



APPLIED RESEARCH LAB
**FOR INTELLIGENCE
AND SECURITY**

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

**Crowdsourced Security and Intelligence
Forecasting Tool (CSIFT) Mitigating Insider Risk
(MInR)
Seedling
Final Report**

September 30, 2021

Monique Beaudoin, PhD^{1*}, Thalia Baumgarten², Sean Kucer², Adam Siegel²

¹Applied Research Laboratory for Intelligence and Security

²Cultivate Labs

*corresponding author, mbeaudoin@arlis.umd.edu

EXECUTIVE SUMMARY

Insider threats are a long-standing concern within the Department of Defense (DoD), where individuals can use inside knowledge of their workplace, employer, or coworkers to threaten the security of U.S. national security interests, whether intentionally or unintentionally. As prior case studies have shown, insider threat events are difficult to predict, due to the complex interaction of multiple factors, combined with limitless unknown variables in an ever-changing world context.

To assist the DoD and USG in its mission to anticipate, mitigate, and prevent such threats, ARLIS performed research and feasibility analyses to determine how to apply the Crowdsourced Security & Intelligence Forecasting Tool (CSIFT) and its quantitative methodology specifically to this critical IC/DoD challenge area. Combining literature research with inputs from government stakeholders and subject matter experts, The ARLIS and Cultivate Labs team identified three unique insider threat/risk use cases and decomposed key issues into specific forecasting questions for use in a future launch of the CSIFT tool for the IC/DoD, to anticipate and mitigate insider risk. A draft of the dashboard design was also created, and suggested forecasting populations were identified.

Contents

EXECUTIVE SUMMARY	2
INTRODUCTION	4
The ARLIS and Cultivate Labs team	4
<i>CSIFT: Technical Approach</i>	5
Use Cases and Key Issues	7
Use Case 1: Domestic Extremism	7
Key Issue 1: National Appeal of Domestic Extremist Organizations	9
Key Issue 2: Organizational Culture	10
Key Issue 3: Security Architecture	11
Key Issue 4: Domestic Extremism Presence	12
Use Case 2: Workplace Violence	13
Who are the stakeholders, and why? Who are the forecasters?	15
Key Issue 1: Mental Health in the Workplace	15
Key Issue 2: External Stressors	16
Key Issue 3: Concerning Behaviors	16
Key Issue 4: Response by Organization	17
Use Case 3: Academic Insider Risk (Espionage)	18
Academic Insider Risk (Espionage) Use Case: Key Issue Decompositions	19
Who are the stakeholders, and why? Who are the forecasters?	19
Key Issue 1: National Threat, Vulnerabilities, and Countermeasures	20
Key Issue 2: University Vulnerabilities	20
Academic Insider Risk (Espionage) Comparative Question Clusters	21
Who are the stakeholders, and why? Who are the forecasters?	21
Design of the Decision-maker Dashboard	22

End User 2: Analyst.....	25
Recommendations.....	27
ACKNOWLEDGEMENTS.....	28
DISCLAIMERS.....	28
ABOUT ARLIS	28
Technical Points of Contact:	28
Administrative Points of Contact:	28
APPENDICES.....	29
A.1: Quantitative methodologies used in aggregative forecasting	29
B1: Aggregation strategies.....	31
Default/Starting Aggregation Strategy	31
Decay Strategies	31
Weighting	31

INTRODUCTION

Insider threat events – especially emergent threats without prior antecedents – are difficult to predict, even using best current practices and methodologies for threat detection and analysis, such as predictive modeling, expert judgment, and historical reflection. It can also be difficult to assess risks and vulnerabilities that make such threats more or less likely. If available, such risk assessments would facilitate proactive actions; enhance the availability of preventative estimations and mitigation approaches; and potentially forestall both familiar and emergent threats.

Most current responses to insider threats focus on trying to identify specific current or near-term threats, helping organizations respond to those threats, and reducing the potential damage that can be caused from an individual event. For example, in the cybersecurity realm, monitoring for anomalies enables system administrators to detect potential insider events as they are happening, while cyber security policies seek to prevent access (e.g. through password complexity requirements) or to limit damages (e.g. having detailed access privileges so that an insider has limited data access). While existing approaches are critical, they do not have the inherent flexibility to predict the wide range of insider threats, such as workplace violence.

Given the need for real predictions to anticipate real threat events across this broad spectrum, ARLIS explored the application of a quantitative methodology (see Appendices A and B) known as crowdsourced forecasting to provide unique insights to the government related to the mitigation of insider risk. This approach had not previously been applied to the insider threat/risk prediction and prevention challenge but was developed and honed throughout more than ten years of basic research investment by the IC¹. The approach leverages the wisdom of well-informed but non-expert pools of individuals to predict the likelihood of events and can be applied to any critical topic of importance to the government.

