sight was obscured by environmental effects, that would also be reported. The model also reports the ground truth location of every agent. Agent perceptions are compared to ground truth to measure the accuracy of their operational picture. The simulation writes this information to large, comma-delimited (.csv) files, and we subsequently analyze this output using Microsoft Excel to produce the indicators shown in this chapter.

Vignette

The baseline vignette is an RPA, equipped with FMV sensors, collecting against an HVI who is in a medium-sized training compound along with 20 other insurgents and 20 civilians.⁶ Here, we assume that the training compound is in a rural environment; finding terrorists within a large urban area that could include thousands of civilians would be much more challenging. The goal is to identify the HVI, track his or her movement, account for the location of all civilians in the compound, and determine when there is an opportunity to strike the HVI in a building without risk of collateral damage. As the RPA circles the compound, the various buildings block line of sight and create shadowed areas depending on viewing angle. Meanwhile, the civilians, insurgents, and HVI occasionally move between buildings. Movement is random and the likelihood of movement is given by a Poisson distribution with average frequency of one

⁶ Though we use a stylized example, it is not inconsistent with news reports about terrorist training camps. See Terri Moon Cronk, "Strike on ISIL Camp Protected National Security, Pentagon Press Secretary Says," *DoD*, February 19, 2016; and Joshua E. Keating, "What Do You Learn at Terrorist Training Camp?" *Foreign Policy*, May 10, 2010.

“activity period.”⁷ (We use the same activity period for all agents.) We vary the vignette by increasing the number of RPA, by impairing line of sight with cloud cover, and by providing other intelligence sources. This vignette was deliberately chosen to be challenging so the effects of varying resources and the environment could be more easily observed

A screenshot of the compound as instantiated within SEAS/SCOPEM is shown in Figure 6.3. The compound in our vignette consists of eight buildings of various sizes and shapes in a roughly rectangular formation, with some open areas between them. The green lines show the outlines of the buildings. The white lines show the movement network through which agents may transit between buildings. The red and white dots show the initial location of a few insurgents and civilians, but most are hidden within buildings in this particular snapshot. The buildings used to template this compound were selected arbitrarily using Google Earth Pro from a city in South Central Asia. As noted above, the compound was not designed to imitate any specific locale but merely to exhibit reasonable size and spacing.

Figure 6.3. Compound




The buildings are assumed to be six to eight meters tall with approximately 10 meters between them. The RPA has a viewing angle of between 40 and 50 degrees. For this analysis, we assume the FMV sensors on the RPA have a footprint sufficient to keep the entire compound within the field of view (FOV) and have image quality sufficient to distinguish individual civilians from one another, and from insurgents, with 95 percent probability when they are

⁷ The Poisson distribution is the only “memoryless” probability distribution. In a Poisson process, events have an average frequency of occurrence, but the timing of each event is independent of the previous. Poisson distributions are used to model everything from bus accidents to radioactive decay. See Georg Rasch, “The Poisson Process as a Model for a Diversity of Behavior Phenomena,” *International Congress of Psychology*, Vol. 2, August 1963.

visible.⁸ We assume the buildings completely block line of sight. For the purposes of this analysis, any line-of-sight sensor may be used—electro-optical (EO), infrared, or other—provided it meets these FOV and resolution requirements.

Cloud cover is represented as a thin grid at a given altitude, where each cell has the same probability of being opaque or clear. For light clouds, we chose 25 percent; for medium clouds, 50 percent. Each cloud cell may change state randomly over the course of the iteration, with an average stability of one activity period.⁹ Cloud cell spacing is narrow—40 meters on a side—to represent spotty coverage. This creates situations in which movements may be partially seen but partially hidden—for example, a civilian may be positively identified as exiting one building, but clouds obscure which building he or she enters next. An opaque cloud cell is assumed to completely block line of sight, but the same cloud cell may block different parts of the compound for different RPA depending on the viewing angle. An illustration of the cloud model is shown in Figure 6.4. Note that we do not simulate how changing patterns of sunlight and shadow affect the ability to identify targets at different times of the day, and the day-night transition is assumed to be seamless. For these reasons, our assumptions regarding target visibility are optimistic.

Figure 6.4. Cloud Grid



⁸ Image quality is represented in the model using the National Imagery Interpretability Rating Scale and the General Image-Quality Equation. The canonical reference for understanding both is Jon C. Leachtenauer, William Malila, John Irvine, Linda Colburn, and Nanette Salvaggio, “General Image-Quality Equation: GIQE,” *Applied Optics*, Vol. 36, No. 32, November 10, 1997.

⁹ For simplicity, we use the same Poisson distribution to govern random movement, randomly shifting cloud cover, and the arrival of external intelligence tips.

In this vignette, we assume restrictive rules of engagement (ROE) under which no civilian casualties are permissible. To call in a strike, the intelligence estimates must therefore indicate that the HVI is in a building with no civilians (other insurgents may be present). Mistakes in civilian accounting or HVI location, however, can lead to unwanted casualties. We measure the quality of the intelligence estimate relative to ground truth using two indicators: the accuracy of civilian accounting and accuracy of HVI location. When a decision to strike is made, we look at the final outcome to determine if the correct building was struck, and whether there were any civilian casualties. We also look at whether there were any missed opportunities, because intelligence was incorrect—and we also give credit where intelligence correctly led to a “no-go” decision because the HVI was not alone. Note that we do not model the weapon strike itself. For our purposes, we assume all strikes are lethal to everyone inside the building. Representing the weapons effects in any greater detail would require more rigorous analysis than we wish to conduct here.

In addition to varying the number of RPA and cloud cover, we also looked at the possibility of having additional intelligence tips on the HVI and civilian locations. We make no assumptions about the means by which this intelligence is acquired, save that it is not subject to the same line-of-sight restrictions as aerial surveillance and it is sufficient to indicate in which building the HVI or civilian is located. These intelligence tips are assumed to arrive randomly, following the same activity period as before, and are assumed to be accurate—but the value of such a “sighting” naturally ages just as does any other sighting.¹⁰ The purpose of looking at this second source of intelligence is to investigate some of the benefits of data fusion.

Results

Run Matrix

The base vignette is a single RPA conducting surveillance of the compound with clear weather and no other sources of intelligence. We vary this vignette by considering up to three RPA working cooperatively, by considering light or medium cloud cover, and by considering additional intelligence tips regarding the HVI only or for both HVI and civilian locations. This leads to a run matrix of 27 cases shown in Figure 6.5. Hundreds of iterations are required for each of these individual cases to establish averages and the distributions of results. In this section, we discuss selected results, grouped by indicator. Tables for all cases are given in Appendix A.

¹⁰ We do not consider the case of inaccurate or misleading intelligence tips. That would be an interesting excursion for further analysis.

Figure 6.5. Run Matrix

Baseline	1 FMV Assets	2 FMV Assets	3 FMV Assets
Clear			
Light Clouds			
Medium Clouds			

Baseline + Intel on HVI	1 FMV Assets	2 FMV Assets	3 FMV Assets
Clear			
Light Clouds			
Medium Clouds			

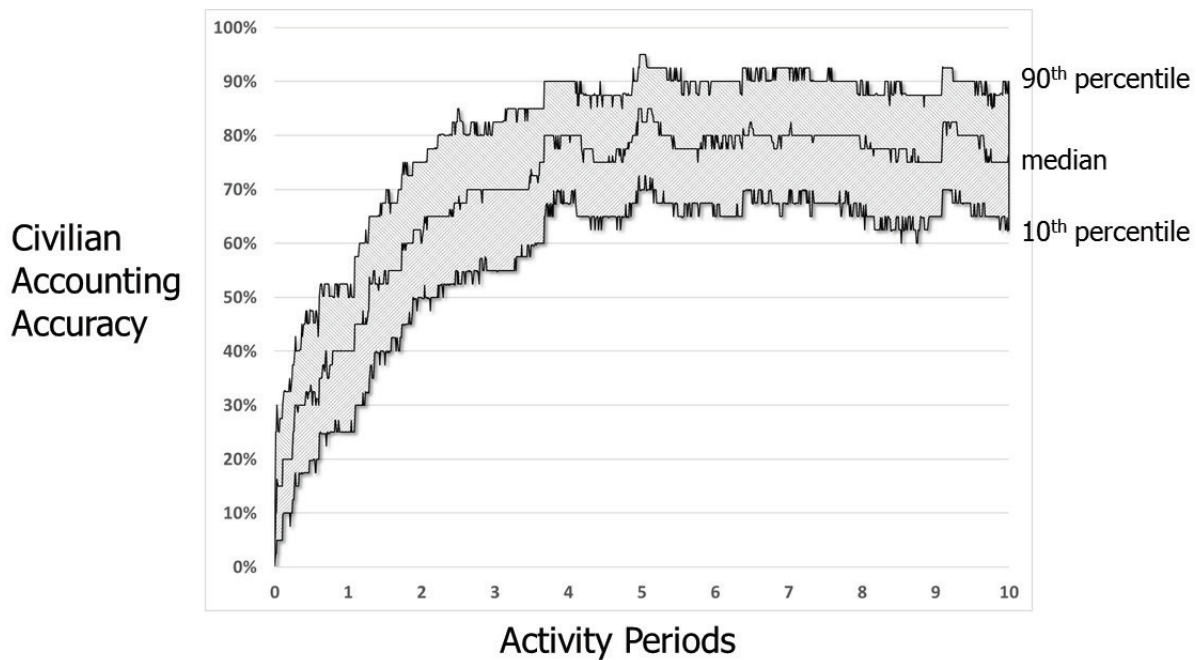
Baseline + Intel on All	1 FMV Assets	2 FMV Assets	3 FMV Assets
Clear			
Light Clouds			
Medium Clouds			

Civilian Accounting

Clear Weather

The most difficult aspect of this ISR challenge is to keep a proper count of where each civilian is located. This simulation allows us to see how good the civilian accounting estimates actually are compared to ground truth. Figure 6.6 shows how the accuracy of the civilian accounting evolves over time after the RPA arrives on station. For each civilian, we marked the location as either correctly known, or incorrectly known or unknown; the vertical axis shows the percentage of civilians that were located correctly. The horizontal axis shows the time, as measured in activity periods. The activity period is important because it is the average time one must wait for any given civilian to move between buildings and therefore be observable from the air.

Figure 6.6. Baseline, Clear Weather, One RPA

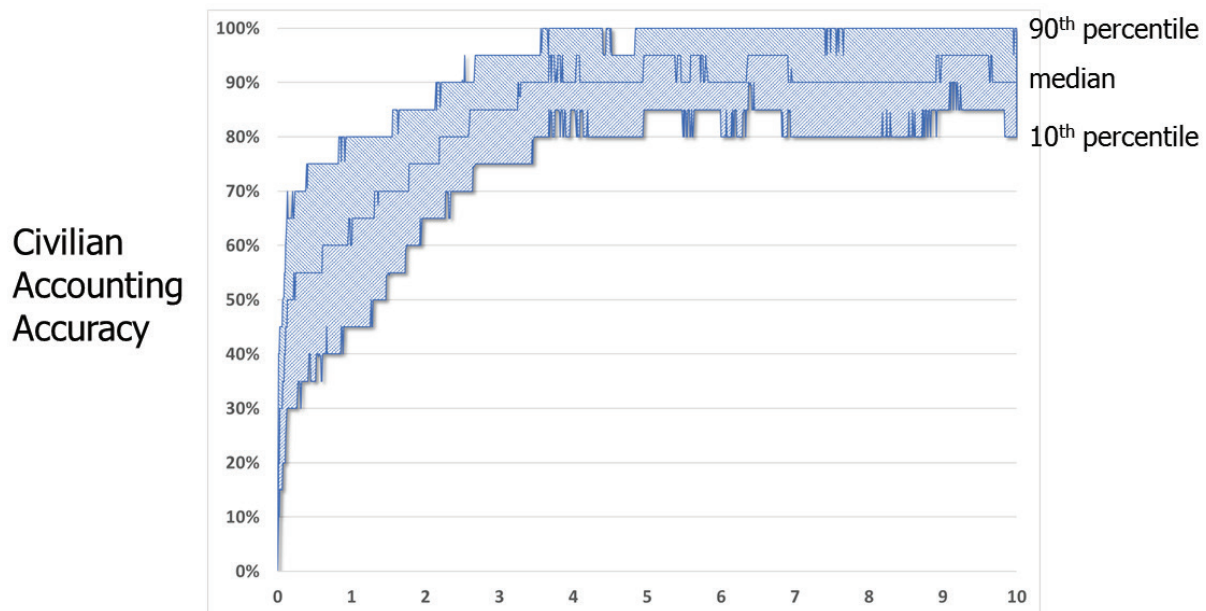


The shaded regions above and below the median show the 90th and 10th percentile of results respectively. This gives a good indication of the variability of the process. It also shows how confidence levels can be read from this kind of data. Note that the median is *not* a high confidence estimate, because, by definition, it is the middling result. The 10th-percentile line is a better way of representing the high confidence estimate: we may say with 90 percent certainty that our civilian accounting is at least as good as the 10th-percentile line.

