

JULIA ARNOLD, ABBY DOLL, RICHARD S. GIRVEN, KARL P. MUELLER, GREGORY SMITH

Infinite Potential—Insights from the Societal Warfare Games

After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises

This research was independently initiated and made possible by RAND's own resources, including income from operations, unrestricted gifts from RAND supporters, and distributions from its long-term investment fund.

Research Report

For more information on this publication, visit **www.rand.org/t/RRA4626-2**.

About RAND

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

Research Integrity

Our research integrity is grounded in RAND's core values of quality and objectivity. Rigorous quality assurance procedures, conflict of interest screening, and transparency in funding ensure that every study is objective and nonpartisan. Learn more at www.rand.org/integrity.

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

Published by the RAND Corporation, Santa Monica, Calif.
© 2026 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

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

RR-A4626-2

About This Report

The RAND Center for the Geopolitics of Artificial General Intelligence is conducting a series of scenario-based games to better understand how governments and other institutions should prepare for and respond to potential artificial intelligence (AI) and artificial general intelligence (AGI) developments. These games are run through the Infinite Potential platform, and this report is part of a series of publications that present findings and insights derived from player feedback.

Each Infinite Potential game presents a plausible future scenario in which advances in AGI outpace governance. Players are asked to craft immediate and longer-term responses and to identify the capabilities and playbooks they would need to effectively respond. This report synthesizes observations from seven games of the Societal Warfare scenario, conducted between June 2025 and April 2026. The Societal Warfare scenario explores how domestic unrest and violence related to advanced AI and AGI could be strategically exploited by adversaries. In the report, we identify the key issues, desired capabilities and playbooks, and other recommendations that participants believed would be necessary to navigate a period when social legitimacy, economic stability, and technical control might collapse. These findings aim to help policymakers, researchers, corporate leaders, civil society organizations, and the public understand the preparations needed before such an event, the decisions likely to matter most once that event begins, and the playbooks required to navigate the event successfully.

Center for the Geopolitics of Artificial General Intelligence

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This effort was undertaken by the division's Center for the Geopolitics of Artificial General Intelligence (AGI), which is committed to helping decisionmakers understand, anticipate, and prepare to navigate the national security and geopolitical implications of AGI. The center convenes leading technologists, strategists, economists, political scientists, and outside experts to consider the feasibility and effectiveness of prospective AGI-enabled capabilities; the domestic and international implications of their use; and the strategies and policies that governments, businesses, and civil society could adopt to respond to new realities. For more information, visit www.rand.org/geopolitics-of-agi.

Funding

This effort was independently initiated and made possible by RAND's own resources, including income from operations, unrestricted gifts from RAND supporters, and distributions from its long-term investment fund. A complete list of donors and funders is available at www.rand.org/about/how-

we-are-funded. RAND clients, donors, and grantors have no influence over research findings or recommendations.

Acknowledgments

We acknowledge the thoughtful and valuable peer reviews of our colleagues, Anita Chandra and Michael Decker. We also thank our players, who took time from their busy schedules to participate in the exercises documented here. Without their participation, this publication would not be possible.

Summary

Infinite Potential is a RAND platform that uses short, scenario-based games to better understand how governments and other institutions should prepare for and respond to potential artificial intelligence (AI) and artificial general intelligence (AGI) crises. In this report, we synthesize observations from seven games of the Societal Warfare scenario, conducted between June 2025 and April 2026. The Societal Warfare scenario explores how domestic unrest and violence related to advanced AI and AGI could be strategically exploited by adversaries. Table S.1 provides an overview of the six key issues that emerged across the games and the actions that the players identified as most important for improving preparedness and response.

Table S.1. Summary of Key Issues, Desired Capabilities and Playbooks, and Other Recommendations

Key Issue 1: Rapid AGI development presents the United States with a difficult trade-off between the strategic and economic benefits of racing ahead and the domestic instability that accelerated deployment can create.	
Desired capabilities	• U.S. government visibility into frontier development across the private sector • Compute resilience and supply-chain resilience • Stronger protection frameworks for frontier AI firms, major data center sites, and related dependencies
Desired playbooks	• Pace and posture • Concentration risk • AI critical-infrastructure protection
Key Issue 2: Foreign information operations that exploit domestic grievances could blur the lines among protest, influence, intrusion, and attack.	
Desired capabilities	• Rapidly attributing hybrid foreign influence campaigns • Better information-sharing across U.S. government agencies, including intelligence, cyber indicators, threat reporting, and other relevant data • Identifying and attributing adversary-backed efforts to influence U.S. policy through domestic organizations or initiatives • Identifying and managing hybrid actors
Desired playbooks	• Hybrid foreign influence campaign characterization • Hybrid actor
Other recommendations	• Develop clearer internal response thresholds for recurring elements of hybrid foreign influence campaigns
Key Issue 3: The United States' toolkit for responding to adversary societal warfare campaigns could prove structurally inadequate.	
Desired capabilities	• Asymmetric responses tailored specifically to hybrid foreign influence campaigns • Structured partnerships with frontier AI developers and major platforms • Escalation management
Desired playbooks	• Asymmetric response • AI incident communications • Escalation management
Other recommendations	• Make more-deliberate investments in precrisis relationships with frontier AI developers and major social media platforms • Develop a clear internal escalation strategy for hybrid campaigns
Key Issue 4: Declining public trust in AI could undermine national resilience faster than counter-disinformation efforts can rebuild it.	
Desired capabilities	• Workforce retraining and economic adjustment programs • Targeted support for communities affected by AI infrastructure and labor disruptions • AI literacy campaigns • Credible safety commitments from frontier AI developers
Desired playbooks	• Rebuilding public trust • AI incident public communications • Frontier AI firm engagement
Other recommendations	• Treat public trust as a national security variable