The *Crowdsourced Security & Intelligence Forecasting Tool (CSIFT) for Mitigation of Insider Risk (CSIFT-MInR)* seedling thus focused on researching how crowdsourced forecasting could assist the DoD and the U.S. Government (USG) in its mission to:

- Predict trends and event outcomes that may increase the risk of insider threat events
- Assess the future likelihood of previously known and documented insider threats and vulnerabilities
- Increase inter- and intra-agency collaboration to improve the diversity of input; highlight and quantify disagreement; and generate consensus viewpoints

The ARLIS and Cultivate Labs team

The Applied Research Laboratory for Intelligence and Security (ARLIS) is a University-Affiliated Research Center (UARC) housed at the University of Maryland, supporting the Office of the Undersecretary of Defense for Intelligence and Security (OUSD(I&S)). A trusted partner of the IC/DoD, ARLIS also possesses the unique and necessary combination of core competencies to innovate solutions related to insider risk/threat, namely: social and behavioral sciences, computer

¹ Walter Frick, “What Research Tells Us About Making Accurate Predictions,” Harvard Business Review (February 2, 2015), <https://hbr.org/2015/02/what-research-tells-us-about-making-accurate-predictions>, Phillip Tetlock, *Super Forecasting: The Art and Science of Prediction* (New York: Broadway Books, 2015), Chapter 3; Puong Fei Yeh, “Using Prediction Markets to Enhance U.S. Intelligence Capabilities,” *Studies in Intelligence*, 50 no. 4 (2006), <https://www.cia.gov/library/center-for-the-study-of-intelligence/csi-publications/csi-studies/studies/vol50no4/using-prediction-markets-to-enhance-us-intelligence-capabilities.html>.

sciences (including AI/ML), and advanced computing. Further, ARLIS is currently providing initial support for the crowdsourced forecasting platform and its software and maintaining a Community of Interest for forecasting across the IC. ARLIS led this seedling with the support of industry subcontractor Cultivate Labs, who were integral to the success of prior IC research efforts establishing the scientific basis for crowdsourced forecasting. Cultivate Labs and its personnel continue to be recognized as crowdsourced forecasting experts in the IC/DoD, where they support other forecasting efforts across a range of topics.

CSIFT: Technical Approach

Crowdsourced forecasting is the ongoing collection of predictions (typically in the form of estimates of event probability) from a large, diverse group of people, which are aggregated into a dynamic “crowd” or “consensus” forecast about an outcome. This consensus forecast shifts over time as new estimates are submitted. Every forecast includes both “true signal” and “noise” – or error. Because the errors are stochastic, they effectively cancel each other out when the probabilities are aggregated, leaving a forecast that is typically more accurate than any individual forecast. This aggregated forecasting stream can be used by policy makers and practitioners as:

- An “early warning system” to proactively signal a risk is becoming more likely
- A challenge to *status quo* assumptions
- To highlight when and how perspectives shift in response to new developments
- A method to quantify disagreement among individuals or cohorts of participant

Crowdsourced forecasting has its origins in the work of a British Statistician named Sir Francis Galton, who discovered the concepts of mean, median, and other basic statistical concepts. In the journal *Nature* in 1906, Galton published findings showing how he used the median to guess almost the exact weight of an ox by soliciting guesses from a large group of people attending a county fair. The crowd prediction was much more accurate than any individual’s (including the veterinarian and the farmer!). In the 1980’s, the Iowa Electronic Market was started at the University of Iowa to crowdsource predictions about Presidential Elections, and in 2011 the U.S. Intelligence Advanced Research Projects Activity (IARPA) funded a multi-million dollar, 4-year study to test the effectiveness of crowdsourced forecasting and identify the most optimal conditions for accuracy. The “wisdom of the crowd” has been replicated many times in situations where opinions diverge widely but are not consistently biased in a specific direction, as well as when many kinds of expertise can provide relevant but incomplete insight into outcomes. Since then, the USG has also proven the accuracy of this approach through multiple research and operational programs at IARPA and the Intelligence Community. In practice, crowdsourced forecasting supplements and enhances alternate estimation processes where modeling may not be effective, or where broad, diverse sets of perspectives would be beneficial.

To begin applying CSIFT methodologies to the area of insider risk mitigation and threat prevention for the IC and DoD, it was first necessary to identify the broad topics, or “key issues,” that would yield the most critical and useful information related to insider risk and threat. The team performed literature research regarding insider threat, met with stakeholders (Table 1), and defined and planned the most critical use cases and key issues.

Initial use cases identified included:

- Workplace violence
- Academic Insider Risk

- Domestic Extremism
- Intellectual Property Theft
- Critical Infrastructure
- Fraud
- Disclosure of classified, sensitive, or personal information

After continued research and further stakeholder and SME engagements, this list of potential was refined to include only three use cases which were determined to be most relevant and critical to the community of practice in insider threat and risk mitigation: Domestic extremism, workplace violence, and academic insider risk (with a focus on espionage).

Table 1: Stakeholder and SME engagements

Use Cases Discussed	Organization
All	DAF Counter-Insider Threat Hub / Policy & Oversight (SAF/AAZ)
All	OPA/PERSEREC
All	OUSD(I&S)
Domestic Extremism & Workplace Violence	UMD START
Academic Insider Risk	Texas A&M University, Research Security Office

Unlike traditional scenario planning, the CSIFT approach involves the identification of the *pivotal factors* that influence each use case, and the *individual signals* that enable the crowd to specifically forecast the outcome of those factors. This process for defining specific forecast questions for the CSIFT platform is called “issue decomposition.” By researching and decomposing each use case in this manner, forecasting questions were developed that would have falsifiable/verifiable and timebound outcomes. The decomposition process (*further described in Table 2*) yields many diverse forecasting questions that provide predictions about both drivers and indicators of insider risk.