Three features of this chart are immediately apparent. First, the accuracy of the civilian accounting varies noticeably from moment to moment. This is due in part to the innate randomness of the simulation and in part also to the favorability of the viewing geometry at any given moment. Second, it takes roughly three to four activity periods to build the picture, after which the accuracy remains fairly steady. (In some excursions, we ran the scenario for three times as long as is shown, and it remains steady.) We refer to this period as the *steady state*. Third, the 10th percentile line is surprisingly poor, with only about 67 percent of civilians identified in the steady state—but this is by design. As mentioned above, the base vignette was deliberately chosen to be challenging. A base vignette which provided high success would only allow us to observe the negative dependencies.

In Figure 6.7, we show the same case, but with two RPAs circling the compound. As expected, the civilian accounting accuracy increases and the uncertainty decreases. The time to build the picture, however, remains about the same. This is a general feature we find, because the opportunities for viewing are determined by enemy and civilian activity.


Figure 6.7. Baseline, Clear Weather, Two RPA



We should note that the angle between the two RPAs is important. Naïvely, one might simply put them 180 degrees apart, but this actually leads to the worst outcome, because, when attempting to peer into an “urban canyon,” both assets may lose line of sight at the same time. On the other hand, placing them too close together obviously invites the same difficulty. So, for this particular case, we choose a random angle of between 90 degrees and 180 degrees for each iteration as part of the stochastic simulation. We believe this gives us a better picture of what two RPA can achieve in any given case, without benefiting or being harmed by accidents of geometry due to our particular choice of laydown for the compound.

The third case, with three RPAs working cooperatively, is shown in Figure 6.8. We see again that the average is improved and the variability is diminished. (In this case, the three RPAs fly 120 degrees apart.) Here, finally, the 10th-percentile line pushes us up to 90 percent accuracy. The operational picture is built up somewhat faster than with one or two RPA, but only marginally. The main effect is to improve the operational picture, not to build it more quickly.


Figure 6.8. Baseline, Clear Weather, Three RPA



We can summarize these three cases by showing how the averages change over time.¹¹ The steady-state averages are shown in gray text to the right of the chart. We see here the general feature that adding RPA improves the steady-state picture but does not bring the vignette to steady state much sooner. We expected to see this feature, but we did not expect to see it emerge as clearly as shown in Figure 6.9. As we will discuss further, the time to strike emerges as a poor indicator for mission success because the timing is chiefly governed by enemy and civilian choices.

¹¹ The average values differ from the median values by 2 percent at most.

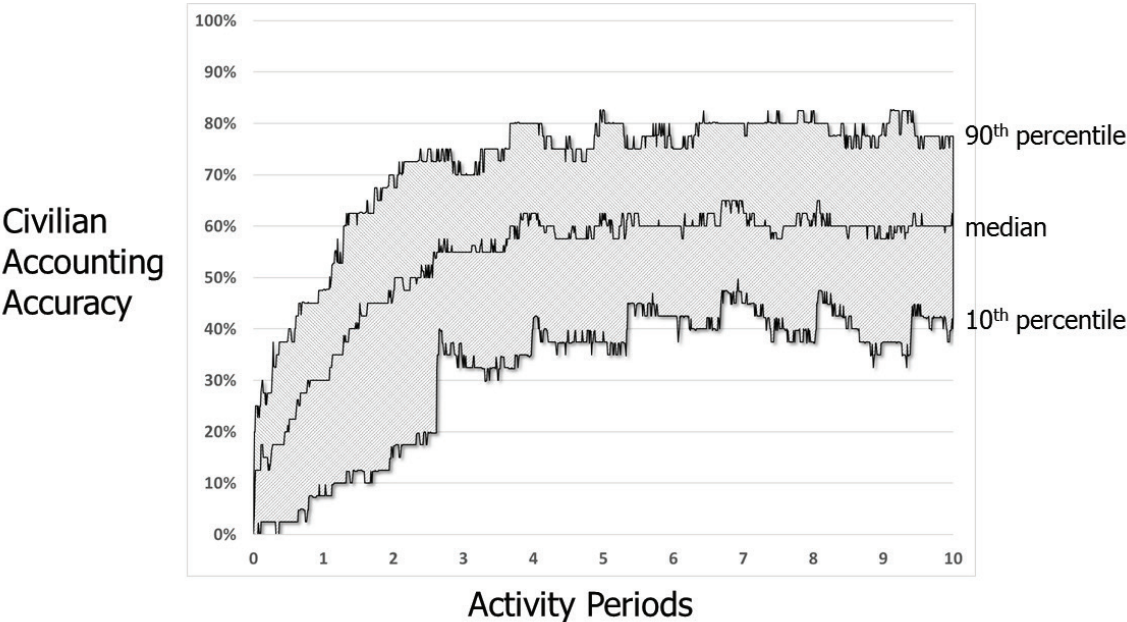
Figure 6.9. Baseline, Clear Weather



Clouds

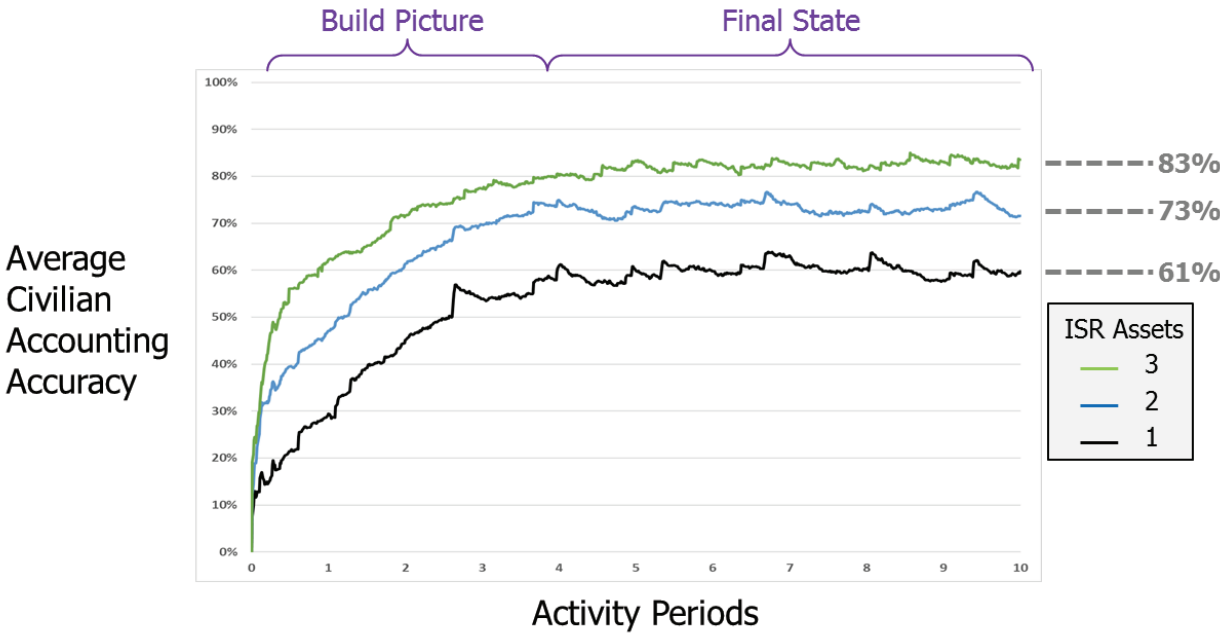
Although we consider many different cases for clouds, a single graph is instructive to demonstrate how degraded weather conditions affect the outcome. Figure 6.10 is the corresponding chart to Figure 6.6, but this time with medium clouds that obscure line of sight. Although the medium cloud cover case we consider assumes 50 percent obscuration, the spotty nature of the clouds allows for multiple opportunities to view civilians moving between buildings, which increases the final accuracy. (If we had instead put a blanket of fog over the compound, reducing visibility by 50 percent, it simply would have reduced all by half.)

Figure 6.10. Baseline, Medium Clouds, One RPA



Compared to the original baseline in Figure 6.6, we see a significantly decreased median and significantly pronounced uncertainty. Under these conditions, we would not expect there ever to be enough confidence to strike. The same pattern of lower, fatter curves emerges for the case of two and three RPA. The final steady state averages for all three are summarized in Figure 6.11.

Figure 6.11. Baseline, Medium Clouds

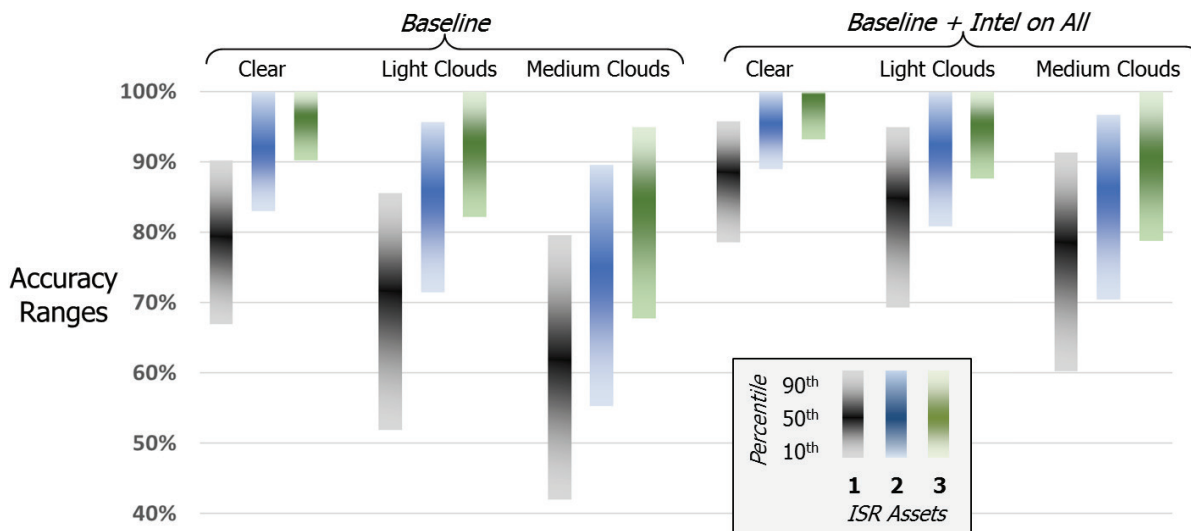


Note that the time to reach steady state with clouds is roughly the same as without. The final operational picture is simply not as good. Again, this is because the timeline is governed by civilian choices, not by the weather.

Summary of Civilian Accounting

In Figure 6.12, we summarize the steady-state civilian accounting accuracy for all cases. (Note that the vertical axis is truncated to highlight the differences.) On the left side, under the “baseline” bracket, the first triplet of bars shows the 10th, 50th, and 90th percentiles for 1, 2, and 3 RPA in the baseline case with clear weather. The second triplet shows the baseline for light clouds, and the third for medium clouds. As discussed above, adding clouds reduces accuracy and increases uncertainty.