Key Issue 5: Existing authorities and federal-state coordination structures may be poorly prepared to respond to hybrid foreign influence campaigns that cross jurisdictions and levels of government.	
Desired capabilities	• Clearer alignment of authorities across departments and agencies for foreign influence campaigns • Federal-state coordination in response to the effects of foreign influence campaigns
Desired playbooks	• Government authorities and coordination • Federal-state coordination
Other recommendations	• Revisit authorities, designations, and coordination mechanisms as a package rather than piecemeal
Key Issue 6: Adversary information operations could increasingly target both the United States and its allies simultaneously.	
Desired capabilities	• Joint attribution, declassification, and public exposure of hybrid campaigns with allies • Allied alignment on frontier export controls and supply-chain diversification • Allied resilience-building against information operations and hybrid campaigns
Desired playbooks	• Joint or coordinated allied attribution and response to information operations • Allied export controls and supply-chain • Allied resilience
Other recommendations	• Dedicate allied institutions to coordinate preparation and response to AGI-era hybrid threats

Contents

About This Report iii
Summary v
Figure and Tables ix

Infinite Potential—Insights from the Societal Warfare Games: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises 1
Applying Gaming to the Problem of Artificial General Intelligence: Infinite Potential's Origin 1
Methodology 2
Key Issues 5
Next Steps 17

Appendix 18
Other Infinite Potential Publications 18

Bibliography 20
About the Authors 21

Figure and Tables

Figure

Figure 1. Infinite Potential Gameplay and Analytic Process 5

Tables

Table S.1. Summary of Key Issues, Desired Capabilities and Playbooks, and Other Recommendations vi
Table 1. Societal Warfare Scenario Summary 4
Table 2. Games and Participants 4
Table A.1. Other Infinite Potential Publications 18

Infinite Potential—Insights from the Societal Warfare Games: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises

Applying Gaming to the Problem of Artificial General Intelligence: Infinite Potential's Origin

Artificial intelligence (AI) has become a growing topic of interest and concern in the U.S. policymaking community, alongside increasing attention to artificial general intelligence (AGI) and its potential to affect areas critical to national security, from economic growth to military power.[1] However, although AGI's potential impact is increasingly recognized, it remains hypothetical, and exactly how to define it is contested. Some definitions describe AGI as an AI system able to perform every intellectual task that a human can perform at human level or better, while others emphasize an AI's ability to rapidly reach superhuman performance. The world has not yet seen AGI, and there is no consensus on what it is.[2] This uncertainty, combined with AGI's potentially far-reaching impacts, means that much of the response to AGI may occur only once the technology, or technologies like it, have been developed and deployed.

But uncertainty does not make analysis less important. On the contrary, AGI's implications must be examined under conditions of uncertainty. One way to explore such uncertainty is through gaming, which can help researchers examine policy problems when the relevant relationships are unclear or when the most important factors are not yet fully understood.[3] In scenario-based tabletop games, players walk through an emergency or crisis situation and discuss how they would respond based on prompts and scenario developments. This type of game allows researchers to explore problems in

---

[1] Jim Mitre and Joel B. Predd, *Artificial General Intelligence's Five Hard National Security Problems*, RAND Corporation, PE-A3691-4, February 2025.

[2] For an exploration of different ways to define AI capabilities, see Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv, arXiv:2311.02462, last updated September 24, 2025.

[3] Edward Parson, "What Can You Learn from a Game?" in Richard J. Zeckhauser, Ralph L. Keeney, and James K. Sebenius, eds., *Wise Choices: Decisions, Games, and Negotiations*, Harvard Business School Press, 1996.

collaboration with players, drawing on their decisions and discussions to identify how a problem is structured, which factors may prove most important, and how those factors interact.

The Infinite Potential platform was launched to use scenario-based tabletop games to analyze a key challenge for AGI policymakers: Because AI is evolving rapidly, and an AGI-related crisis could have uncertain, widespread, and severe impacts, it is difficult to address them through preemptive policymaking alone. Without adequate preparation, AGI policy might evolve through response to specific crises and incidents, which can lead to reactive, piecemeal, and short-term solutions that fail to address underlying risks or prevent future crises. Proactive contingency planning, by contrast, enables more effective, coordinated, and timely responses, which can reduce disruptions, minimize negative impacts, and support better long-term outcomes for society.[4] To support this kind of contingency planning, the Infinite Potential games seek to answer the following three key questions to provide insights on how the U.S. government and other institutions could better prepare for an AI- or AGI-facilitated crisis:

1. What key issues were players trying to make sense of in the face of AI or AGI-related crises, and what information did they seek to help them make judgments about those issues?
2. What capabilities did players wish they had available or wished to develop in response to the crisis?
3. What playbooks did players indicate need to be written and validated to adequately respond to the crisis presented to them?

The answers to these questions point to actions that might be worth taking in advance of a similar AI or AGI crisis. The issues that players identified, the judgments they sought to make, and the questions they asked imply areas of uncertainty that could be clarified before a crisis occurs. The capabilities and playbooks they desired likewise suggest areas for investment and development by the U.S. government or other actors. In this report, we synthesize findings across multiple games and scenarios to highlight recurring issues, capabilities, and playbooks that might help inform future preparation and response. It concludes with an appendix listing other Infinite Potential publications, which directs readers to additional reports and resources from the Infinite Potential platform.

Methodology

Game Format

Infinite Potential uses two-hour tabletop games to examine how players respond to plausible AI- and AGI-related crises. The games draw on RAND's "Day After" methodology, which was designed to support strategic decisionmaking under uncertainty by confronting players with crises and asking those players to reason through their implications, consequences, and possible responses.[5] Each game

[4] Joel Predd, "Riding the Coming Wave: The Urgent Need for Contingency Plans for AGI Futures," RAND Corporation, WR-A3691-1, 2025.