Table 2: Decomposition Framework

Key Issues →	Pivotal Factors →	Individual Signals
What overarching themes or trends will impact the mission significantly? <i>e.g., Organizational Culture: Social and Behavioral Health</i>	Key issues can be broken down into various pivotal factors, which help to assess the directional impact of the issue. <i>e.g., Prevalence of negative behaviors: substance use, alcoholism, DUIs, gambling</i>	Signals are concrete ways to measure and track pivotal factors over time. They may be metrics, events, or other forecastable outcomes. These become the forecast questions. <i>e.g., How many people will be referred to preventive/healthy lifestyle counselling for</i>

		<i>"Alcohol Use" as part of their Periodic Health Assessment between 1 January 2022 and 1 July 2022?</i>	For example, a forecasting question surrounding domestic extremism could
--	--	--	--

be: "How many current or former employees of the US Department of Defense will be arrested before 1 January 2022 due to suspected involvement in the Boogaloo movement?" This question is timebound and falsifiable; individual forecasters may enter (and update) their probabilities (informed guesses) at any time during the open forecasting period for that specific question (in this case, ending on 1/1/2022). At that point, the outcome of the forecasting question is knowable, and individual forecasts may be "judged" across the platform. By asking falsifiable questions such as these, the approach can measure the performance of estimations made at levels of individuals, cohorts of forecaster pools, and the full crowd. Even during an open forecasting period, estimated probabilities at any time point can also be informative for practitioners and decision makers, particularly when crowd forecasts diverge from initial expectations or change abruptly.

Use Cases and Key Issues

Described below is the list of the use cases and associated key issues relating to insider threats that were developed in decompositions. These use cases and issues were specifically chosen as a focus of our methodology because they are currently being addressed by the USG and are challenging to predict with traditional methods.

Use Case 1: Domestic Extremism

The Office of the Director of National Intelligence (ODNI) defines domestic extremists as "US-based actors who conduct or threaten activities that are dangerous to human life in violation of the criminal laws of the United States or any state; appearing to be intended to intimidate or coerce a civilian population; and influence the policy of a government by intimidation or coercion, or affect the conduct of a government by mass destruction, assassination, or kidnapping."² The number of extremist attacks has been rising over the past two decades, with a doubling of attacks from 2002 to 2017 and a sharp increase in domestic terrorist investigations.³ The January 2021 attack on the US Capitol by individuals seeking to prevent certification of the 2020 election results – including supporters of Donald Trump and members of domestic extremist groups such as the Proud Boys and Oath Keepers – catapulted this issue to the forefront of the nation's consciousness.⁴ In March 2021, ODNI released a report underscoring the heightened threat posed by domestic extremists throughout the US.⁵

² from [ODNI](#), as per [18 U.S. Code 2331 \(5\)](#)

³ [Understanding Domestic Radicalization and Terrorism | National Institute of Justice, Opinion | Does the FBI have the right culture to fight domestic terrorism?](#)

⁴ [Researchers: More Than a Dozen Extremist Groups Took Part in Capitol Riots](#)

⁵ [\(U\) Domestic Violent Extremism Poses Heightened Threat in 2021 01 March 2021](#),

Of particular concern is the affiliations some of the members of these groups have with the US military. Nearly 1 in 5 of the individuals arrested in connection with the Capitol breach had ties to the military.⁶ A March 2021 report from the Department of Defense (DOD) to the House and Senate Armed Services Committees highlighted how domestic extremist groups were targeting the military by seeking to recruit military members, encouraging their members to join the military.⁷ The following month CSIS released a report analyzing domestic terrorism data from 1994 to 2021. They found that the percentage of attacks and plots perpetrated by active-duty military and reservists has increased by 6.4% in 2020 (compared to 2019).⁸ The numbers of attacks are low but growing. In 2020, active duty military and reservists perpetrated seven domestic terror attacks.⁹ The trends identified in the report are troubling given that “individuals with a military or law enforcement background have skills the extremists want -- such as proficiency in firing weapons, building explosive devices, conducting surveillance and reconnaissance, training personnel, [and] practicing operational security.”¹⁰ At the same time, the military is increasingly becoming a target for attacks by domestic extremists: attacks against the military, government, and police constituted 38% of attacks in 2022, the most of any category.¹¹

These two trends point to the importance of addressing insider threats emanating from domestic extremist groups. Identifying these threats is difficult because individuals often “radicalize independently by consuming violent extremist material online and mobilize without direction from a violent extremist organization.”¹² In addition, information about this risk is distributed widely, with many different organizations studying the issue and warning signs appearing in many different, unconnected places. For example, DHS, CISA, and NCSC all work on relevant topics, yet information is siloed among certain individuals at these organizations.

The insider threat from domestic extremism use case was decomposed into four **key issues**: *National Appeal of Domestic Extremism, Organizational Culture, Security Architecture, and Domestic Extremism Presence*. In addition to researching each key issue, a potential stakeholder was also identified for each issue area, initial forecasting questions were developed, and relevant forecasting pools for future CSIFT platform launch were identified.

⁶ [Military Veterans Overrepresented In Those Charged In Jan. 6 Capitol Riot](#), [In the Battle for the Capitol, Veterans Fought on Opposite Sides](#), [A Veteran's Journey From Army to Storming the Capitol | Time](#)

⁷ Office of the Under Secretary of Defense for Personnel and Readiness, Report to Armed Services Committees on Screening Individuals Who Seek to Enlist in the Armed Forces (Washington, DC: U.S. Department of Defense (DoD), March 2021), <https://media.defense.gov/2021/Mar/02/2002592042/-1/-1/0/REPORT-TOARMED-SERVICES-COMMITTEES-ON-SCREENING-INDIVIDUALSWHO-SEEK-TO-ENLIST-IN-THE-ARMED-FORCES.PDF>.


⁸ [The Military, Police, and the Rise of Terrorism in the United States](#)

⁹ Ibid.