Figure 6.12. Civilian Accounting Summary



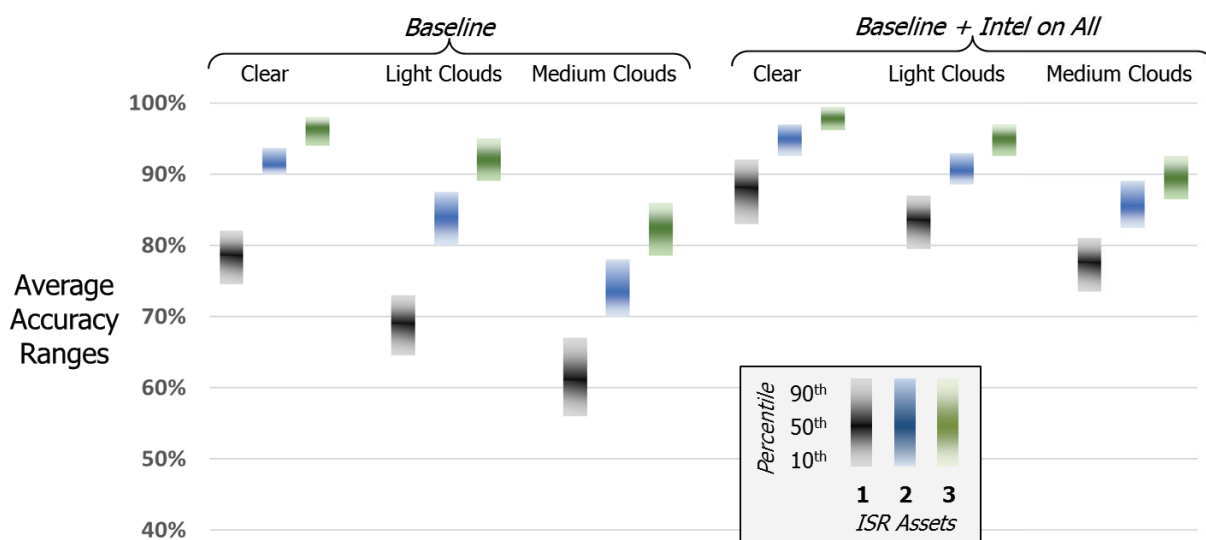
The right side, under the “baseline + intel on all” bracket, shows a separate set of cases. These parallel the baseline, except we allowed for additional, periodic intelligence tips on civilian locations, as discussed above. (We also allowed tips for HVI locations, hence the notation “all,” but only the civilian accounting is reflected in this chart.) The main result here is that the primary effect of fusing FMV with some other form of intelligence is to reduce uncertainty, hence risk. The timing does not change significantly, however, as it remains governed largely by civilian choices.¹²

¹² As indicated earlier the other intelligence tips have the same activity period as the rest of the vignette. If this activity period were significantly shorter, it *would* speed up the time to build the operational picture. If it were longer, it would have little effect on the time—the shortest activity period usually governs the results.

HVI Accounting

Just as we looked at the civilian accounting, we can also look at how well we were able to locate and track the HVI. In this case, we cannot plot percentiles for each moment because there is only one HVI in the vignette, so the answer is either 0 or 1: Either we know the location of the HVI correctly or we do not. However, we can average over the entire steady state period to find what percentage of the time the HVI was correctly located—and from this, we can also find variances. We show these averages in Figure 6.13. (Note that the vertical axis is truncated to highlight the differences.)

Figure 6.13. HVI Location Summary








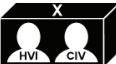





















The results are very similar to the previous case, as we would expect, because the same mechanisms used to track and identify the HVI were used for the civilians. The primary difference is that the uncertainties appear smaller, but this is due to the fact that, as discussed above, this represents the variation of the mean rather than the full variation—which, as discussed above, would technically vary from 0 percent to 100 percent every moment.

Strike Outcomes

Civilian and HVI accounting accuracy are reasonable indicators of the quality of the intelligence, but the outcome of the decision to call in a strike depends on both of these. There are several possible outcomes that could arise, based on the accuracy or inaccuracy of these various counts. We summarize all possible outcomes in Figure 6.14. We will then discuss each of the cases in turn.

Figure 6.14. Strike Outcome Possibilities

				Blue strike decision		
				HVI believed to be at Building X		HVI location unknown
				HVI believed alone	HVI believed not alone	
				Strike	No strike	No strike
						
Ground truth	HVI is at Building X	HVI is alone				
		HVI is not alone				
	Building X is empty	HVI is elsewhere, alone				
		HVI is elsewhere, not alone				
	Building X has civilians	HVI is elsewhere, alone				
		HVI is elsewhere, not alone				

NOTE: CIV = civilian.

The various row and column headers show all possible combinations of ground truth and intelligence calls relative to the decision to strike or not to strike. The colored circles denote the possible outcomes at these intersections of fact and belief.

Green circles represent correct decisions made for the right reasons. The most desirable case is a correct decision to strike the HVI that leads to no civilian casualties; the other case is a correct decision not to strike based on the presence of civilians. While the latter is not a successful strike mission, it is an intelligence success—one that is rarely counted as a “win” because it is difficult to measure outside of the simulation context.

Yellow circles represent an intelligence failure that does not lead to civilian casualties: we correctly observe that there are no civilians present, but we are looking at the wrong building. Red circles represent intelligence failures that *do* lead to civilian casualties: those in which we decide to strike but are mistaken in our belief that there are no civilians present.¹³


¹³ We do not attempt to count the actual number of civilian casualties because, as discussed above, we do not explicitly model the lethality of the strike; we assume all civilians in the building struck may be casualties. For less-restrictive ROE, examining the level of civilian casualties would be of interest.

Ivory circles indicate missed opportunities—situations in which the HVI was alone, but our intelligence was never sufficient to make that call. Light green circles indicate that we ultimately made the correct decision not to strike, but we did so for the wrong reason, usually because we were looking at the wrong building. (We name these cases “broken clocks” because “even a broken clock is right twice a day.”) We include these cases for completeness and to indicate that we did *not* count them as either successes or failures in subsequent figures.

Strike Decision

A successful strike is one that hits the building containing the HVI and also satisfies the restrictive ROE that require zero civilian casualties. We follow the same coloring scheme in these figures as in Figure 6.14 above, with the exception that we single out the worst possible combination of outcomes in dark red: civilians are hit and the HVI escapes. (These are the lower two red circles in Figure 6.14.) Figure 6.15 shows the probabilities of these categories of outcome for the baseline case with clear weather.


Figure 6.15. Strike Outcomes, Baseline, Clear Weather



The green bar shows that the probability that a strike decision results in a clean, successful hit. For a single ISR asset, the likelihood of success is surprisingly low at 56 percent. These figures are lower than either the civilian accounting and HVI location accuracy because *both* must be correct. The most likely mistake is the bright red bar indicating that there are civilians present, which reflects that any unaccounted-for civilian runs the risk of being in the wrong place at the wrong time. There is also about 9 percent chance that the HVI has been mislocated, which could lead to striking an empty building or, in the worst case, civilians. As expected, the green bar success rate jumps up—to 83 percent and 93 percent respectively, in this case—when a second and third asset can be employed.

The situation looks even bleaker when there is cloud cover. Figure 6.16 shows the same strike outcomes for the baseline case, but when the confounding factor of cloud cover is added to the mix.


Figure 6.16. Strike Outcomes, Baseline, Light Clouds



We see here that even light clouds make any strike decision too risky under these ROE. (The result with heavy clouds would be even worse.) The uncertainty in the civilian accuracy and HVI location are compounded to make a bad outcome unacceptably likely, even with three RPA providing intelligence. It is worth noting, however, that the operational picture is still built in the same time and still reaches a steady state—so being able to build a steady-state picture is not enough. The problem is that, without knowledge of ground truth, the sensor operator knows only that a steady state has been reached, and not how good the operational picture actually is.

Fusing this information with intelligence tips from another source, however, can make a difference. Figure 6.17 shows the same case, but with intelligence on HVI and civilian locations from intelligence tips.


Figure 6.17. Strike Outcomes, Baseline Plus Other Intelligence, Light Clouds



We see here that adding in these intelligence tips effectively counteracts the deleterious effects of cloud cover. This makes sense, since the primary advantage of these tips is that they are not blocked by line of sight—and thus are not affected by the cloud cover. However, it is worth noting that they do not counteract the effects completely, and the advantage is diminished with additional RPA. Because these tips are independent of any other collection, they do not always provide new information. The better the existing intelligence picture, the more likely that any uncorrelated intelligence will be redundant. The best way to increase the success rate would be more targeted collection, but that would require a more complicated model than we use here.

In Figure 6.18, we show the likelihood of a successful strike without civilian casualties for all 27 cases we considered. In this stacked-column chart, the black bar shows the result for a single RPA, while the stacked blue and green bars show the additional probability of success attributable to adding a second and third RPA. Here we distinguish between receiving intelligence tips on the HVI only and receiving them for the HVI and all civilians.

Figure 6.18. Strike Outcomes, All Cases




Three features are worth noting. First, the second cluster of three bars looks much like the first. In other words, receiving extra intelligence tips on the HVI alone does little to increase the likelihood of a successful strike without civilian casualties. This makes sense because most mistakes occur when the civilian accounting is inaccurate, and providing extra intelligence tips for the HVI does not affect that count. Second, as we have seen repeatedly, there are diminishing returns with each additional asset—but a second or third asset is almost always required to get above the minimum thresholds one would expect based on restrictive ROE. Finally, medium-dense cloud cover presents unacceptable risks in virtually any case under these ROE. Additional intelligence sources or collection methods would be required to overcome this deficit.

Decision Not to Strike

In addition to the results of strike decisions, we can also assess the decision *not* to strike. Figure 6.19 shows the chance that the decision not to strike was made correctly. This shows the probability that the HVI was correctly located and correctly known not to be alone, out of all cases in which the HVI was actually not alone.


Figure 6.19. Nonstrike Outcomes, All Cases



The results are qualitatively similar to that of the strike decision, but with one important difference: here, the value of additional intelligence tips can be seen for both HVI and civilians. In other words, to get the most benefit, intelligence tips are needed on all.

The counterpart to this are missed opportunities. In this case, we count the number of times the HVI was actually alone, and could have been struck, but our intelligence estimate either missed that fact or we were looking at the wrong building entirely. Figure 6.20 shows these results.


Figure 6.20. Missed Opportunities



Whereas intelligence tips are required on both the HVI and civilians to make a correct “no-go” decision, having intelligence tips on the HVI alone is what mitigates the risk of a missed opportunity. The risk of a missed opportunity can be quite high, but that is an appropriate and expected outcome whenever information is poor and ROE are strict.

Timing

Finally, we look at one potential indicator which turns out to be an unexpectedly poor indicator of ISR value in this case: the time one must wait until the first opportunity to strike. Intuitively, one might expect that having additional ISR assets or intelligence tips would shorten this time, but this proves not to be the case. Figure 6.21 indicates what the distribution of strike times looks like for all 27 cases.



This complicated graph requires some explanation, but the ultimate message is fairly simple. The vertical axis shows the time until intelligence indicates there is a strike opportunity. The bars show the median, 25th, and 75th percentiles. The number of ISR assets, the cloud cover, and the possible addition of other intelligence sources are shown in the labeling of the cases. The red lines show the ground truth of when the actual time to first strike occurs in the model. (The time varies because it is a stochastic model.)

The main thing one sees is simply that there is large variance and little pattern. As indicated above, this is because the timing of the scenario is governed by enemy and the civilian activity, not by our choices. A closer inspection, however, reveals a curious feature: The time until intelligence indicates there is a strike opportunity is almost always less than the actual time for a genuine strike opportunity. This is because the timing errors are largely one-sided here. In all cases, one is most likely to strike correctly, so mistakes that might happen after that genuine first

opportunity usually do not matter because the strike has already occurred. On the other hand, any mistakes that might occur before the genuine first opportunity cause a premature strike based on overconfident intelligence estimates.

Time to strike is therefore not just a poor metric in this case, but in some cases it is a counterindicator to success. One of the advantages of being able to simulate many cases, as we have done here, is that we can test the usefulness of different indicators for ISR assessment.

Conclusions

Through this specific analysis of a stylized intelligence problem, we have demonstrated how computer simulations can be used to support notional assessments about the value of adding RPA for the purposes of increasing the accuracy of civilian and HVI accounting. In this case, the assessment suggests additional value for each airborne ISR asset, but with diminishing returns for each. We have also shown that the key indicator governing the number of RPA required for the missions (in this case, at least) is set by the *threshold of risk for collateral damage*. If, for example, 90 percent confidence of no civilian casualties is required for this complex scenario, two or three will be required, depending on the specific case, although changing ROE or adding other forms of intelligence could in theory reduce that. We have also shown that time to strike is a poor indicator for ISR assessments, and that poor weather adds uncertainty that only a separate form of intelligence can mitigate.