[5] Marc Dean Millot, Roger C. Molander, and Peter A. Wilson, *"The Day After . . ." Study: Nuclear Proliferation in the Post-Cold War World*, Vol. I: *Summary Report*, RAND Corporation, MR-266-AF, 1993.

is designed to be easily accessible and repeatable, requiring no advance player preparation. The games typically include between ten and 16 players and last two hours.

The games unfold over two turns. In Turn 1, players receive intelligence and open-source information about an emerging AI- or AGI-related crisis. After a subject-matter expert briefs the group, players ask clarifying questions and discuss policy objectives, response options, and possible consequences. In Turn 2, the scenario advances—usually through escalation or a major change in circumstances—and players reassess the situation in light of the new developments.

The Games

Infinite Potential scenarios are generally set 12 to 18 months in the future. This time horizon keeps the technological and geopolitical environments broadly recognizable while allowing for plausible near-term advances in AI. Unless otherwise specified, players are asked to assume that existing institutions, laws, and international relationships remain largely intact. In this report, we synthesize observations from seven games of the Societal Warfare scenario that were conducted between June 2025 and April 2026.

The Societal Warfare scenario explores how domestic unrest and violence related to advanced AI and AGI could be strategically exploited by adversaries. In this scenario, players learn of a coordinated People's Republic of China (PRC) information operations (IO) campaign that exploits U.S. societal vulnerabilities and growing public distrust of AI to disrupt AGI development and deployment.[6]

The seeds of such unrest can be seen today. For example, a 2025 Pew Research Center survey found that a majority of adults (51 percent) are more concerned than excited about AI's impact on the country,[7] and other polls report that 77 percent of Americans surveyed do not trust businesses to use AI responsibly.[8] As the United States seeks to maintain leadership in AI development, sophisticated information warfare campaigns could seek to exploit or even manufacture these types of societal friction points. Recent U.S. threat reporting has described foreign actors using AI to spread misinformation. For example, OpenAI recently disrupted a covert IO campaign that attempted to use generative AI for an influence campaign, indicating that the set of tools and techniques used for influence campaigns is becoming faster, cheaper, and more tailored.[9]

Exploring these dynamics through a structured scenario helps anticipate vulnerabilities, identify early-warning indicators, and develop courses of action to build national resilience against efforts to erode public trust, stall innovation, or incite division at critical junctures in technological competition. Table 1 provides a brief summary of the Societal Warfare scenario.

---

[6] IO campaigns involve a coordinated set of activities designed to gain a decisive advantage in the information environment by influencing adversary perceptions, disrupting their information systems, and protecting friendly information assets.

[7] Brian Kennedy, Eileen Yam, Emma Kikuchi, Isabelle Pula, and Javier Fuentes, "How Americans View AI and Its Impact on People and Society," Pew Research Center, September 17, 2025.

[8] Julie Rey, "Americans Express Real Concerns About Artificial Intelligence," Gallup, August 27, 2024.

[9] Ben Nimmo, Kimo Bumanglag, Michael Flossman, Nathaniel Hartley, Jack Stubbs, and Albert Zhang, "Disrupting Malicious Uses of AI: An Update," OpenAI, October 2025.

Table 1. Societal Warfare Scenario Summary

Scenario Title	Scenario Synopsis
Turn 1: Early 2028	The intelligence community reports an ongoing PRC IO campaign linking AI integration to safety incidents, specifically exploiting two events: a Texas power outage accompanied by a viral hoax attributing the outage to an unsafe AI system and a parallel set of false unemployment-fraud accusations blamed on AI agents. The campaign operates through ostensibly organic domestic groups, including by funding a U.S. policy think tank, which capitalizes on genuine domestic qualms about AI that has arisen through technologically enabled disruption.
Turn 2: + five months	The PRC-enabled IO campaign escalates from information activity that stokes domestic dissent into physical attacks. A riot-turned-armed-attack strikes a major AI research and development facility in Boston and a sabotage attack hits a data center in Texas. Intelligence indicates that there is additional possible involvement by Russia's Unit 29155 (a clandestine military intelligence unit) that has been channeled through domestic actors, but attribution remains uncertain, and the line between foreign-directed sabotage and foreign-influenced domestic attacks is unclear.

Players were selected based on their experience, institutional knowledge, and expertise. For games that include RAND researchers, our player-selection process ensures a diverse mix of expertise and experience—including but not limited to law, economics, technology policy, security studies, and other relevant fields. Players assumed the roles of National Security Council Principals Committee members and had to confront multiple overlapping challenges, including a PRC IO campaign exploiting genuine domestic grievances about AI, AI-driven economic disruption and labor displacement, foreign-influenced domestic movements operating within First Amendment constraints, and the tension between countering adversary operations and sustaining the pace of domestic AI development. Players were asked to engage in "light" roleplay: They represented their assigned positions but were not expected to mimic any specific individual. Table 2 lists when and where each game was conducted.

Table 2. Games and Participants

Date	Location	Participants
June 3, 2025	Washington, D.C.	RAND researchers
June 10, 2025	Washington, D.C.	RAND researchers
June 24, 2025	Washington, D.C.	RAND researchers
July 1, 2025	Washington, D.C.	Former federal government officials
December 16, 2025	Washington, D.C.	RAND researchers
April 1, 2026	Boston, Mass.	Law students
April 9, 2026	Boston, Mass.	Academics and economists

Data Collection and Analysis

To ensure reliable findings, a standardized data collection methodology is used for all Infinite Potential games. Each game is observed by a dedicated notetaker who uses a uniform template to capture player discussions in detail. This approach allows us to systematically track the key issues that are raised; the capabilities, playbooks, and other recommendations suggested; and any areas of disagreement or uncertainty. Immediately after each game, the notes are reviewed by the notetaker, a RAND analyst, and an Infinite Potential project lead to confirm accuracy and completeness.

After all the games referenced in this after-action report were completed, the notes were compared across games to identify recurring themes and recommendations. This cross-game analysis made it possible to highlight the most important issues and solutions that emerged repeatedly rather than those unique to a single game. Figure 1 illustrates this process, from data collection during each game to the synthesis of findings across the games included in this after-action report. The key issues, capabilities, playbooks, and other recommendations identified through this process are presented in the following sections.