¹⁰ Ibid.


¹¹ Ibid.

¹² [\(U\) Domestic Violent Extremism Poses Heightened Threat in 2021 01 March 2021](#), pg. 2

Figure 1: Sample Use Case Decomposition: Domestic Extremism**Key Issue 1: National Appeal of Domestic Extremist Organizations**

This key issue captures the strength of domestic extremist groups nationally. When their appeal is stronger, recruitment will be easier, and organizations must be more alert to potential insider threats. When their appeal goes down, recruitment will be more difficult. This issue is about the environment -- how conducive the environment is to domestic extremism, internally. The environment is more conducive to domestic extremism if the nationwide appeal of domestic extremist groups is high.

Figure 2: Sample Key Issue Decomposition: Domestic Extremism National Appeal



Who are the stakeholders, and why? Who are the forecasters?

This key issue was developed with the Weapons Training Battalion, Camp Lejeune, NC as the stakeholder. It is located in a state with high extremist activity, according to the Southern Poverty Law Center.¹³ Additionally, Camp Lejeune is home to US Marine Corp units, and according to a recent analysis completed by the National Consortium for the Study of Terrorism and Responses to Terrorism (START), individuals affiliated with the Marine Corp were the second highest group (preceded by the Army) of United States Extremists with Military Backgrounds between 1990-2021.¹⁴ Because this key issue is focused on national trends and data, the optimal forecaster population would be a public pool, such as from Good Judgement Open.

Key Issue 2: Organizational Culture

Many of the case studies examined during this study suggested that mental and social health played a key role in most individuals' paths to radicalization. Key indicators of an individual's susceptibility to DE recruitment are loneliness, unaddressed mental health issues, lack of an outlet for workplace and personal grievances, and a history of domestic violence. This key issue assesses whether an organization's culture exacerbates or alleviates those issues.


Because we focused on forecasting the risk of extremism within an organization – not on evaluating an individual's risk for extremism – this key issue focuses on organizational culture. Is the culture inclusive and does it adequately address the social and mental health needs of its members? Like

¹³ <https://www.splcenter.org/hate-map>

¹⁴ https://www.start.umd.edu/sites/default/files/publications/local_attachments/Extremism%20In%20the%20Ranks%20and%20After%20-%20Research%20Brief%20-%20July_13_2021%20Final%20accessible.pdf

the “National Appeal of Domestic Extremism” key issue, this is an environmental issue; however, it is organization-specific rather than national. The environment is very conducive to domestic extremism internally if the organization’s culture is rife with social or mental health challenges and will be less conducive if the organization’s culture is inclusive and mental health issues are addressed effectively.

Figure 3: Sample Key Issue Decomposition: Organizational Culture – Social and Behavioral Health



Who are the stakeholders, and why? Who are the forecasters?

The decomposition of this key issue was developed using any unit in the US Army Installation at Fort McCoy, Wisconsin, in mind as a potential stakeholder. This group was selected because of the geographic and functional diversity of this installation relative to others. Army units also have higher rates of domestic extremism.¹⁵ Because this key issue is at the organizational level, the forecaster population would be a cross-section of everyone in the unit, as well as unit support staff such as mental health and social services (i.e. legal personnel).


Key Issue 3: Security Architecture

While the previous two key issues of Organizational Culture and the National Appeal of Domestic Extremism capture the willingness of individuals to engage in extremist behavior, they do not assess the ability of extremists to acquire information, access materials, or access sensitive areas to execute an attack. The way in which organizations protect their information and their physical assets significantly impacts the ability of extremists to appropriate information, steal material, or

¹⁵ <https://www.csis.org/analysis/military-police-and-rise-terrorism-united-states>

plan attacks. This key issue looks at how strong an organization's security architecture is. It looks at security policies, team compliance with those policies, and cultural components of security. Organizations that have high compliance and a culture that prioritizes physical and information security are going to be less at risk of a domestic extremist insider threat.

Figure 4: Sample Key Issue Decomposition: Security Architecture




Who are the stakeholders, and why? Who are the forecasters?

The decomposition of this key issue was focused on Naval Station Mayport, largely due to the geographic and political diversity within the region, as well as due to the large number of bases in Florida. The aim was to demonstrate the flexibility of this approach and how it can be used by a diverse set of stakeholders. An ideal forecasting pool for this key issue would be a cross-section of the unit, as well as security forces and the IT support squadron.

Key Issue 4: Domestic Extremism Presence

The first three sets of key issues focus on the environment – whether the national and organizational environments are vulnerable to domestic extremism. This key issue is focused on indicators that suggest that a domestic extremist group has penetrated the organization. For example, if there are members of the organization who have joined an extremist group and are actively engaged in activities that undermine the organization and/or give the extremist group access to critical information, technology, etc. This key issue is about tracking behaviors. It assesses what is happening within the organization and whether there are any indicators that there is an active insider threat.

Figure 5: Sample Key Issue Decomposition: Domestic Extremism Presence



Who are the stakeholders, and why? Who are the forecasters?

We chose the 377th AFB Wing at Kirkland AFB and Sandia National Laboratories. The ideal forecasting pool would include a cross-section of the unit and people from the internal threat office. This would ensure that a wide range of relevant insights and knowledge are incorporated into the consensus forecast.