We note that the results shown here are, of course, specific to the vignette. If we had chosen a different laydown for the compound, for example, this would change some of the details of the results. However, we believe this would be a minor effect, and the primary results described here would remain.¹⁴ We also did not change the population of the compound over time; doing so could increase the uncertainty. Further, we did not account for the possibility of adversary denial and deception. Finally, we also made the heroic assumption that the PED process introduced no additional errors and that information from different RPA was always correlated correctly. Relaxing these assumptions would make the results less favorable than shown.

There are a number of additional ways in which the vignette presented could be usefully varied, including

- increasing the order of magnitude for the number of civilians to simulate an urban environment
- giving greater consideration to the day-night cycle, including shifting levels of illumination and varying resolution of sensors

¹⁴ There is a clear axis of orientation in the compound design we chose, running northwest to southwest, which allows for the best viewing angles. Some of the jumps and choppiness visible in the results are due to this choice. However, unless the buildings were perfectly spaced and circular, there will always be some viewing angles that are more favorable than others, so we believe this kind of irregularity is actually an important feature of the model.

- including personnel exiting and entering the compound (not just moving between buildings on the compound as we do currently)
- varying the layout of the compound, for example making it larger or smaller and with additional or fewer buildings
- reducing visibility so certain individuals can be detected but not identified
- adding specificity about additional sources of intelligence information and RPA capabilities.

There are also some ways in which this simulation concept could be expanded to consider ISR support in different roles. For example, a simulation of this nature could be adapted to explore the amount of activity needed to be observed by a particular number of ISR assets (potentially collection different forms of information) to correctly identify indicators of impending adversary activity and provide sufficient warning time to allow friendly forces time to make decisions.

In this chapter, we have broadly shown that computer simulation can be used to provide insight into the expected outcomes of scenarios, investigate risks and uncertainties, and identify likely good indicators for ISR assessment. In particular, agent-based modeling is useful because it allows us to compare perceptions to ground truth to determine the value of intelligence in situations when there are “no go” decisions or to find missed opportunities. These insights can be important for establishing reasonable decision points and for asking the right questions when deciding to allocate ISR assets.

7. Step 5: Apply Results

The final step in any ISR assessment process is applying the results for informing ISR resourcing and employment decisions. If assessment steps 1–4 have been conducted effectively, the assessment product will be relevant and accurate, and it should not be difficult to apply the resulting information to ISR resourcing and employment decisionmaking.

However, as our expert interviews and site-visit observations revealed, ISR assessments are not always used or relevant to decisions being made within or about the USAF ISR enterprise. This is because assessments do not always address the most relevant (or correct) questions, provide implications for upcoming resourcing and employment decisions, and reach all audiences who would benefit from the information. These issues are largely the result of problems related to steps 1–4 and discussed earlier in this report. However, there are a few additional steps that can be taken beyond what has already been mentioned in order to help ensure that ISR assessments are accurate and informative. In this way, assessments themselves can have value.

In this chapter, we discuss two focus areas for USAF to concentrate on for improving progress towards the last assessment step:

- supporting decentralized assessment activities by providing brief, overarching assessments guidance and harnessing expertise in data science and computer programming
- ensuring that procedures are in place to communicate the results of ISR assessments to appropriate audiences.

Our insights and findings in this chapter are drawn from expert interviews and observations made during site visits, as well as from our examination of assessment practices in other industries and organizations (see Appendix C).

Supporting Decentralized Assessments

ISR assessments are generally conducted within the organizations or teams that will use the information to support specific resourcing and employment decisions. Within USAF, these organizations can range from USAF Headquarters to individual teams of airmen, flights, and squadrons. Sometimes organizations use assessments as a means for communication with others in addition to informing specific resourcing and employment decisions. Sometimes, airmen or teams perform what amounts to assessment without having these activities recognized as such (including among themselves). As we have emphasized elsewhere, the local nature of some assessments is a primary reason why the analytic shell we have proposed here must be reconsidered and adapted for particular teams or organizations. In the case of Joint communities

and the IC, which we did not examine in great detail or at all, our approach may not be relevant or useful.

Within the USAF airborne ISR community, assessments currently lack consistency in approach and understanding of terms. For this reason, a centralized ISR assessments capability might seem tempting. However, there are several reasons why enabling a local, decentralized assessments capability might be better, especially one that is supported with additional guidance, access to data, and use of data science and computer programming skills.

First, on-site SMEs—whether they are at headquarters and are interested in broad, strategic questions or are located within small units and have very tactical needs—have the most direct knowledge of the problems they are attempting to help resolve through the use of assessments. In many cases, SMEs interviewed for this research literally “lived” the bad situations they were trying to resolve through assessments. Having an in-depth understanding of the problem(s) and the information sources from which data can be drawn is of great value for assessments.

Although some might argue that being too close to the problem might cause a failure to “see the forest for the trees,” our research team did not observe this problem. On the contrary, most of the challenges being examined through assessment required deep site-based expertise to design an assessment. Many of the assessments activities we observed were not doctrinal requirements, but were borne out of a specific need, which might include needing to stretch existing resources further or digging deeper into a problematic issue that has arisen multiple times. In addition, there are a number of details that must be understood, sometimes minutely, that may only be readily observable and understandable within an individual organization or team. For example, someone distant from a theater of operations might focus on the number of airborne platforms and PED analysts available at any given time, neglecting the reliance of platforms on the availability of pilots and ground logistics and the fact that there may be combinations of specific PED personnel types required for mission execution. Further, on-site SMEs may also best understand how the required actions implied by assessments can best be implemented. By contrast, if assessments are conducted off-site, the SMEs who need to implement the results may not fully understand the findings.

Importantly, it should be recognized that having a decentralized assessments capability – that is, one which enables many different organizations to conduct analyses to meet their needs – does not mean that there is not a need for some types of centralized assessments. There are many ISR assessment questions that are broad in nature. For example, questions guiding acquisition decisions or comparing demands on PED across different theaters. These will necessarily require oversight and execution by USAF organizations, such as Analysis, Assessments and Lessons Learned, with viewpoints broader than any particular mission or operational unit.

The second reason why decentralized assessments may be preferred is that organizations and teams are more likely to both understand and accept the results of assessments that they conduct themselves. It is important to consider that organizations will probably be more responsive to acting upon an assessment if there is not the perception or impression that an outside

organization is making judgements upon them. In our limited interviews across the enterprise, this appeared to be a fairly common concern with asking for help with assessment from an outside organization. Of course, there will be times when assessments covering missions of multiple organizations are needed. In many cases, there will be an appropriate umbrella organization that could handle this. Otherwise, leadership of different organizations will have to coordinate an effort and ensure that results are disseminated in a way that they are useful to different components.

The third reason in favor of decentralized assessments is that a large amount of manpower is needed to conduct assessments for the entire enterprise, and standing up a new organization to do this would require hiring a large number of personnel. As we have suggested, there are many assessment questions across the enterprise, and the requirements for answering them are varied such that there are limited—if any—opportunities to achieve an economy-of-scale situation. Providing analysis to all the organizations and teams that require it in a timely fashion would require a huge investment in manpower if assessments capability were to be housed within one organization.

Finally, a centralized approach within USAF would fail to recognize the fact that not only does the Service require assessments, it must also organize, train, and equip airmen to conduct assessments as they serve the various combatant commands. Focusing only on Service assessment needs in a centralized approach would not serve USAF's interests since much or all of the data and analyses the service-centric assessments (e.g., in support of long-term force sizing) would rely on are generated by airmen serving in different theaters around the world. If those data and other assessments don't meet the standards of the centralized assessment organization, then it will be impossible (or nearly so) to conduct many types of assessment for the benefit of the Service.

Although analysts within the AOCs and DCGS, ACC, 25th AF, NASIC, AFSOC, and elsewhere are already tasked with many things, our research team did find that many people are still managing to conduct assessments of one form or another, some of them formally required,¹⁵ but others simply in response to emerging problems that need addressing. There is a large workforce of full- and part-time assessments specialists already in existence. We argue that it will be most productive for USAF to provide these existing personnel conducting assessments with broad guidance and access to databases and tools to enable their activities. This also permits USAF to leverage airmen in various positions that have skills in data science and coding that may or may not be relevant to their primary jobs.

Naturally, leaving assessments to individual organizations requires ensuring that people have the proper guidance, skills, tools, and time to conduct these activities. Brief, overarching guidance in the form of a short document, briefing, or infographic that can be posted and

¹⁵ That is, part of documented processes conducted by analysts sitting in billets dedicated to assessments, though these are not necessarily always filled due to budgetary or other constraints.

periodically updated will be helpful for standardizing some aspects of assessments procedures as well as data gathering and analysis. This guidance can include the basic assessment steps and questions (such as those suggested in this report), a lexicon of assessment terms, example approaches for identifying and analyzing assessment indicators, an ISR taxonomy, and a list of database, simulation, and tools available to aid assessment analysts.

Infusing the right skills for conducting assessments into the ISR enterprise will likely involve multiple efforts, including recruitment and retention strategies, as well as career development and training. These topics are probably well-suited for longer-term planning with additional analysis required to flesh out the implications of different approaches, not only for ISR assessment but for other aspects of the ISR enterprise and USAF. For this reason, we do not address them in more detail here.

In the shorter term, it would help to give personnel access to, or to build expertise in, data science and computer programming. In previous chapters we have shown that collecting, storing, analyzing, and even modeling relevant data can greatly enhance the relevance and accuracy of ISR assessments. USAF should consider several approaches to making these capabilities available, whether individually or in combination. First, USAF can leverage airmen who already possess these skills. Some already exist within the ISR community, but may not be performing tasks that leverage these technical skills. Leaders within the enterprise should identify airmen with these skills and seek opportunities for allowing them to devote time to solving problems having to do with data access or analysis. In addition to already being within the enterprise, these airmen also will have a good sense of the problem they are trying to solve and how to do it from a process perspective because they are embedded within the affected organization.

Of course, ISR airmen have a variety of tasks to perform and thus should not be overburdened with too much additional work. Although they might be relied upon for small data science and computer programming projects outside of their named position, bigger jobs such as integrating several data sources, preparing for a large database overhaul, or expanding a prototype tool should be given to scientists within USAF, contracted to external organizations, or both. Outsourced activities should always keep the ISR airmen that the improvements are intended to serve in the loop to ensure that solutions serve the intended purposes.

The last factor to consider is time. ISR airmen have many demands on their time. Effective assessments require analysts with time to put into them. Since many analysts are already conducting assessments, additional guidance, data access, and other enablers might in fact begin to save time. However, the recognition by leadership that assessments are beneficial and inherent to any organization's success will be helpful in carving out time to conduct them. This will mean that something else will have to be left out of the workday. Perhaps assessments will be able to inform which activities have the potential for greater efficiency.

Communicating Assessment Results

Assessments that are not delivered in the right format to the right people at the right time are not useful. Doing so is a challenge for some parts of the USAF ISR enterprise. This is because of (1) a lack of consistency in guidance for how assessments results should be communicated under different circumstances, including how to derive implications for specific ISR resourcing and employment decisions, and (2) limited visibility into and accessibility of some assessment activities.

Getting Information to the Right People, in the Right Format, at the Right Time

Some organizations and teams may not recognize their analytic activities as assessments that may be of use in providing feedback to others, or may not elect to share this information for various reasons. The enterprise should consider examining which organizations should share assessments, whether to support resourcing and employment decisions at different echelons or to improve assessment processes in different parts of the enterprise. Specific assessment information needs should also be communicated so that analysts know what is needed, for what purpose, and within what timescale or geographic scope, as alluded to in assessment step 1.

Having particular assessment product templates, already in use in parts of the community, will help analysts make consistent presentations of ISR assessments. In addition, including a “product slate” of assessment types in the overarching guidance described above may also be useful. For example, this could clarify when briefing slide “storyboards” should be used as assessment products, and when they should not. Storyboards are briefings in which a series of closely connected events are described in a “play-by-play” approach to highlight what happened under a particular circumstance (usually in a positive light). These are perhaps helpful for assessing what went particularly right or wrong in an ISR mission in a way that is compelling for senior leaders. This is not an effective type of product, however, for informing force-sizing decisions, or even for deciding how to employ ISR platforms and PED in the next ATO. In these examples, an assessment might take the form of a brief written report with graphs summarizing data analysis and explanations of factors that require consideration in the interpretation of results. Assessments working to compile other assessments in order to discern similarities and differences between organizations (e.g., in order to better understand the range of missions PED supports) might develop a template to help ensure an “apples-to-apples” comparison can be made.