Figure 1. Infinite Potential Gameplay and Analytic Process

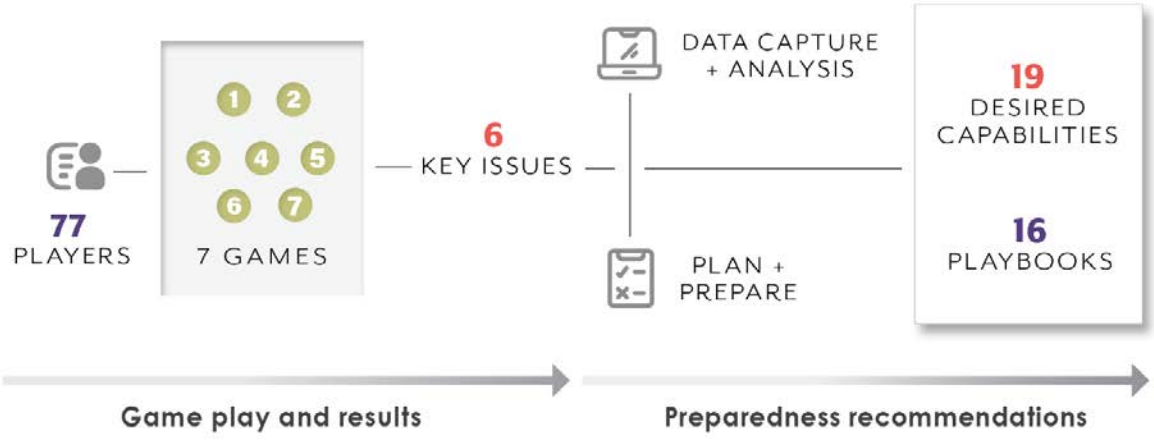

Key Issues

This section summarizes the key issues that emerged across the games regarding how the U.S. government might prepare for and respond to an AGI-related crisis that is part of a foreign influence or societal warfare campaign.[10] Although the issues raised are closely linked, they fell into six distinct problem areas: contending with the pace and pressures of AI development, characterizing and

[10] A *foreign influence campaign* is an effort by a foreign state or state-linked actor to shape public opinion, policy, or behavior. In this context, a foreign influence campaign is one that combines domestic grievances or dissent with foreign influence or infiltration. Foreign powers can exploit existing social or political tensions by amplifying, manipulating, or directing domestic groups to advance their own strategic objectives, often blurring the lines between a legitimate protest and foreign interference.

attributing hybrid foreign influence campaigns,[11] responding with a toolkit that might be structurally inadequate, sustaining public trust in AI, aligning authorities and federal-state coordination, and coordinating with allies that are facing similar campaigns.

Under each key issue, we identify the related capabilities, playbooks, and other recommendations that were raised across games. In this context, the term *capabilities* refers to things that would need to be built, established, or made available to enable effective action, such as systems, mechanisms, or other operational supports. The term *playbooks* refers to written guidance for action, including strategies, plans, protocols, regulations, or standard operating procedures. *Other recommendations* capture important findings or proposed actions that do not fit neatly into either of these categories.

Key Issue 1: Rapid AGI Development Presents the United States with a Difficult Trade-Off Between the Strategic and Economic Benefits of Racing Ahead and the Domestic Instability That Accelerated Deployment Can Create

Across all seven games, discussion returned repeatedly to the question of whether the United States should accelerate, slow, or *selectively gate* frontier AI development.[12] In some games, players framed the choice as a strategic dilemma: Moving quickly could ensure U.S. technological leadership, but moving more deliberately could buy time to manage labor-market disruption, strengthen public trust, and harden vulnerable infrastructure. In other games, players emphasized that accelerated AI development was a strategic necessity. They argued that slowing down in response to unrest would cede an advantage and suggested that emerging domestic unrest was itself a national security vulnerability that adversaries could exploit.

A related concern was *concentration risk*: the risk exposure created by heavy dependence on a limited number of frontier AI firms, compute resources, and infrastructure sites. The scenario's depiction of a single company as a near–single point of failure for U.S. frontier capability—and of data center buildout as a similarly concentrated piece of physical infrastructure—drew strong reactions in multiple games. Players suggested that a frontier ecosystem dominated by a small number of firms and sites would be easier for adversaries to target. Some also suggested that concentration might have shaped adversary targeting choices in Turn 1 and, in retrospect, made the targets of the physical attacks in Turn 2 predictable.

Key Issue 1: Desired Capabilities

Players repeatedly identified the need for the United States to have meaningful **visibility into frontier development across the private sector**. For example, players felt that the government would need to detect dangerous capability jumps and monitor concentration risk to inform timely

[11] A *hybrid* foreign influence campaign is one that includes multiple strategies, such as IO, cyber operations, and physical campaigns.

[12] In the context of frontier AI development, to *selectively gate* means to control access to the most advanced AI capabilities so that only certain organizations or users can deploy them, often under strict conditions.

decisionmaking. However, they also noted that such monitoring would need to be accomplished without driving development offshore or to less cooperative firms. During several games, players emphasized the importance of the U.S. developing **compute resilience and supply-chain resilience**, including diversifying fabrication, energy sources, and siting for large training and inference facilities.

Players also pointed to the need for **stronger protection frameworks for frontier AI firms, major data center sites, and related dependencies**. Across the games, players repeatedly treated key AI labs, compute clusters, data centers, utilities, and cloud dependencies as strategically significant but not yet fully integrated into existing protections for critical infrastructure. They discussed the need for stronger cybersecurity expectations, better mapping of physical and digital dependencies, clearer liability and disclosure expectations, and more-systematic arrangements for threat reporting and defensive coordination between government and the private sector.