Use Case 2: Workplace Violence

Workplace violence is defined as “a verbal, written, or physically aggressive threat or attack intended to intimidate, cause injury or death to others in a place of employment”¹⁶, and is a fundamental threat to the productivity of the USG. Roughly two million American workers are victims of workplace violence every year.¹⁷ That number does not include the many incidents of workplace violence that are thought to go unreported.¹⁸ Workplaces with any degree of aggression suffer lost productivity, diminished ability to fulfill the mission, and an unpleasant and potentially dangerous work environment. This seedling is focused on workplace violence perpetrated by insiders of the USG.

Workplace violence is the third leading cause of occupational fatalities¹⁹, and the military is not exempt from it or its negative consequences. In fact, in 2015 the Inspector General found that the

¹⁶ <https://www.cisa.gov/workplace-violence>

¹⁷ <https://www.osha.gov/sites/default/files/publications/factsheet-workplace-violence.pdf>

¹⁸ <https://www.osha.gov/sites/default/files/publications/osh3148.pdf>

¹⁹ <https://www.osha.gov/workplace-violence>

DoD did not have an effective comprehensive program to combat workplace violence, and recommended they do so.²⁰

Crowdsourced forecasting can be applied to help better understand—and ultimately reduce—the risk of workplace violence. Information about this risk is distributed widely, with many different organizations studying the issue. The CSIFT platform is a way to bring all the information in an organization together in one location, where it is continuously updated as events unfold, providing decision makers and end users with a way to quickly ascertain the status of various risks.


CSIFT could help the USG address workplace violence by providing an accurate readout of the risks of workplace violence in specific locations, such as military installations. The focus of this research is on key issues involving Mental Health in the Workplace, External Stressors, Concerning Behaviors, and Response by Organizations. This includes forecasting on topics such as depression in the workplace, financial stressors, formal repercussions for infractions in the workplace, and the maintenance of efficient safety procedures. By using an informed, diverse crowd to forecast on these questions, the CSIFT platform would allow for more informed decision making, to improve planning, identify threats, and mitigate the risks of future violence.

CSIFT has the power to have both quantitative and qualitative impacts on the issue of workplace violence. It will reduce the number of violent incidents in the workplace and will give decision makers accurate and continuously updating probabilities of various risks related to workplace violence. Qualitatively, CSIFT will improve the overall work environment and productivity as decision makers better mitigate various risks. It will also improve the quality of decision-making processes.

the problem of workplace violence was decomposed into four key issue areas: Mental Health in the Workplace, External Stressors, Concerning Behaviors, and Response by Organization.

Figure 6: Sample Use Case Decomposition: Workplace Violence

²⁰ <https://www.hsdl.org/?view&did=822368>



Who are the stakeholders, and why? Who are the forecasters?

The outputs of CSIFT could be presented to any number of stakeholders interested in mitigating the risks of workplace violence. The Carl R. Darnall Army Medical Center at Fort Hood (TX) was selected as the stakeholder for all key issues within this use case, in part because of the Army's higher rates of domestic extremism²¹. However, many of the forecast questions posed could apply equally well to other military installations and other organizations.


The forecasting pool for this use case and stakeholder would be composed of staff at the medical facility (from all positions, e.g. frontline, admin, supervisory), people served by the medical facility, and people charged with oversight of the medical facility. Anyone who has personal knowledge or experience with the facilities could be a candidate for forecasting.

Key Issue 1: Mental Health in the Workplace

This key issue is designed to capture ongoing personal characteristics involving mental health, such as patterns in the overall mental health of both patients and employees, rather than mental health challenges caused by particular circumstances. When there are more people with consistently poor mental health at a given organization, the rate of workplace violence is likely to increase.

Figure 7: Sample Decomposition: Mental Health in the Workplace


²¹ <https://www.csis.org/analysis/military-police-and-rise-terrorism-united-states>



Key Issue 2: External Stressors


This key issue is designed to capture the extent of external stressors that could be affecting workers' mental health. These are broken down into professional, personal, and financial stressors. Stressors inside and outside the workplace increase the likelihood of an incident of workplace violence occurring. Examples in this key issue include family stressors and officers not respecting soldiers.

Figure 8: Sample Key Issue Decomposition: External Stressors




Key Issue 3: Concerning Behaviors

This key issue is designed to capture employee behaviors that indicate the prevalence of workplace violence or that workplace violence is more likely to occur. Certain behaviors--such as compulsive gambling or frequent, heavy drinking--are cause for concern. Also included in this key issue are direct indicators of workplace violence such as employees experiencing violence at work.

Figure 9: Sample Key Issue Decomposition: Concerning Behaviors

Key Issue 4: Response by Organization

This key issue is designed to capture how well the organization manages risks and episodes of workplace violence. How comfortable people feel reporting things to management and how well management is set up to mitigate risks is a significant factor in the overall amount of workplace violence. Examples in this key issue include: completion of GSA's workplace standard response plan checklist, passing various workplace security checks, as well as employee ratings of organizational response.