Although doctrine dictates some pathways of communication (e.g., within the AOC) and practice has strengthened other pathways (e.g., between the AOC and DART), this communication does not always take place and inhibits sharing of assessment information. It will take consistent leadership around the ISR enterprise to encourage assessments information being shared.

The more teams and organizations that openly share assessments, the more likely it is that there can be willing participation (everyone is sharing information, not just a few sharing with many). This could also help facilitate the sharing of new methods and data sources, as well as ensure that consistency as specified in overarching guidelines is maintained. Identifying the right medium for doing so will be important; if analysts find their experience using a common SharePoint site, for example, to be frustrating, they will be less likely to take the time to share their work.

Conclusion and Suggested USAF Actions to Improve Assessment Step 5

This chapter has suggested that ISR assessments can be more effectively communicated by doing the following, based on our site observations, interviews, and examination of other assessment practices detailed in Appendix C:

- adopting a consistent approach to ISR assessments, such as the cost-benefit approach described in steps 1–4
- supporting decentralized assessment via guidance and access to data science and computer programming skills
- enabling better communication of assessments by improving guidance on assessment formats and ensuring that communication of assessments is prioritized by leadership.

Now that we have examined each of the five assessment steps in detail, this report turns to overarching conclusions and recommendations to USAF for enhancing ISR assessment capability.

8. Improving the USAF's ISR Assessments Capability: Conclusions and Recommendations

Conclusions

The USAF ISR enterprise conducts missions with important strategic, operational, and tactical implications every day around the world. ISR capabilities are key enablers that are in high demand: Users invariably request more, not less, ISR. While the enterprise helps close knowledge gaps for USAF, Joint Community, and other users of airborne intelligence information, there is a growing urgency to close what could be argued is its own most important knowledge gap: the ability to consistently, systematically, and accurately assess ISR contributions, employment, and operational efficiency.

Conducting ISR assessments is not simply about “bean counting”; that is, understanding how many assets and analysts of what type are being employed for how long. This report has suggested that assessment is about contextualizing *how* ISR were used, in what *amount*, and for what *purpose*. With this information in hand the enterprise can make—and support others in making—better choices about ISR acquisition, force structure, allocation, distribution, and employment. Such an approach to assessment would apply equally to different parts of the ISR enterprise, ranging from individual missions to weeks-long operations to years of employment within an operation or against a strategic mission set. The enterprise and its partners are not yet ready, however, to conduct the types of assessments described in this report. Building such a capability requires changes to guidance, management of assessments, empowerment of individual organizations and teams to conduct assessments, data and databases, feedback mechanisms and other factors. Importantly, the approach suggested in this research is a framework intended for the USAF airborne ISR community, and aspects of it may not apply to other parts of the Joint Community or IC. For example, this approach would require substantial review and filling in with appropriate analytic methodologies before it could be potentially considered for the review of major intelligence failures.

Not all of the required changes are within the USAF ISR community's power to implement. As discussed in Chapters 1 and 2, the ISR enterprise is a complex web of organizations working to address different concerns at different echelons and timescales. However, the USAF ISR community can do much to improve ISR assessments. The recommendations below represent the best ways forward for the USAF A2 office to improve ISR assessments based on the research team's interviews, analyses, and reviews of other studies and practices. These recommendations should be considered as inputs to a broader USAF process aimed at enhancing ISR assessments, and weighed alongside inputs from other organizations within and outside the USAF that have

also rigorously examined this topic. These recommendations should also be considered alongside the issue of limitations in manning for ISR assessments, which was not examined in this study.

Recommendations for the USAF to Improve ISR Assessments

The following recommendations articulate steps forward in policy oversight; guidance; and tactics, techniques, and procedures. While the recommendations are designed to help HQ AF/A2, who sponsored the project, they should be relevant to other HQ USAF staffs, ACC, 25th AF, the 480th ISR Wing, and other organizations. Implementing many of these recommendations will require communication with National and Joint partners, and in some cases will require their partnership, or even leadership in affecting necessary changes. These recommendations are collectively derived from our examination of the five assessment steps in the preceding chapters. Like the insights discussed therein, these recommendations derive from information collected during interviews and site visits, guidance and other documents, use of operational databases and analysis of operational data, an ISR simulation, and examination of assessment approaches used by other organizations and within other industries.

1. **Adopt and provide guidance to implement a consistent approach to ISR assessment.** Guidance should cover what constitutes an assessment (see example questions and assessment objectives in Chapter 1), the basic steps for conducting assessments (such as the five-step approach outlined in Chapter 1), and broad key questions about benefits and costs that should be answered as part of any assessment (such as those detailed in Chapter 2). This guidance should integrate broadly relevant lessons from ongoing and recent studies and other efforts focused on improving USAF and Joint ISR assessments. It would be beneficial for this guidance to be informed by and coordinated with assessment organizations within the Joint Community and IC, including the Joint Staff and COCOM staffs. This guidance should also establish or reaffirm the expectation that organizations will use assessments to support their own decisionmaking about ISR (e.g., within an ATO) as well as Joint or IC decisionmaking under certain circumstances (e.g., an AOC feeding assessments to a COCOM for planning purposes). It might be useful to disseminate guidance in a short format like an information graphic that can be widely distributed and displayed.
2. **Direct the development and adoption of a common lexicon for use in assessments, requirements articulation, collection management, and PED.** The lexicon should include a granular taxonomy of ISR roles (such as the one described in Chapter 3) and clearly defined terms such as *benefit*, *cost*, and *indicator*. This lexicon should be coordinated with the Joint Community, in some cases leveraging or providing a cross-walk to existing terms and definitions.¹⁶
3. **Guide and oversee short-term and long-term plans for improving the ISR data environment.** This includes increasing recorded data fidelity and discoverability for the

¹⁶ Note that some of the terms; definitions; and tactics, techniques, and procedures needed for ISR assessments are not available in Joint or Air Force doctrine. Developing these in coordination with the broader Joint and intelligence will be very important so that more than one lexicon is not developed.

purposes of assessment. A two-phased approach, explained in detail in Chapter 4, requires:

- a. In the short-term: Recording the granular ISR roles within requirements databases, the ATO, UNICORN, and other databases; and making limited adjustments to UNICORN to add more data fields, make the data queryable by additional types, and implement backend checks to scrub existing data.
 - b. In the long-term (where implementable): Implementing automatic machine-to-machine transmission of data directly from aircraft and sensors to data repositories, tagging collection to produce searchable metadata, storing historical data and making it easily retrievable, and creating a new, indexed database for ISR assessments that enables data discoverability and analysis. This should also include an ability to track intelligence citations, which will require setting and enforcing standards for when and how to use citations.
4. **Enhance airmen skills and resources to perform assessments.** A more analytically rigorous and sophisticated assessment capability will require analysts with comparable skills who can regularly execute a variety of analyses and assessments needed to ascertain ISR's value and drive future decisionmaking. USAF should consider several options to bolster the analytic force, such as those discussed here.¹⁷ Emphasize data science and computer programming for recruiting and developing ISR airmen; examine whether the Reserve Component can be used to enhance ISR assessments capabilities. Take advantage of computer coding and data science expertise resident within USAF by encouraging local commanders to select computer and data science savvy airmen for temporary task reassignments, and finding ways to engage more existing USAF scientists in the ISR community (e.g., through a program that organizes opportunities for visits and temporary assignments). This could include further taking advantage of partnership opportunities with USAF and DoD organizations designed to harness and explore innovative technologies. Explore whether more contractors could be used for short-term outsourcing of coding or database development. Hire temporary consultants from technology companies to advise on how to best proceed with technological innovation plans. These topics are discussed in additional detail in Chapter 4.
 5. Coordinate with IC, Joint, and Industry partners to **identify opportunities for gathering feedback** about USAF ISR products. As described in Chapter 4, two types of feedback to institute include active (e.g., an automated window asking users to select a score) and passive (e.g., statistics on how often products are accessed) approaches. The Analysis, Assessments and Lessons Learned community could be an important partner in this effort.
 6. Continue to **weigh the advisability of changing PED force presentation** and the impacts this would have on assessments. As discussed in Chapter 7, PED teams that are organized around collection platforms might provide assessments that are best suited for examining platform execution in a narrow context. Shifting PED analyst operations to more a "problem-centric" approach might enable these analysts to better communicate

¹⁷ The service should weigh the pros and cons of these, and perhaps other options as well before determining any specific course of action.

their findings in the context of progress towards solving problems, which would directly benefit the ability to conduct assessments.

7. **Identify cases when simulations are needed to support assessments.** Simulations may be especially useful at developing tactics, techniques, and procedures for informing ISR employment in different contexts for which operational real-world data may not be available or otherwise obtainable, situations where “ground truth” in the real-world is unreliable (either because too little data has been collected or there are many confounding factors), and cases where risk must be understood in mathematical detail, such as exploring options for minimizing collateral damage. This is discussed in more detail in Chapter 6.
8. **Continually refine and update processes and guidance for ISR assessments.** Host regular, periodic discussion forums to refine guidance as necessary, update the ISR roles taxonomy as needed, as well as identify the latest assessments challenges and possible areas of “low hanging fruit” to target. **This last recommendation is one that will be needed in order to ensure that the previous seven recommendations move forward.** These forums could also be used to communicate and update a list of ISR assessment priorities for the Air Staff related to emerging (e.g., ISR in new operating environments, pending investment or divestment decisions) and enduring issues (e.g., ISR benefits in support of long-term national strategic priorities) so that communities within the enterprise can identify existing assessments and/or datasets within their purview that are relevant and pass them to an appropriate point of contact.

Although some of the recommendations may pose budgetary and organizational challenges or require cultural shifts within USAF, the potential benefits of having better-managed ISR resources are surely worth the effort.

Appendix A. HVI Simulation Analysis Results

Here we show tables of the results. Table A.1 shows the civilian accounting accuracy for all relevant cases for different numbers of FMV assets and different weather conditions.

Table A.1. Civilian Accounting Accuracy in the Steady State—Averages

Vignette	Environment	Average Accuracy (%)			
		Assets:	1	2	3
Baseline	Clear		78.8	91.6	96.0
	Light Clouds		69.9	84.7	91.4
	Medium Clouds		61.2	73.4	82.6
Baseline + Intel on Civilians	Clear		87.5	95.1	97.5
	Light Clouds		82.7	91.0	94.8
	Medium Clouds		77.0	85.1	89.6

NOTE: The case of additional intelligence on the HVI is not shown because this does not affect the civilian count.

Table A.2. Civilian Accounting Accuracy in the Steady State—Percentiles

Vignette	Environment	Percentiles:	Accuracy (%)								
			1 Asset			2 Assets			3 Assets		
			10th	50th	90th	10th	50th	90th	10th	50th	90th
Baseline	Clear		66.9	79.3	90.1	83.0	92.1	99.9	90.2	96.5	100
	Light Clouds		51.9	71.5	85.5	71.5	85.9	95.7	82.2	92.8	99.9
	Medium Clouds		42.4	61.9	79.6	56.1	75.9	90.0	67.5	84.4	95.0
Baseline + Intel on Civilians	Clear		78.1	88.3	95.8	88.9	95.5	100	93.2	99.5	100
	Light Clouds		68.6	84.5	94.8	81.0	92.4	99.9	87.7	95.3	100
	Medium Clouds		60.7	78.4	91.4	70.7	86.7	96.7	78.1	90.6	99.9

NOTE: The case of additional intelligence on the HVI is not shown because this does not affect the civilian count.

Table A.3. HVI Location Accuracy in the Steady State—Percentiles

Vignette	Environment	Percentiles:	Accuracy (%)								
			1 Asset			2 Assets			3 Assets		
			10th	50th	90th	10th	50th	90th	10th	50th	90th
Baseline	Clear		74.5	78.5	82.0	90.0	91.3	93.7	94.0	96.5	98.0
	Light Clouds		64.5	69.0	73.0	80.0	84.0	87.5	89.0	92.0	95.0
	Medium Clouds		56.0	61.0	67.0	70.0	73.5	78.0	78.5	82.5	86.0
Baseline + Intel on HVI	Clear		83.0	88.0	92.0	92.5	95.0	97.0	96.1	97.8	99.4
	Light Clouds		79.5	83.5	87.0	88.5	90.5	93.0	92.5	95.0	97.0
	Medium Clouds		73.5	77.5	81.0	82.5	85.5	89.0	86.5	89.5	92.5

NOTE: The case of additional intelligence on civilians is not shown because this does not affect the HVI location.