Key Issue 1: Desired Playbooks

Players pointed to the need for a **pace-and-posture playbook** that would give senior decisionmakers a tool for weighing the risks and benefits of an accelerated pace of AI development against slowdown options at specific decision points rather than treating acceleration as a default. Such a playbook could link metrics, such as frontier capability indicators, adversary behavior, and domestic resilience,[13] to a graduated set of policy levers, including export controls, safety commitments, infrastructure investment, and public communications.

Players also suggested the need for a **concentration-risk playbook** to identify and manage firms and sites that are or risk becoming single points of failure or significant sources of risk. This playbook could include pre-identified continuity arrangements, hardening priorities, and conditions under which the U.S. government would treat private AI infrastructure as a protected national asset.

Players also expressed a need for **an AI critical infrastructure protection playbook** that would clarify when and how frontier AI firms, major data center sites, cloud dependencies, or related facilities should receive heightened protection-related treatment. Such a playbook could cover prioritization criteria, security expectations, disclosure obligations, government support mechanisms, and the implications for coordination with operators, suppliers, and local authorities.

Key Issue 2: Foreign Information Operations That Exploit Domestic Grievances Could Blur the Lines Between Protest, Influence, Intrusion, and Attack.

The IO campaign in the scenario included false or misleading claims about AI tied to real-world events, covert amplification of public anxiety about AI, involvement by domestic groups, adversary-linked influence entities operating in U.S. policy spaces, and possible participation by multiple foreign actors. These incidents and dynamics touched on areas that the U.S. government usually treats separately, including protected speech, foreign influence, cyber intrusion, and physical attack. As a

[13] Highlights from player discussions about domestic resilience are presented in the "Key Issue 4" section.

result, in every game, players struggled to identify when and to what extent foreign actors amplified and exploited genuine domestic grievances. Several challenges emerged:

- **Policy and disclosure gaps.** Players repeatedly noted that counterintelligence and disclosure frameworks might be insufficient to address the nuanced threats posed by foreign-funded entities operating within U.S. policy and public spaces. Several players suggested that adversaries were deliberately exploiting these gaps as part of a broader strategy of societal warfare.
- **Attribution and recognition.** The involvement of domestic players complicated attribution and recognition decisions. Although some foreign involvement, such as PRC activity, was considered plausible, players often lacked the clarity and timely information they wanted to support decisive public messaging or policy action. In several games, players debated whether the United States could or should act on partial attribution. Some players were particularly attentive to the risk that if the government publicly made an attribution claim that later was proven to be more complicated or disproven, that lack of clarity would further erode trust in government information.
- **Balancing response and civil liberties.** The presence of *hybrid actors*—domestic groups with genuine grievances but affected by foreign infiltration or influence—further complicated attribution. These actors' presence blurred the lines between protest and foreign influence, creating uncertainty about how the U.S. government should respond. Players highlighted the tension between protecting free speech and responding to foreign influence, noting the risks of both over-securitizing domestic protest and underestimating the threat posed by foreign actors exploiting these grievances.
- **Fragmented information-sharing and coordination.** The blurred boundaries among protest, influence, intrusion, and attack meant that incidents in the scenario often fell across multiple agencies or jurisdictions. Players recognized the struggle that agencies would have sharing information and coordinating responses because existing systems and authorities are not always equipped to handle threats that cross traditional lines.

Key Issue 2: Desired Capabilities

Player discussions pointed to the need for **rapid attribution of hybrid foreign influence campaigns**, meaning the ability to reach defensible decisions within hours or days rather than weeks or months or after the fact. That, in turn, would require **better information-sharing across U.S. government agencies**, including intelligence, cyber indicators, threat reporting, and other relevant data, so that information about risk can be assessed as a coordinated effort rather than in separate stovepipes.

Players also emphasized the need for **capabilities to identify and attribute adversary-backed efforts to influence U.S. policy**, including the ability to detect when ostensibly domestic organizations or initiatives are being supported, steered, or amplified by state-linked actors. They also identified a need to be able **to identify and manage** hybrid actors without collapsing the distinction between protected protest and foreign direction.

Key Issue 2: Desired Playbooks

Players pointed to the need for **a hybrid foreign influence campaign characterization playbook** that would walk decisionmakers through the categorization, attribution, and threshold judgments in cases in which foreign campaigns exploit domestic vulnerabilities through multiple forms of activity. Such a playbook could link confidence levels in attribution to graduated public, diplomatic, and operational responses and could specify when partial attribution is sufficient to act. Players also suggested the need for a **hybrid actor playbook** for domestic groups with genuine grievances that are affected by foreign infiltration or influence. In these cases, a playbook could aid in distinguishing protected protest from foreign-directed activity and provide suggestions of proportionate responses that do not collapse or muddle the distinction.

Key Issue 2: Other Recommendations

Discussion across the games suggested that the United States might need **clearer internal response thresholds for recurring elements of hybrid foreign influence campaigns.** Examples of such activities include foreign funding of domestic influence vehicles, covert amplification of domestic groups, and IO targeting critical infrastructure. Several games suggested that without articulated thresholds, deterrence is difficult to establish, and adversaries are incentivized to continue probing.

Key Issue 3: The United States' Toolkit for Responding to Adversary Societal Warfare Campaigns Could Prove Structurally Inadequate

Across multiple games, players struggled to decide how to react because available U.S. response options did not appear to match the problem in front of them. Conventional diplomatic démarches, sanctions, indictments, and cyber responses were considered, but players felt that they were too slow, revealed too much, would not provide effective deterrence, or were too aggressive and risked escalating the situation too quickly. In some games, players speculated that adversaries were exploiting exactly these weaknesses by operating in ways that moved faster than the U.S. government agency response process.