Figure 10: Sample Key Issue Decomposition: Response by Organization

Use Case 3: Academic Insider Risk (Espionage)

In January for 2020, the chairman of Harvard's Chemistry Department was arrested by the FBI for lying on federal grant applications about the fact that he was receiving millions of dollars of financial contributions from China.²² This high profile case was just one of a string of investigations, arrests, and convictions that vaulted the threat posed by academic espionage to the forefront of the public's consciousness.²³ The openness and academic freedom of US universities and the larger research establishment offer foreign adversaries seeking to gain access to cutting edge research on topics ranging from artificial intelligence to robotics. Although much of the attention has focused on China, specifically its Thousand Talents Recruitment Plan²⁴, a range of US adversaries seek to use non-traditional collection methods to gain access to information, intellectual property, and critical technologies through insiders working at US universities and research institutions.²⁵ Mike Orlando, the acting Director of the National Counterintelligence and Security Center, explains that "we've seen a pivot to these non-traditional collectors, which are students, researchers . . . people who have legitimate jobs, who act as proxies or surrogates for the intelligence service. They generally don't have formal training, but they're able to operate using their cover or using their legitimate cover to then work for [or provide information to] the intelligence service."²⁶ A 2019 FBI report underscores the risk: "Foreign state adversaries seek to illicitly or illegitimately acquire US academic research and information to advance their scientific, economic, and military development goals. By doing so . . . they can achieve generational advances in technology."²⁷

This type of infiltration poses a significant threat to the interests of the USG, as our adversaries seek to erode and eclipse our competitive edge. Breaches can include theft of data, lab equipment, software (source code), as well as information about access protocols to classified areas, equipment specifications, prototypes and blueprints of new inventions and tacit knowledge about how to operationalize key concepts.²⁸ Identifying insiders in academic environments that might be providing this information to our adversaries is difficult given the norms of inclusiveness and transparency that underpin our research establishment. It is even more difficult to predict when and where a threat might materialize because universities employ fewer vetting and monitoring techniques than the military and intelligence services do. As such, it makes this issue a prime candidate for risk assessment using the CSIFT tool.

²² <https://www.justice.gov/opa/pr/harvard-university-professor-and-two-chinese-nationals-charged-three-separate-china-related>

<https://www.cnbc.com/2020/06/09/harvard-professor-indicted-for-false-statements-about-work-in-wuhan.html>

²³ <https://cis.org/Report/Stopping-Chinese-Infiltration-US-Educational-and-Research-Institutions>, <https://www.nbcnews.com/news/china/american-universities-are-soft-target-china-s-spies-say-u-n1104291>

²⁴ <https://www.hsgac.senate.gov/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China%27s%20Talent%20Recruitment%20Plans%20Updated2.pdf>

²⁵ <https://www.cbsnews.com/news/foreign-espionage-threats-u-s-intelligence-matters-podcast/>

²⁶ Ibid.

²⁷ <https://www.fbi.gov/file-repository/china-risk-to-academia-2019.pdf/view>


²⁸ Ibid.

We have developed two models to assess Academic Insider Risk: The first is a traditional decomposition, broken down into two key issues focused primarily on Academic Insider Risk challenges at the national and university levels. The second develops comparative question clusters to help predict where threats are originating, and what technologies are being targeted.

Academic Insider Risk (Espionage) Use Case: Key Issue Decompositions

Academic espionage is when foreign governments or adversaries “improperly [attempt] to obtain sensitive or proprietary information” from academia, without the knowledge of the originator of the information.²⁹ For this case study, we focus on predicting Chinese academic espionage at Stanford’s Artificial Intelligence Lab. China was chosen as the foreign actor because recent high-profile arrests and investigations suggest that China is actively attempting to gain information through academic insiders. Stanford’s Artificial Lab was chosen because China’s military fusion strategy emphasizes artificial intelligence as a key technology that will propel China into a leadership role in the coming century³⁰, and Stanford’s Artificial Intelligence Lab is one of the best in the nation. Questions derived from the decomposition process were customized to focus on China, artificial intelligence, and Stanford University. However, the questions can be customized to focus on other foreign adversaries, alternative technologies, and different universities.

Figure 11: Sample Use Case Decomposition: Academic Insider Risk Use Case



Who are the stakeholders, and why? Who are the forecasters?

The key stakeholders for this study would be Stanford’s Artificial Intelligence Lab, the federal grant making entities who are funding that research (NSF), and the Counterintelligence directorate focused on academic espionage.


²⁹ <https://www.nrc.gov/docs/ML1726/ML17262A027.pdf>

³⁰ <https://www.cbsnews.com/news/foreign-espionage-threats-u-s-intelligence-matters-podcast/>

Key Issue 1: National Threat, Vulnerabilities, and Countermeasures


This key issue is designed to capture the extent of the foreign threat to universities and how effectively US policies -- at the national level -- are countering that threat. It looks at vulnerabilities and countermeasures in our immigration system, within federal grantmaking and research institutions such as the NSF and the NIH, and our export control policies.

Figure 12: Sample Key Issue Decomposition: National Threat, Countermeasures, and Vulnerability



Key Issue 2: University Vulnerabilities

This key issue is designed to capture how vulnerable a specific university and department is to insider threats. Pivotal factors within this key issue area include training and awareness, the efficacy of the University's insider threat program, cyber security, institutional connections to the University has to the foreign actor, and direct indicators of academic espionage investigations and arrests.

Figure 13: Sample Key Issue Decomposition: University Vulnerabilities

Academic Insider Risk (Espionage) Comparative Question Clusters

One of the interesting findings to emerge from our interviews with experts in academic espionage was that two of the key challenges in countering this threat was understanding where the threat was coming from (which foreign actors were targeting academic institutions in particular technologies) and what was being targeted (which technologies were specific actors focusing on). In response to this insight, we developed two comparative clusters for academic insider risk. The goal of these clusters is twofold: First, to provide decision makers with forecasts to assess which states will pose the largest threat of academic espionage in particular subject areas to aid in resourcing decisions and crafting a strategic response to counter the threat. Second, to provide decision makers with forecasts to assess which subject areas are the target for academic espionage by particular actors in order to help them understand the priorities of our adversaries when crafting a strategic response.

Who are the stakeholders, and why? Who are the forecasters?