Table A.4. Strike Outcomes in the Steady State

Vignette	Environment	Probability (%)											
		1 Asset				2 Assets				3 Assets			
		H	C	H+C	E	H	C	H+C	E	H	C	H+C	E
Baseline	Clear	55.9	3.1	35.2	5.8	82.5	0.5	15.4	1.5	93.2	0.0	6.5	0.3
	Light Clouds	42.6	9.6	37.7	10.1	66.8	2.1	26.9	4.3	80.0	1.0	18.0	1.1
	Medium Clouds	27.8	24.2	36.2	11.8	48.3	9.7	32.9	9.1	59.9	4.5	30.1	5.5
Baseline + Intel on HVI	Clear	57.3	1.3	39.6	1.7	82.5	0.2	16.5	0.8	93.5	0.0	6.5	0.0
	Light Clouds	43.9	5.8	47.6	2.7	70.9	1.3	25.3	2.4	81.6	0.5	17.2	0.6
	Medium Clouds	34.4	10.2	50.3	5.1	50.5	5.0	40.7	3.8	62.7	2.6	32.0	2.8
Baseline + Intel on Civilians	Clear	73.4	1.0	22.5	3.1	91.5	0.2	8.0	0.3	95.4	0.0	4.6	0.0
	Light Clouds	63.6	3.9	27.0	5.5	79.5	0.9	17.1	2.6	90.6	0.2	8.2	0.9
	Medium Clouds	53.1	7.2	29.5	10.1	66.8	3.8	25.2	4.2	77.3	1.5	18.8	2.4

NOTE: H = HVI only; C = civilians only; H + C = HVI and civilians; E = empty building.

Table A.5. Nonstrike Outcomes in the Steady State


Vignette	Environment	Probability (%)					
		1 Asset		2 Assets		3 Assets	
		Missed Opportunity	Correct "No-Go"	Missed Opportunity	Correct "No-Go"	Missed Opportunity	Correct "No-Go"
Baseline	Clear	35.6	56.5	11.3	85.4	4.8	94.8
	Light Clouds	47.3	40.6	25.4	70.1	14.6	83.4
	Medium Clouds	64.6	26.9	41.6	47.7	29.9	64.9
Baseline + Intel on HVI	Clear	17.1	68.9	8.3	89.6	1.6	96.3
	Light Clouds	25.2	55.6	12.9	78.4	7.7	88.9
	Medium Clouds	31.2	41.5	20.0	62.2	14.6	73.1
Baseline + Intel on Civilians	Clear	18.0	75.5	5.4	91.9	3.0	96.5
	Light Clouds	24.0	68.4	12.4	81.8	6.4	91.1
	Medium Clouds	33.8	57.5	21.2	71.5	14.2	80.9

Appendix B. Mathematical Model of Fixed Point Security

Overview

In this appendix, we examine a mathematical model of fixed-point security to illustrate some basic aspects of ISR allocation tradeoffs. We treat it as a simple, random search for an intruder entering an area of size M , using a sensor with FOV a , that is being pushed like a broom by an aircraft moving overhead at speed v . This model assumes no prior knowledge of enemy timing, location, path, or intent, and looks only at the “find” portion of the targeting process—it does not attempt to capture the additional value of ISR to track and engage subsequently. The main indicator we use here is the confidence that we can detect an enemy intruder within the time it would take him to reach the base we are trying to protect, what we call the “penetration time” for this perimeter. This search problem was first investigated mathematically by Koopman in 1946 in the context of a naval operations over a wide area; we adapt here it to the context of fixed-point security.¹⁸ Figure B.1 shows the overall situation.

Figure B.1. Search Area



For a random search, of this area, the cumulative probability of detection P_D is found to be,

$$P_D(t) = 1 - (1 - q\alpha)e^{-\frac{q}{T}t},$$

where t is the search time, T is the average revisit time based on the velocity and search area, q is the instantaneous probability of detection if the target is within the FOV, and α is the likelihood

¹⁸ Bernard O. Koopman, “Search and Screening,” Operations Evaluation Group Report No. 56, U.S. Navy, 1946.

that the target is within the FOV at any given time—the fraction of the protected area within the FOV ($\alpha = \frac{a}{M}$). The derivation is based on taking the continuum limit of repeated snapshots.


Single Point

For multiple, identical ISR assets, the cumulative probability of detection within a given penetration time τ is given by the related formula

$$P_n(\tau) = 1 - (1 - qn\alpha)e^{-\frac{qn\tau}{T}}$$

We show the results in Figure B.2 for the illustrative case where $\alpha = 0.1$, $T = 1$, and $q=1$.

Figure B.2. Confidence of Detection



We see that the marginal increase in confidence of detection from adding additional aircraft ($\Delta P_1, \Delta P_2, \Delta P_3, \dots$) decreases with each aircraft, no matter the penetration time associated with protecting that particular fixed point. Although we show only one example, here the results are robust. The fact of diminishing returns results from the exponential form of the relationship and is a very general feature. The actual shape of the curve will depend on the sensor FOV, the revisit interval, and probability of detection if within sensor FOV.


Allocation to Multiple Points

If n identical ISR assets are available to protect m points ($n \geq m$), and the k th point has coverage α_k , revisit time T_k , and penetration time τ_k , the P_D for any given point is a function of the number of assets deployed n_k to that point:

$$P_D(n_k) = 1 - (1 - qn_k\alpha_k)e^{-\frac{qn_k\tau_k}{T_k}}$$

When deploying assets to provide maximum security, some criteria is needed to weigh the value of protection given to each. If we use a simple maximin criteria, wherein we try to maximize the minimum value of P_D at any point, we end up with a situation analogous to the one shown in Figure B.3.

Figure B.3. Allocation Across Many Points



In Figure B.3, we see the additional value of adding an asset to a given point as a block, and we may choose to fill it or not. The maximin criteria is akin to filling up the confidence levels so the levels of liquid in the glass are as even as possible. This criteria leads to a kind of “peanut-butter spreading” where one asset is assigned to each unit, and additional ISR assets are assigned preferentially to the points with the worst (lowest) ratios of penetration time to revisit time ($\frac{\tau_k}{T_k}$).

As before, this is a general feature. If one point is particularly difficult to protect, it will require more assets. There are no surprises here. The purpose is to show that a simple mathematical model can capture many aspects of ISR allocation issues in a sensible way. This is another approach to modeling when the problem is simple enough that an agent-based simulation is unnecessary. However, as soon as we start looking at other aspects of the problem—such as tracking and capturing the intruder—a mathematical approach of this kind quickly finds its limitations.

Appendix C. Takeaways from Review of Selected Intelligence Community and Commercial World Assessment Practices

This appendix describes some key takeaways from a limited-scope review of metrics development and use practices within selected communities and industries, including the intelligence community, the healthcare provider Kaiser Permanente, “smart markets” and their role in Uber’s ride-sharing service, online (internet) search, and the academic community. Our intent with this examination was to develop a sense of some kinds of assessment and valuation practices used outside the USAF ISR enterprise. We did not widely survey practices across many organizations or summarize them in any great detail. We elected to look at one IC agency’s practices because of similarities with some missions USAF analysts support and because one project team member had years of professional experience to draw upon for the review. The other examples we examined were selected for diversity, based on project team member prior research experience, and because conversations with experts initially interviewed for the research suggested that there might be some useful takeaways from looking at these industries. Overall, these examples revealed useful insights about exploring information value, developing metrics, and prioritizing resources; determining data requirements and facilitating collection and information communication; and executing decisions through the use of data and metrics. Our examination did inform analysis for some of the assessment steps described in the main body of the report, including through cautionary tales; however, we did not sufficiently review practices to recommend (or not) that the USAF ISR enterprise explicitly adopt any given technology or approach used in these other industries and communities. Below, in no particular order, we provide a short summary of each set of practices reviewed and takeaways that influenced some of our findings on ISR assessment.

Intelligence Community

The USAF ISR enterprise is part of the IC, but for a variety of reasons (e.g., differences in missions, authorities, and resources) does not necessarily always employ the same intelligence practices, including for assessment. Though we did not have the research scope to survey intelligence assessment practices across the IC in its entirety, we explored one team member’s experiences within the IC and supporting special operations,¹⁹ which highlighted the use of

¹⁹ This team member participated in airborne ISR operations as a collection manager and targeting officer in the IC and as an embedded analyst for Special Operations forces. The team member’s ISR experience includes missions in the Middle East, the Horn of Africa, and Central Asia between 2004–2015, where he helped develop and implement several of the approaches to ISR described in this example.

assessment to support collection management and operational decisionmaking that provided some informative inputs to the research team's analyses.²⁰ We do not specify names of specific organizations to preserve anonymity and because it is not relevant for the key takeaways about assessment discussed below.

In the practices reviewed, ISR missions focus on groups of related targets, or "intelligence threads." The term *thread* connotes the interwoven nature of the targets, and within each "thread" are different targets of varying types and importance. The articulation and organization of missions around these "threads" enables ISR activities to be evaluated in terms of fulfilling specific collection objectives and also to be judged by whether and to what degree each collection advances understanding that supports execution of higher-order operational objectives. These assessments are used to support decisions about future collection activities and operations.

Some measures used for assessment are only relevant to specific collection objectives or at certain times, and can be tracked using binary responses (yes/no): "Did we capture all vehicle loadups and loadouts?" "Did we arrive on target and establish a baseline prior to the meeting?" "Did we maintain the correct look angle on all entrances, or were we masked?" These can be instructive if they focus on measuring activities that are within the control of IC mission teams rather than outcomes that are dependent upon the adversary's cooperation. These types of measures should be considered prior to making operational decisions, since they suggest which important factors are known and which ones are not.

Other measures for assessment are based on whether new collection "advances the narrative"—an analogy to building upon an existing storyline—meaning it should fill knowledge gaps and help refine courses of action. The emphasis here is on bringing the "story" to completion by contextualizing new developments within an existing body of knowledge. This contextualization may be different for different "threads." The advantage of making decisions based on assessments using "threads" is flexibility and adaptability when presented with new developments, presuming an ability to process and react to that information. Stressing that new collection must relate its findings to previous intelligence also appears to help reduce "wild goose chases" and redundant collection efforts, and increases the likelihood that new collection adds value and enables future decisionmaking opportunities.

There are still other types of assessment measures linking intelligence to decisionmaking, such as the incidence of noncombatant deaths when this results from poor intelligence activities. Having assessment information about meeting collection objectives and advancing intelligence narratives (discussed above) helps decisionmakers identify whether and how intelligence activities failed.

²⁰ The insights discussed in this section are not necessarily reflective of current best practices or policy in the IC, which may evolve rapidly. Additionally, this discussion is at the level of objectives, approaches, and institutional procedures, rather than that of sensitive sources and methods.

Using this approach to intelligence assessment enables the organizations examined to inform resource requirements based on the specific collection objectives for each intelligence thread and an assessment of the minimum number and asset capabilities required to meet those objectives. This is followed by an examination of the marginal utility of allocating more assets to the thread. Initial asset allocation among multiple intelligence threads is determined by the relative value expected from different allocations, with the recognition that optimizing initial collection for all threads is usually not possible. In most cases, the expectation is that assets will later be reallocated or “surged” to support new developments that constitute an opportunity to make progress on other collection and operational objectives.

Taking this approach appears to facilitate dynamic asset allocation, increasing efficiency and effectiveness. Mission teams engage in near-continuous reevaluation of collection objectives and each operation’s progress. Asset allocation decisions are not necessarily made in reference to the absolute value of a target, but are instead generally based on how much progress has already been made toward achieving the mission objectives relevant to each intelligence thread and what value can be expected from more collection. This approach to collection management tends to prioritize intelligence threads that are closer to the next or final operational decision point (e.g., to strike a target), and it deemphasizes or delays collection on threads that are still far from a decision point. For instance, if a high value target is being pursued, but it is not expected that much progress will be made toward the operational objectives during present ISR missions, then that high-value target might temporarily receive minimal or no ISR assets, with an understanding that resources can be surged to it later (assuming that the existing progress on other objectives will not be irrevocably lost by diverting assets).