A recurring asymmetry in tools available to the United States and its adversaries was also discussed. Across games, players noted that the PRC and Russia operate with a wider permissible toolkit, blending covert influence, deniable intermediaries, cyber activity, and pressure on domestic fault lines in ways that the United States could neither mirror nor readily answer. Players struggled to identify U.S. responses that were both effective and stayed within legal and ethical boundaries. In several games, players discussed asymmetric responses—responses that do not copy the adversary's tactics but instead impose costs through channels better aligned with U.S. legal authorities and comparative strengths. Examples discussed across games included exposing adversary-linked influence networks, disrupting financial flows and support to influence vehicles, and coordinating with allies to increase pressure without inviting direct retaliation. Players noted, however, that these options require prepositioned authorities, intelligence, and partner alignment that are not yet in place.

Across multiple games, players argued that the U.S. government's existing ad hoc relationship with frontier AI developers and major social media platforms was inadequate for responding to fast-

moving AI incidents. Some players favored public-private partnerships for information-sharing. They argued that to preserve public trust in such a relationship, there should be publicly articulated rules governing the information that could be shared, the conditions under which it could be used, and the oversights that would apply. These players warned that a more direct approach to information-sharing could undermine the cooperation needed to act at speed during fast-moving incidents. However, other players favored a more compulsory approach, arguing that the government could not rely solely on voluntary cooperation in a crisis. Regardless of this debate, players broadly agreed that AI-related incidents would require prearranged systems for communication and coordination between the U.S. government and AI developers and social media platforms rather than a reactive, ad hoc approach.

Key Issue 3: Desired Capabilities

Players repeatedly identified the need for the U.S. government to have **asymmetric response capabilities tailored specifically to hybrid foreign influence campaigns**. Such capabilities would allow the United States to respond effectively without simply copying adversary tactics or escalating too quickly. For example, players repeated noted wanting tools for disrupting the finances of foreign influence operations, such as blocking or freezing assets linked to adversary-linked influence vehicles; coordinating with allies to publicly expose foreign campaigns, including sharing information and issuing joint statements to reveal adversary activities and reduce their impact; and establishing rapid declassification procedures so the government can quickly release credible information to support attribution and build public trust.

Discussion also pointed to the need for the U.S. government to **work more closely with frontier AI developers and major platforms**, including the establishment of prearranged information-sharing mechanisms, joint incident-response mechanisms, and transparency standards for crisis conditions. Finally, players described **escalation management capabilities**, particularly the abilities to identify likely adversary countermoves, define off-ramps in advance, and maintain a shared internal framework for deciding when and how to escalate.

Key Issue 3: Desired Playbooks

Players desired playbooks that would support the capability gaps identified. Players pointed to the need for an **asymmetric-response playbook** that outlines diplomatic, financial, informational, cyber, and law-enforcement response options indexed to attribution confidence and campaign severity. Players also suggested the need for an **AI incident communications playbook** detailing how the U.S. government and developers and platforms would communicate during fast-moving IO events. In addition, players described an **escalation management playbook** that would map likely adversary reactions, identify off-ramps, and help senior leaders decide when to absorb friction, when to impose costs, and when to shift toward de-escalation.

Key Issue 3: Other Recommendations

Player also suggested that the United States should more deliberately invest in **precrisis relationships with frontier AI developers and major social media platforms** rather than rely on reactive coordination during incidents. In several games, players also expressed that the lack of a **clear**

internal escalation strategy for hybrid campaigns left senior decisionmakers without a shared frame of reference, leading to repeated relitigation of basic questions about objectives, thresholds, and acceptable risk during the games themselves.

Key Issue 4: Declining Public Trust in AI Could Undermine National Resilience Faster Than Counter-Disinformation Efforts Can Rebuild It

Across all games, players treated the domestic information environment as a strategic variable, not a backdrop. The *domestic information environment* refers to the overall climate of public trust, understanding, and media narratives, as well as the population's susceptibility to false or misleading claims. Players suggested that the hybrid foreign influence campaign appeared effective not because its claims were convincing but because the targeted population was already inclined to distrust the government and AI firms and to believe that they would not benefit from or might even be harmed by AI. Several players observed that to simply counter false narratives with messaging would be ineffective in an atmosphere of deep public mistrust. Instead, they emphasized that any credible response would need to address the underlying economic grievances being exploited by adversaries. In these discussions, players saw addressing substantive issues as a requirement for successful public messaging, not just an alternative or addition to it.

Players also discussed several reasons for diminishing public trust in the context of rapid AI development. They specifically noted that the pace of AI innovation is outstripping the country's ability to absorb its broader impacts, such as workforce dislocation, regional economic shocks, and uneven public understanding of AI. Players saw these impacts as vulnerabilities that could easily be exploited by foreign influence campaigns. Players noted a mismatch between the speed of AI advancement (months to a few years) and the slower timelines for retraining the workforce, adapting infrastructure, and rebuilding public trust (all of which could take years to a decade or more).

In nearly every game, players identified workforce dislocation and regional economic shocks as major drivers of public distrust toward both the U.S. government and frontier AI firms. Players noted that claims that AI firms profit at the expense of displaced workers resonate with the public, tapping into existing anxieties about AI and job loss. Across all games, players agreed that U.S. counter-messaging should highlight concrete government actions, such as worker retraining programs and support for affected communities, to address these concerns. Without substantive responses, they felt that counter-disinformation efforts risk being seen as mere defensive posturing.

The behavior of frontier AI developers was also identified as shaping public trust. Players emphasized that credible safety commitments from these firms, backed by transparency and independent verification, are essential for restoring public confidence in the AI ecosystem. Players suggested that without such commitments, government counter-messaging is likely to be discounted by a skeptical public. Over the long term, players identified public AI literacy as a critical investment for national resilience, noting that the public's ability to distinguish synthetic media from authentic media and understand AI capabilities is lagging deployment. They suggested that AI literacy should be prioritized alongside cybersecurity awareness, recognizing that benefits would be realized in years rather than weeks.