Stakeholders for these clusters would include the Counterintelligence Department at ODNI, as well as DOD, Department of State, NSC, and NIC strategists focused on specific countries/regions. The forecasting pool would be like that recommended for the academic insider risk decomposition: Public forecasters making predictions on a public forum and counterintelligence/FBI professionals forecasting on a classified forecasting platform. We would also recommend that a pool of forecasters be recruited to participate on their own platform or on a public platform from researchers and students in the top 10 universities for the relevant technology.

Figure 14: Sample Decomposition: Technology Comparison




Note: for both comparative frameworks, questions and resolution data has the potential to be sensitive or classified, so we would recommend this platform be hosted on an appropriate network.

Design of the Decision-maker Dashboard



The decomposition method for each of the three selected use cases (domestic extremism, workplace violence, and academic insider risk (espionage)) generated over 100 different questions each. Responses to these questions will have the potential to provide extensive insights. However, looking ahead towards future use of the CSIFT platform by the IC/DoD, this amount of quantitative data could be extremely time-consuming for a decision-maker to sift through, to quickly assess overall level of threat... especially because each of the 100 forecasting question streams changes daily. Forecasting data must therefore be synthesized and visualized for the decision-maker in a way that shows changes to the larger threat landscape, while also providing insight as to what is prompting the change, so that they can craft a strategy to respond proactively.


To this end, the ARLIS and Cultivate Labs team sought to design a dashboard for the decision-maker, which consolidates and visualizes changing data in an easily digestible format.

To develop the wireframe of the dashboard, the team considered the potential for two different end users of the dashboard: *policymakers*, who steer long term policy, and *operational decision makers*, who are on the ground, making operational decisions, such as commanders. We expect these types of users to be more interested in a high-level view of the CSIFT outputs, which are in the “By Issue” views of the dashboard, shown below in Figures 15 & 16.

Figures 15 - 18: Dashboard wireframe – Policymaker / Operational Decision Maker view

In this view, each square represents one forecast question, and users can hover over individual squares to obtain further information about that question. Users are alerted to questions with significant increases or decreases in probability (see label "A" above), can see how many questions are getting better or worse (label "B"), and can scroll/play to see snapshots of previous dates' heat maps (label "C"). They can select questions which are open or resolved, view percent change in consensus probability in the last 7 or 30 days, sort questions by numerical or color intensity (where deeper colors represent higher probabilities, label "D") and filter by different cohorts of forecasters (label "E"), which will update the probabilities accordingly.








End User 2: Analyst

The second category of executive dashboard users are traditional *analysts*, who analyze and track information at a deeper level. We felt that these users would be better suited with a view providing a more thorough, detail-oriented analysis of the individual signals on the CSIFT platform, which we display in a “By Signal” view as shown in Figures 19-22.

Figures 19 - 22: Dashboard wireframe – Analyst view







Recommendations

With a problem as complex and multifaceted as the estimation and anticipation of insider threat events, decision makers may benefit from insights that crowdsourced forecasting may provide. Further, focused crowd-sourced forecasting could help the IC/DoD guide resource allocations decision-makers.

Crowdsourced forecasting can supplement and enhance alternate estimation processes when modeling may not be effective, or when broad, diverse sets of perspectives would be beneficial. While the approach is primarily known for its ability to increase forecast accuracy, it has other notable benefits:

- **Explainability** - quantitative forecasts are accompanied by qualitative, written rationales and risk assessments that can be shared and analyzed in simple, digestible formats
- **Speed** - initial forecasts by a large crowd can be produced in hours vs. weeks or months
- **Track Record** - the platform creates an audit trail of individual and consensus forecasts over time, allowing for improved methods and identification of topics and question frameworks that perform particularly well
- **Focus** - forecasting skill is distinct from - and complementary to - domain expertise, providing an additional perspective on unanticipated and emergent threats
- **Culture** - with quantitative forecasts that easily scale across the US government, industry, and academia, the platform enables unprecedented diversity and breadth of judgments, and can highlight both consensus and important disagreements within and across different organizations.

We recommend deployment of the CSIFT platform for the mitigation of insider risk, focused on the identified Use Cases and Key Issues, and leveraging the identified forecasting pools. Further, given the potential sensitivity of the issues decomposed in the CSIFT platform, we recommend careful consideration of which network is most appropriate for each use case.

ACKNOWLEDGEMENTS

This report was prepared for [Office of the Undersecretary of Defense for Intelligence and Security (OUSD(I&S)), United States Department of Defense (DoD)] under the following agreement:

HQ003420F0655, *University of Maryland*, “Insider Threat and Personnel Vetting.”

DISCLAIMERS

Any views, opinions, findings, conclusions, or recommendations expressed in this publication do not necessarily reflect the views of an official United States government position, policy, or decision. Additionally, neither the United States government nor any of its employees make any warranty, expressed or implied, nor assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, product, or process included in this publication.

Certain commercial entities, equipment, or materials may be identified in this document to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the Applied Research Laboratory for Intelligence and Security (ARLIS), the University of Maryland, or the United States government, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

ABOUT ARLIS

Applied Research Laboratory for Intelligence and Security (ARLIS) is a UARC based at the University of Maryland College Park and established in 2018 under the auspices of the OUSD(I&S). ARLIS is intended as a long-term strategic asset for research and development in artificial intelligence, information engineering, acquisition security, and social systems. One of only 14 designated United States Department of Defense (DoD) UARCs in the nation, ARLIS conducts both classified and unclassified research spanning from basic to applied system development and works to serve the U.S. Government as an independent and objective trusted agent.