The key takeaways from this example that helped inform our research on ISR assessments include the following:

- It is important to articulate and monitor not only how many ISR assets are being used, but how they are being used—i.e., are they being employed for the right activities?
- Intelligence activities often result in making incremental progress, which can be tracked by articulating and monitoring operational narratives and progress towards objectives.
- Assessments can be used to directly influence future intelligence collection and operational activities if measures are calculated using a consistent process and delivered to decisionmakers in a timely fashion.
- Assessments have the potential to help enhance both the efficiency and effectiveness of intelligence collection because they can enable more dynamism in collection management decisions.

Kaiser Permanente

Metrics are heavily used in the healthcare industry. Like the USAF ISR enterprise, healthcare providers sometimes struggle to identify what measures best reflect the quality of patient care and other factors these organizations care about. In 2006, Kaiser Permanente initiated the

development of a performance improvement system.²¹ This included the development and tracking of different types of metrics, including the *hospital standardized mortality ratio* (a clinical effectiveness measure), the *serious reportable event composite* (a safety measure), total delivery costs, and hospital rating (a service measure). Metrics were identified by first considering high-level organizational goals and then what activities are needed to accomplish those goals.

Kaiser Permanente uses a database to store the information collected in a standardized format, which enables access across the organization. The data are accessible via a web-based dashboard in which hundreds of performance measures are distilled into a few high-level metrics. The goal is to enable decisionmakers to track whether the organization is improving and to identify systemwide needs for improvement. The organizational leadership communicates performance expectations using a lexicon in line with language used to define metrics. Assessment results are translated into tactical improvements through the employment of assessment teams integrated with local providers.

However, Kaiser Permanente has experienced ongoing challenges with collecting data as part of daily operations. Individual personnel have differing habits for inputting data to the system. Some decisionmakers also find it difficult to link quantitative assessments to conceptual goals in order to determine what improvements (if any) need to be made.

Overall, this example highlighted the following which influenced the project team's thinking about USAF ISR assessments:

- Multiple metrics may be used in a “dashboard” when there are several goals that all need to be accomplished to achieve organizational success.
- Having a centralized database and common lexicon is necessary for communicating assessment results and implementing resulting decisions.
- Standards for inputting assessments data must be communicated and maintained.
- Linking overarching goals to intermediate objectives and tactical tasks is important for supporting decisionmaking using metrics.

“Smart Markets” and Uber’s Ride-Sharing Service

Intelligence information can, perhaps, be thought of as an economic good, but one which is difficult to intrinsically value due to rapidly changing and difficult to characterize supply and demand. Electric utilities and ride-sharing services (e.g., Uber), among others, share this challenge and have overcome it by leveraging a concept known as a “smart market.” Here, we first describe these markets and how they differ from traditional ones, and then give an example of the Uber ride-sharing service use of a “smart market” to match supply and demand in NRT.

²¹ Lisa Schilling, James Dearing, Paul Staley, Patti Harvey, Linda Fahey, and Francesca Kuruppu, “Kaiser Permanente’s Performance Improvement System, Part 4: Creating a Learning Organization,” *The Joint Commission Journal on Quality and Patient Safety*, Vol. 37, No. 12, December 2011.

Markets, in the economic sense, are a mechanism for individuals to reallocate goods at prices that are beneficial for both the seller and buyer. Prices in this framework function in two complementary ways: as signals between buyers and sellers, and as measures of value. Because of this, prices are used in a market as a decision metric by both buyers and sellers. These prices are a reflection of supply and demand, and their use as a metric allows participants to efficiently evaluate whether they wish to buy or sell their goods, or produce more.

“Smart markets” and related mechanisms are similar to traditional markets, except that the prices are generated by centralized computation which also allocates goods, in place of the conventional distributed market mechanisms for pricing goods and services. In smart markets, prices are generated based on market-participant bids on goods as allocations are decided.²²

Smart markets can be constructed to maximize producer profit, consumer savings, or overall economic surplus. They can also be used in places where coordination among producers is needed, but competition is desired as well. Smart markets can also be used to provide information that is inaccessible in traditional markets, like unfilled demand or the value of providing additional supply.

Because of this flexibility, smart-market-like mechanisms are used in a wide variety of areas. A common use of “smart markets” is for electric and water utilities, for which they function as an auction where suppliers and consumers each provide bids for different prices and quantities. From a computational perspective, the “smart market” is a linear program, which is efficient to solve to find prices and quantities that are both compatible with the mechanics of a particular system, and acceptable to all parties. The “smart market” also provides implicit prices for both producers and consumers to consider.

Unlike most traditional markets, smart markets allow a high degree of visibility into the behavior of the market for the market operators (or producers). In some cases, there is also greater visibility for market participants (or consumers). This visibility is enabled through the use of economic “shadow prices” that are generated by the systems or by simulating the system with variations to the input variables.²³ These can show the effective price or overall benefit impact of a change in any of the constraints or requirements.

For example, high shadow prices can inform participants or monitors that the system is operating at near capacity or is highly constrained by a single part of the system. This can lead to system operators or participants deciding to provide more of whatever factor is limiting supply.

Additionally, “smart markets” enable users to quantify their needs, which is especially useful in systems where resources are otherwise not priced or well-tracked. For example, “smart markets” can incentivize users to reflect on how much of a resource they truly need.²⁴ Because

²² This means that the techniques can be used even in contexts where traditional markets cannot.

²³ Frederick S. Hillier, *Introduction to Operations Research*, New York: McGraw-Hill Education, 2012.

²⁴ John F. Raffensperger and Craig Bond, “Creating a Smart Market for California Water,” *Orange County Register*, September 11, 2015.

many markets involve frequent changes in supply and demand, “smart markets” are frequently used to generate resource allocations over time and corresponding dynamic prices. Periodic auctions are rapidly rerun to produce new prices.

Example: Uber Ride-Sharing Service

Very frequently updated markets, like the one used for the Uber ride-sharing service, match individual riders and drivers in NRT, and adjusts prices to ensure overall supply matches demand. Uber was founded in 2009, and is widely known for its smartphone application that serves as matching platform between consumers seeking rides and drivers making their personal vehicles available for hire. Uber has proven very popular and has been growing very fast since launching in San Francisco in 2010. Here, we briefly touch upon Uber’s use of an electronic platform to perform multiple functions, including collecting data, as well as how data drives automated decisionmaking about ride pricing. This summary draws from a body of literature examining Uber.²⁵

Uber records and stores data about all actions taken on its application, including whether or not a user decides to order a ride. This electronic platform also handles communication, performance feedback, and payment processing mechanisms.

Uber’s base pricing system is somewhat similar to standard cab pricing systems that define the fare as a function of city and type of product, as well as price per mile, price per minute, a fixed fee and a minimum total fare. In addition to these standard parameters, Uber also utilizes a dynamic pricing system, called surge pricing. Uber’s surge algorithm monitors rider demand from ride requests and number of users opening the app, as well as the available driver supply. It then institutes a multiplier on the base price in a given geographical area when demand outstrips supply at the base price.

The “smart market,” and Uber’s specific use of this concept, yielded some insights that fed into our reflections on the USAF ISR enterprise’s potential for enhancing assessments:

- NRT time matching of supply and demand can provide the potential for close-to-instantaneous valuation of a service.
- User participation in markets (i.e., input and feedback) is critically important for pricing and apportionment of resources, and in calculating unmet demand.
- Centralized data collection facilitated by standardized procedures and applications to gather user input enables rapid, consistent calculation of value to support decisionmaking.

²⁵ See Le Chen, Alan Mislove, and Christo Wilson, “Peeking Beneath the Hood of Uber,” *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 495–508; M. Keith Chen, “Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform,” *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016; and Peter Cohen, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, *Using Big Data to Estimate Consumer Surplus: The Case of Uber*, National Bureau of Economic Research, No. w22627, 2016.

- Both actively supplied input (e.g., “bids” on services or ordering of a ride) and passively recorded activity (e.g., opening an application) can be valuable for determining the value of an asset at a particular point in time.

Online Search

Quickly finding the right information is important to the USAF ISR enterprise and those that depend on it. The commercial web contains vast amounts of information, which search engines such as Google, Safari, and Internet Explorer allow users to rapidly identify and select content from ranked lists of results. These engines employ algorithms to search webpages and rank results according to how well websites match a query. The more similar webpage titles, metatags, subtitles, hyperlinks, anchor text, and body content are to search queries, the more highly ranked they are in the list returned by the search algorithm. Ranking also improves based on how updated, popular, authoritative, hyperlinked and easy-to-read a site is.

Search algorithms can only be effective if web content is discoverable and indexed appropriately. Web content is “discovered” through the use of web crawlers or web spider algorithms that rapidly scour the Internet and record information, and through document feeds that track real-time information streams from news or social media sites.

After documents are found through the web crawler or received through a document feed, the search engine converts the information from various data formats (e.g., HTML, XML, PDF, Microsoft Word) into a single, compact, easy-to-access format. This format includes both text and metadata and is stored in a database for quick access.

With data indexed and stored, a search engine is able to rapidly search and return results to a user’s search query. After receiving a query, a search engine uses algorithms to look through its indexed data and provides webpages that best answer the search. PageRank and similar algorithms rank results according to popularity (how many visits a page receives) and authority (how many other sites link to it). Search terms that appear in the titles, subtitles, metatags, anchor text, and associated links are given greater weight and are ranked more highly. A search engine also searches the full text of a web page and returns pages with a high occurrence of the search query. It also ranks results according to number of links and “freshness” or how recently the page was updated. The interested reader can refer to several sources for additional information.²⁶

This example is useful for considering ISR assessment in the following ways:

- Data discoverability and indexing is essential for establishing a database that can be queried.
- With the right algorithms and databases in place, basic data management and retrieval can be automated to enable continuous and consistent data storage and retrieval processes.

²⁶ See W. Bruce Croft, Donald Metzler, and Trevor Strohman, *Search Engines, Information Retrieval in Practice*, New York: Pearson, 2015; and BBC, “How Do Search Engines Work?” webpage, undated.

- Metadata, or the ways in which information is tagged, is very important for influencing how it can be discovered and used.

Academic Citation Indexes

Academic researchers can experience similar “data deluges” to those working in the USAF ISR enterprise. There is a tremendous amount of academic knowledge communicated in peer-reviewed journal articles contained in a huge variety of publications. One class of metrics has been developed over the last few decades to help the academic community decide which research has a high level of importance, and should thus be given priority for attention. Citation impact indicators provide information on the impact of scientific publications based on the citations (or references) they have received. The indicators may focus on the impact of individual publications or the impact of researchers, research groups, research institutions, countries, or journals (i.e. different research units). The main purpose of the different indicators is to provide information for evaluating the quality and quantity of research. The body of academic literature around citation impact indicators has grown rapidly in recent decades, and can be found in bibliometrics, scientometrics, and research evaluation journals.²⁷

Example types of citation impact indicators include:²⁸

- total number of citations, or average number of citations received per publication. The most popular indicator based on citation averages is the journal impact factor, which counts the average number of citations received by recent articles published in a given journal. Other research units can also be used; examples include the average number of citations received by articles published by a given research institution, or the total number of citations that the articles in a journal receive in a given year divided by the number of articles published (the immediacy index).
- total number or proportion of highly cited publications. For these measures, a threshold needs to be established to determine whether a publication is counted as highly cited or not. The i10-index reported by Google Scholar is based on this idea, as it counts the number of publications by a given scholar with at least 10 citations.
- the H-index (or Hirsch index), which is often used to evaluate and compare individual researchers. For example, a researcher is given a H-index of 3 if three of her publications each have at least three citations and the other researchers in her department or other unit have no more than three citations each.

Sometimes citation impact indicators also consider the level of resources used to conduct the research being published—for instance, the amount of funding allocated to a given research unit.