Key Issue 4: Desired Capabilities

Players repeatedly identified domestic programs as national security tools rather than as simple solutions to domestic policy issues. They noted that these programs can build public resilience and reduce the effectiveness of adversary narratives that exploit fears of displacement, unfairness, and loss of purpose. For example, players repeatedly highlighted the need for **workforce retraining and economic adjustment programs**. They said that such programs would help workers displaced by AI technologies quickly transition to new roles and industries, minimizing the impact of job loss and reducing society's vulnerability to adversary narratives about unfairness and loss of purpose. Players also emphasized the importance of faster retraining pipelines and benefit structures that are responsive to the changing labor market.

In addition to workforce-focused measures, players also expressed the need for **targeted support for communities affected by AI infrastructure and labor disruptions**. The types of assistance that players discussed included regional assistance for areas hosting major AI buildouts and flexible programs to help communities adapt to economic shifts caused by AI deployment. They noted that this capability would also require better reporting on AI's labor market effects to target the assistance to areas most in need, along with stronger federal, state, and local mechanisms for delivering support.

Discussion also pointed to a need for **AI literacy campaigns** in schools, workplaces, and public communications, and for mechanisms for credible safety commitments from frontier developers, including transparency, independent verification, and public reporting. Players described needing to develop long-term capacity to explain what AI systems can and cannot do, how they fail, how synthetic media can be identified, and what safeguards are in place in high-consequence settings.

Finally, players expressed concern that without visible and independently verified assurances, public defenses of AI systems would be perceived as self-serving or politically motivated. To address this, players identified a need for mechanisms to secure **credible safety commitments from frontier AI developers**. In several games, players suggested raising assurance standards, implementing benchmarking, conducting neutral auditing, increasing public reporting and transparency, and in some cases, using model provenance or verification to make claims about safety and reliability more credible to the public.

Key Issue 4: Desired Playbooks

The games also pointed to the need for a **playbook for rebuilding public trust** that would match substantive policy actions with counter-disinformation messaging. This playbook could help officials determine how to acknowledge legitimate grievances without validating false claims, identify the policy actions needed to reassure the public, and communicate openly about uncertainty and corrective action.

Discussion also pointed to the need for a **public communications playbook to guide responses to AI-related incidents**. This playbook could coordinate messaging across all levels of government—federal, state, and local. It could clarify leadership roles, set standards for public statements, and provide guidance on translating technical findings for nonexpert audiences.

Finally, players also called for a **playbook to guide engagement with frontier AI firms**. This playbook could include commitments for safety, transparency, and public reporting during periods of heightened concern. It could also specify the information that developers should share with

governments during incidents, identify credible forms of third-party validation or independent assessment, and provide guidance on communicating assurance measures to the public without creating additional security or competitive risks. During several games, players expressed that such a playbook could help manage the tension between working cooperatively with AI firms and imposing regulatory and oversight requirements by clarifying expectations in advance rather than improvising them during a crisis.

Key Issue 4: Other Recommendations

Player discussions also suggest that **public trust should be treated as a national security variable**, not just a residual or political issue managed through communications. They suggested that public trust merits senior-level attention, along with specific indicators and targeted investments. Because adversary campaigns might exploit existing distrust, players emphasized that building resilience before a crisis through proactive measures is likely to be more effective than relying solely on communication efforts during an incident.

Key Issue 5: Existing Authorities and Federal-State Coordination Structures May Be Poorly Prepared to Respond to Hybrid Foreign Influence Campaigns That Cross Jurisdictions and Levels of Government

Across multiple games, players described the U.S. legal and institutional architecture as poorly prepared to respond to the hybrid foreign influence campaigns presented in the scenario. Players repeatedly noted overlapping and convoluted authorities across departments and agencies, gaps between domestic and foreign-facing tools, and unclear lines of responsibility for hybrid threats. Players observed that several of the most consequential decisions in the scenario, including how to characterize a foreign-influenced domestic actor, whether to invoke emergency authorities, and how to coordinate with state and local actors, did not have clear owners.

Federal-state coordination emerged as a recurring concern in player discussion. Players observed that state and local actors would often be the first to encounter the visible effects of a foreign influence campaign, such as online activity, protests, and threats or attacks against AI infrastructure. However, they noted that existing federal coordination mechanisms are not well adapted to campaigns moving quickly across governmental jurisdictions and domains. As a result, they recognized that local officials might be forced to manage AI-related disruptions, public order concerns, and infrastructure risks without clear lines of communication to federal agencies or sufficient technical context about the systems involved. Coordination issues are also heightened by this being an AI-driven crisis. Players repeatedly discussed how AI's impacts in the scenario cut across labor displacement, misinformation, critical infrastructure, and foreign influence, which made it harder to assign clear ownership and authorities. Players also treated government dependence on frontier labs and social media firms, along with increasingly low levels of trust in AI, as complicating elements of the coordination problem itself.

Key Issue 5: Desired Capabilities

Players repeatedly identified the need for the **clear alignment of authorities across departments and agencies for foreign influence campaigns**, especially those that cross multiple domains. They emphasized the importance of having predesignated leads for hybrid campaigns. Across the games, players expressed uncertainty about who would lead public attribution, who would address domestic groups with possible adversary ties, what legal tools were available against adversary-linked influence efforts, and where the boundaries lay between criminal investigation, intelligence activity, and public-order response. They suggested a need not only for clearer doctrine but also for practical interagency mechanisms to quickly map responsibilities and operate effectively across domestic, foreign, cyber, and law-enforcement boundaries.