Technical Points of Contact:

PI: Adam Russell, D.Phil.
Chief Scientist, ARLIS
301.226.8834; arussell@arlis.umd.edu

Task Lead: Monique Beaudoin, PhD
Associate Research Scientist, ARLIS
301-226-8869; mbeaudoin@arlis.umd.edu

Co-PI: Kelly Jones, Ph.D.
Assistant Research Scientist, ARLIS
301.226.8850; kjones@arlis.umd.edu

Administrative Points of Contact:

Ms. Monique Anderson
Contract Officer, Office of Research Administration
Assistant Director, ARLIS
301.405.6272; manders1@umd.edu

APPENDICES

A.1: Quantitative methodologies used in aggregative forecasting

A Brier Score is a way to judge and score the accuracy of probabilistic forecasts. For instance, if a forecaster says that there is a 90% chance that the Cubs will win the World Series in 2017 (a probabilistic forecast), and then the Cubs subsequently win the World Series, the forecaster was "mostly" correct. But the forecaster also implicitly said that there was a 10% chance that the Cubs would *not* win, so that 10% allocation was "wrong." How do we quantify *how* "correct" the 90%/10% forecast was? A Brier Score gives us a way to score these forecasts so that we can compare their accuracy.

Calculating a Brier Score

A Brier Score will often be referred to as the "mean squared error of a forecast." Below is a discussion of what this means.

Assume that one week before the end of the World Series, a forecaster estimates a 90% chance that the Cubs win (and, the World Series ends with the Cubs winning). To calculate a Brier Score, we start by calculating the "squared error" for each of the 7 days that this forecast was valid. To calculate the squared error, we simply square the difference between my forecast and the actual outcome (100% if the Cubs win, 0% if they lose):

$(0.9 - 1.0)^2 = \mathbf{0.01}$ (where 0.9 captures the 90% win forecast and 1.0 captures the *actual* 100% win outcome)

And for the forecaster's implicit "10% Cubs lose" forecast:

$(0.1 - 0.0)^2 = \mathbf{0.01}$ (where 0.1 captures the 10% lose forecast, and 0.0 captures the *actual* 0% lose outcome)

Therefore, the forecaster's total error for the day:

$0.01 + 0.01 = \mathbf{0.02}$

To calculate the mean squared error, calculate the squared error for each day that my forecast was running and average them:

Day	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Forecast (win, lose)	0.9, 0.1	0.9, 0.1	0.9, 0.1	0.9, 0.1	0.9, 0.1	0.9, 0.1	0.9, 0.1
Error (only calculable after forecasting period has ended and	0.02	0.02	0.02	0.02	0.02	0.02	0.02

*question
has
resolved)*

We then calculate the average (a.k.a. mean) daily squared error to obtain a final score. In this case, since all the days' errors are the same, the forecaster's mean squared error is equal to **0.02**.

Updating a Forecast

What would have happened to the forecaster's Brier Score if he/she had updated their forecast mid-way through the week? Maybe the Cubs won a game on Day 3, leading them to update their forecast to 95% on Day 4. For any days where they had a 95% forecast, the squared error would then be:

$$(0.95 - 1.0)^2 = 0.0025$$

And their new, implicit Cubs-lose forecast of 5%:

$$(0.05 - 0.0)^2 = 0.0025$$

So, their total error for those days:

$$0.0025 + 0.0025 = 0.005$$

And their daily squared error would be:

Day	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Forecast	0.9, 0.1	0.9, 0.1	0.9, 0.1	0.95, 0.05	0.95, 0.05	0.95, 0.05	0.95, 0.05
Squared Error	0.02	0.02	0.02	0.005	0.005	0.005	0.005

Expanded, then, the calculation of forecaster's mean squared error after updating their forecast halfway through the week would be:

$$(0.02 + 0.02 + 0.02 + 0.005 + 0.005 + 0.005 + 0.005) / 7 = 0.0114$$

In the world of Brier Scores, *a lower Brier Score (i.e., less error) is more desirable/better*; therefore, by updating their forecast mid-way through the week, the forecaster improved their Brier Score from **0.02** to **0.0114**.

Further reading:

[Cultivate Labs | What are net Brier Points and how are they calculated?](#)

B1: Aggregation strategies

Default/Starting Aggregation Strategy

When a platform first launches or doesn't have the critical mass mentioned above, the following aggregation strategy has proven sufficient:

- Mean - average of the most recent forecast from each user;
- Establishing a minimum number of forecasts before applying any forecast decay;
- Setting a maximum number of forecasts to be included in any aggregation;
- Setting a minimum number of days' worth of forecasts to include in aggregation, never decaying away very recent forecasts; and
- Once past the minimum number of forecasts, only keep a subset of the latest forecasts for the current aggregation calculation.

Decay Strategies

Given a set of forecasts, a decay strategy will eliminate some subset of those forecasts prior to applying the aggregation strategy to calculate a consensus. Decay options include:

- Constant - weights newer forecasts more heavily and reduces the influence of older forecasts
- N-Recent - keeps the most recent N forecasts, discards the rest
- Percent Recent - keeps the most recent x% of forecasts, discards the rest
- No Decay - keeps all forecasts

Weighting

Forecasts can be weighted more or less heavily in the consensus calculation based on:

- Accuracy - historical accuracy/Brier score of the forecaster
- Frequency - how often the forecaster updates his/her forecast in the question
- Training - regardless of whether the forecaster has completed a training course