²⁷ See John Panaretos, and Chrisovaladis Malesios, “Assessing Scientific Research Performance and Impact With Single Indices,” *Scientometrics*, Vol. 81, No. 3, 2009, p. 635; Leo Egghe, “The Hirsch Index and Related Impact Measures,” *Annual Review of Information Science and Technology*, Vol. 44, No. 1, 2010, pp. 65–114; Ludo Waltman, “A Review Of The Literature On Citation Impact Indicators,” *Journal of Informetrics*, Vol. 10, No. 2, 2016, pp. 365–391.

²⁸ Waltman, 2016.

This enables measurement not only of citation impact but also of publication productivity.²⁹ However, such indicators have received only limited attention in the literature. Other metrics besides the citation impact indicators discussed above include article views, downloads, or mentions in social media.

There are several potential factors to consider when using citation impact indicators. Analysts evaluating research impact need to consider which publications and citations to include in indicator calculations. Indicators can also be normalized to enable fairer comparisons between academic fields, document types, and time periods. How credit of coauthored work is allocated to individual authors can also impact indicator values. Choosing to explicitly include or exclude some types of citations (e.g., those made long after an article is published) can impact index values and the interpretation. There is no generally applicable rule for choosing citation windows, or the time period within which citations are eligible for includes in the calculation of indices.³⁰ Other potential problems include inconsistent inclusion of different types of documents (e.g., editorials or non-English language publications), the lack of uniformity and accuracy of data across indices and databases, differences in the amount of referencing between open access and paid subscription services.

Some important takeaways that we used in our examination of the USAF ISR assessment process included the following:

- Tracking source material can be helpful in comparing information value.
- There are alternative ways to value information using citations, each with strengths and weaknesses.
- Using common citation formats helps index and track citations.

²⁹ See Giovanni Abramo and Ciriaco Andrea D'Angelo, "How Do You Define and Measure Research Productivity?" *Scientometrics*, Vol. 101, No. 2, 2014, pp. 1129–1144.

³⁰ Jian Wang, "Citation Time Window Choice for Research Impact Evaluation," *Scientometrics*, Vol. 94, No. 3, 2013, pp. 851–872.

Bibliography

- Abramo, Giovanni, and Ciriaco Andrea D'Angelo, "How Do You Define and Measure Research Productivity?" *Scientometrics*, Vol. 101, No. 2, 2014, pp. 1129–1144.
- Air Combat Command, ISR Assessment Framework (Draft), September 20, 2016.
- Air Force Doctrine Annex 2-0, *Global Integrated Intelligence, Surveillance, and Reconnaissance Operations, Introduction to Global ISR*, Washington, D.C.: Headquarters, U.S. Air Force, January 29, 2015.
- Air Force Doctrine Document 2-0, *Global Integrated Intelligence, Surveillance, and Reconnaissance*, Washington, D.C.: Headquarters, U.S. Air Force, January 16, 2012.
- Air Force Instruction 13-1AOC, Vol. 3, *Operational Procedures—Air Operations Center (AOC)*, Washington, D.C., May 18, 2012.
- BBC, "How Do Search Engines Work?" webpage, undated. As of February 4, 2017:
<http://www.bbc.co.uk/guides/ztbjq6f>
- Box, George E. P., and Norman R. Draper, *Empirical Model-Building and Response Surfaces*, New York: John Wiley and Sons, 1987.
- Caplan, Ezra, *Recommendations for Mission-Centric Airborne ISR Employment and Assessment*, Maxwell Air Force Base, Ala.: LeMay Center for Doctrine, March 2017.
- Chen, Le, Alan Mislove, and Christo Wilson, "Peeking Beneath the Hood of Uber," *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 495–508.
- Chen, M. Keith, "Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform," *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016.
- Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe, *Using Big Data to Estimate Consumer Surplus: The Case of Uber*, National Bureau of Economic Research, No. w22627, 2016.
- Croft, W. Bruce, Donald Metzler, and Trevor Strohman, *Search Engines, Information Retrieval in Practice*, self-published edition, 2015.
- Cronk, Terri Moon, "Strike on ISIL Camp Protected National Security, Pentagon Press Secretary Says," *DoD*, February 19, 2016
- Egghe, Leo, "The Hirsch Index and Related Impact Measures," *Annual Review of Information Science and Technology*, Vol. 44, No. 1, 2010, pp. 65–114.

- Headquarters, Air Combat Command, “ISR Assessments Deliverable Review #2,” briefing slides, August 29, 2016.
- Hein, James L., *Discrete Mathematics*, 2nd ed., Sudbury, Mass.: Jones & Bartlett Publishers, 2003.
- Heinze, Clint, *Modelling Intention Recognition for Intelligent Agent Systems*, Canberra, Australia: Defence Science and Technology Organisation Systems Science Laboratory, No. DSTO-RR-0286., 2004.
- Hillier, Frederick S., *Introduction to Operations Research*, New York: McGraw-Hill Education, 2012.
- Hodges, Ashley, “UNICORN 101,” 480th Intelligence Wing, briefing slides, April 29, 2008.
- National Geospatial-Intelligence Agency, “GEOINT Standards,” webpage, undated. As of August 31, 2018:
<https://www.nga.mil/ProductsServices/geointstandards/Pages/default.aspx>
- Joint Functional Component Command for Intelligence, Surveillance and Reconnaissance Assessments Division, “Measures of Effectiveness,” briefing slides, April 20, 2016.
- Joint Publications 1-02, *Department of Defense Dictionary of Military and Associated Terms*. Washington, D.C.: Joint Chiefs of Staff, November 8, 2010.
- Joint Publication 2-0, *Joint Intelligence*, Washington, D.C., October 22, 2013.
- Joint Publication 3-0, *Joint Operations*, Washington, D.C., January 17, 2017.
- Joint Publication 3-30, *Command and Control of Joint Air Operations*, Washington, D.C., February 10, 2014.
- Keating, Joshua E., “What Do You Learn at Terrorist Training Camp?” *Foreign Policy*, May 10, 2010.
- Koopman, Bernard O., “Search and Screening,” Operations Evaluation Group Report No. 56, U.S. Navy, 1946.
- Leachtenauer, Jon C., William Malila, John Irvine, Linda Colburn, and Nanette Salvaggio, “General Image-Quality Equation: GIQE,” *Applied Optics*, Vol. 36, No. 32, 1997, pp. 8322–8328.
- Lingel, Sherrill, Lance Menthe, Brien Alkire, Scott Grossman, Robert A. Guffey, Keith Henry, Lindsay D. Millard, Christopher A. Mouton, George Nacouzi, Edward Wu, and John Gibson, *Methodologies for Analyzing Remotely Piloted Aircraft Effectiveness in Future Roles and Missions*, Santa Monica, Calif.: RAND Corporation, DB-637-AF, 2012. As of November 5, 2018:
https://www.rand.org/pubs/documented_briefings/DB637.html

- Lingel, Sherrill, Carl Rhodes, Amado Cordova, Jeff Hagen, Joel Kvitky, and Lance Menthe, *Methodology for Improving the Planning, Execution, and Assessment of Intelligence, Surveillance, and Reconnaissance Operations*, Santa Monica, Calif: RAND Corporation, TR-459, 2008. As of November 5, 2018:
https://www.rand.org/pubs/technical_reports/TR459.html
- Menthe, Lance, Amado Cordova, Carl Rhodes, Rachel Costello and Jeffrey Sullivan, *The Future of Air Force Motion Imagery Exploitation: Lessons from the Commercial World*, Santa Monica, Calif.: RAND Corporation, TR-1133, 2012. As of November 5, 2018:
https://www.rand.org/pubs/technical_reports/TR1133.html
- Menthe, Lance, Myron Hura, and Carl Rhodes, *The Effectiveness of Remotely Piloted Aircraft in a Permissive Hunter-Killer Scenario*, Santa Monica, Calif.: RAND Corporation, RR-276-AF, 2014. As of November 5, 2018:
https://www.rand.org/pubs/research_reports/RR276.html
- Menthe, Lance, and Jeffrey Sullivan, *A RAND Analysis Tool for Intelligence, Surveillance and Reconnaissance: The Collections Operations Model*, Santa Monica, Calif.: RAND Corporation, TR-557-AF, 2008. As of November 5, 2018:
https://www.rand.org/pubs/technical_reports/TR557.html
- Munsell, Mark, “A Focus on Data Management; Transforming the Way We Do Business,” *CIO Review*, May 30, 2017. As of January 23, 2018:
<https://aerospace-defense.cioreview.com/cioviewpoint/a-focus-on-data-management-transforming-the-way-we-do-business-nid-24539-cid-5.html>
- Obama, Barack, *Executive Order 13732 on the U.S. Policy on Pre- and Post-Strike Measures to Address Civilian Casualties in U.S. Operations Involving the Use of Force*, Washington, D.C., July 1, 2016.
- Office of the Director of National Intelligence, *Summary of Information Regarding U.S. Counterterrorism Strikes Outside Areas of Active Hostilities*, Washington, D.C., July 1, 2016.
- Office of the Director of National Intelligence, Air Combat Command, *ISR Assessment Framework*, white paper, Washington, D.C.: September 27, 2016.
- Office of the President of the United States, *Procedures for Approving Direct Action against Terrorist Targets Located Outside the United States and Areas of Active Hostilities*, May 22, 2013.
- Panaretos, John, and Chrisovaladis Malesios, “Assessing Scientific Research Performance and Impact With Single Indices,” *Scientometrics*, Vol. 81, No. 3, 2009, p. 635.
- Raffensperger, John F. and Craig Bond, “Creating a Smart Market for California Water,” *Orange County Register*, September 11, 2015.

- Rasch, Georg, “The Poisson Process as a Model for a Diversity of Behavior Phenomena,” *International Congress of Psychology*, Vol. 2. 1963.
- Rhodes, Carl, Jeff Hagen, and Mark Westergren, *A Strategies-to-Tasks Framework for Planning and Executing Intelligence, Surveillance, and Reconnaissance (ISR) Operations*, Santa Monica, Calif.: RAND Corporation, TR-434-AF, 2007. As of December 30, 2017: https://www.rand.org/pubs/technical_reports/TR434.html
- Schilling, Lisa, James Dearing, Paul Staley, Patti Harvey, Linda Fahey, and Francesca Kuruppu, “Kaiser Permanente’s Performance Improvement System, Part 4: Creating a Learning Organization,” *The Joint Commission Journal on Quality and Patient Safety*, Vol. 37, No. 12, 2011, p. 532-AP5.
- Thaler, David E., *Strategies to Tasks: A Framework for Linking Means and Ends*, Santa Monica, Calif.: RAND Corporation, MR-300-AF, 1993. As of January 11, 2018: https://www.rand.org/pubs/monograph_reports/MR300.html
- USCENTCOM—See U.S. Central Command.
- U.S. Central Command, *ISR Roles*, white paper, undated.
- Waltman, Ludo, “A Review Of The Literature On Citation Impact Indicators,” *Journal of Informetrics*, Vol. 10, No. 2, 2016, pp. 365–391.
- Wang, Jian, “Citation Time Window Choice for Research Impact Evaluation,” *Scientometrics*, Vol. 94, No. 3, 2013, pp. 851–872.
- White House Office of the Press Secretary, “Fact Sheet: Executive Order on the U.S. Policy on Pre- and Post-Strike Measures to Address Civilian Casualties in U.S. Operations Involving the Use of Force and the DNI Release of Aggregate Data on Strike Outside Areas of Active Hostilities,” Washington, D.C., July 1, 2016.



The changes in intelligence, surveillance, and reconnaissance (ISR) and processing, exploitation, and dissemination (PED) capabilities over the past two decades have led to ever-increasing demand from warfighters. Commanders, planners, and operators across the U.S. Air Force (USAF) ISR enterprise face difficult decisions about how to best meet ISR needs at the strategic, operational, and tactical levels. Yet USAF currently lacks a consistent, quantitative, empirically grounded method of assessing the value that the service's airborne ISR provides—which is essential to good resourcing decisions. This report presents an approach to ISR assessments that seeks to articulate the costs and benefits of USAF airborne ISR in specific operational contexts. Though aspects of this may be applicable across different USAF ISR organizations, this work focused primarily on the Distributed Common Ground System and the operational theaters it does or could support. The assessment methodology is designed to be flexible enough to support ISR resourcing decisions at different echelons, yet consistent enough to foster feedback, standardize data collections, and make use of empirical analysis methodologies.

\$43.00

www.rand.org

ISBN-10 1-9774-0693-9
ISBN-13 978-1-9774-0693-4