Federal-state coordination capabilities were repeatedly described as underdeveloped, especially given that the effects of foreign influence campaigns would be felt initially and most intensely by state and local actors. Players pointed to the need for prearranged points of contact, federal-state information-sharing mechanisms, and coordination mechanisms that could function even when the relevant infrastructure was privately owned or federal authority was limited. In some discussions, players also suggested a need for better state- and local-government understanding of the role of AI systems in critical infrastructure to prevent officials from responding to AI-related crises without the technical context needed for sound decisionmaking. Players suggested a number of possible entities to lead coordination across governance levels. On the federal side, players noted that the National Institute of Standards and Technology's Center for AI Standards and Innovation and the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency could lead efforts to develop and coordinate standards, ensure system reliability, and communicate information to the public. They also noted that the U.S. President would play a key role in setting national priorities and publicly endorsing the coordination effort. At the state and local levels, union leaders, mayors, veterans' groups, independent experts, academics, frontier labs, social media firms, and local media were identified as key players whose involvement would be essential for giving credibility to and facilitating the implementation of any coordination efforts.

Key Issue 5: Desired Playbooks

Players pointed to the need for a **government authorities and coordination playbook** to guide AI incident response involving domestic organizations. This playbook would map responsibilities to the appropriate departments, agencies, and state actors. It would also help decisionmakers distinguish between protests, sabotage, criminal conspiracy, adversary-backed influence activity, cyber intrusion, and hybrid campaigns, while ensuring civil liberty safeguards. It could also clarify which institutions should lead in ambiguous or fast-moving situations. Across games, players repeatedly expressed uncertainty about the sufficiency of existing authorities, when escalation to emergency measures was warranted, and which agency should lead in unclear cases. This playbook would not eliminate the need for judgment, but it could help leaders reach decisions faster and be more consistent.

Discussion also pointed to a **federal-state coordination playbook** for hybrid incidents targeting AI infrastructure or AI-related public order concerns. This playbook would include protocols for U.S. government agencies sharing threat information with governors' offices, state homeland security entities, public utility commissions, emergency managers, and local law enforcement; guidelines for

coordinating public messaging across levels of government; and procedures for managing incidents in which adversary influence, domestic protest, and infrastructure risk intersect.

Key Issue 5: Other Recommendations

Player discussions also suggest that **authorities and coordination mechanisms might need to be revisited as a package rather than addressed piecemeal**, given how often discussions surfaced interactions among them. Players noted that gaps or overlaps in jurisdiction, unclear lines of responsibility, and inconsistent escalation procedures can undermine the effectiveness of the response to hybrid campaigns. Addressing these issues together would help ensure that the roles, legal tools, and communications channels are aligned, enabling a more unified and agile response to complex cross-domain threats and incidents.

Key Issue 6: Adversary Information Operations Could Increasingly Target the United States and Its Allies Simultaneously

Across multiple games, players highlighted that the IO campaign that the PRC and Russia were running against the United States was comparable with campaigns being run against allies. They noted that the U.S. response would therefore be more effective if coordinated with partners rather than mounted unilaterally. Given the speed and the sensitivity of the underlying intelligence, players identified Five Eyes, the Group of Seven, and emerging AI governance forums as the venues most likely to support coordinated responses.

Coordinated response and exposure of PRC tactics was one of the most frequently discussed allied tools. Players suggested that joint public attribution of hybrid foreign influence campaigns, backed by a coordinated release of supporting information, would enhance credibility. They also noted that joint disruption of adversary-linked influence vehicles would be more effective and difficult for adversaries to dismiss than unilateral U.S. action. The scenario's foreign-linked think tank was sometimes discussed as an example of the kind of influence channel that might also appear in allied policy ecosystems.

Export controls and supply-chain coordination were discussed as potential additional tools for allied action. Players suggested that the effectiveness of frontier AI export controls depended on partner alignment and that adversary efforts to fracture allied positions—including through targeted economic pressure on individual allies—were recurring challenges. In several games, players highlighted that building allied resilience-building, including the diversification of compute, energy, and supply chains, was a precondition for sustained alignment on export controls. They noted that strengthening allied resilience and coordination would not only support export controls but also enhance the collective ability to resist societal warfare and IO targeting multiple allied countries simultaneously.

Key Issue 6: Desired Capabilities

Players repeatedly identified the need to develop **allied capabilities for joint attribution, declassification, and public exposure of hybrid campaigns and adversary-linked influence vehicles**.

Players expressed a need for practical mechanisms to share intelligence rapidly, align evidentiary standards, and decide what information can be released publicly without unacceptable cost to sources and methods.

Discussion also pointed to a need for **allied alignment on frontier export controls and supply-chain diversification**. Across games, players linked resilience against foreign pressure with reducing dependence on vulnerable suppliers, addressing concentrated compute bottlenecks, diversifying sources of critical minerals, and minimizing reliance on cloud service providers that adversaries could exploit politically or operationally. For example, players discussed allied cooperation on reindustrialization, decoupling in selected sectors, and shared approaches to protecting strategically important AI inputs and infrastructure.

Players also discussed a need for **allied resilience-building against IO and hybrid campaigns**. Across the games, players noted that democratic societies face similar vulnerabilities: low public trust, economic anxiety, platform manipulation, and difficulty responding quickly without overreaching. They discussed useful allied resilience-building strategies, including shared monitoring arrangements, shared technical investments in detection and authentication, and the ability to exchange lessons on public messaging, platform coordination, and response to synthetic media at speed. For all the capabilities requiring coordination among allies, players noted that standing channels, such as Five Eyes, the Group of 7, and emerging AI governance forums, are the most promising venues for coordination at speed.

Key Issue 6: Desired Playbooks

Players pointed to the need for a playbook to **guide joint and allied attribution and response to IO**. This playbook could outline plans for joint attribution, declassification, and disruption of adversary-linked influence vehicles, including prearranged roles and authorities. It could also specify the evidentiary standards needed for public responses, identify which governments should lead when incidents span multiple allied jurisdictions, and provide guidance on synchronizing public statements to ensure that responses are mutually reinforced rather than fragmented.

Players also discussed the need for playbooks on **allied export controls** and **supply chains** covering frontier compute, energy, and inputs. Such a playbook would identify where shared vulnerabilities exist, what diversification steps are most urgent, what measures should be coordinated in advance of a crisis, and how to respond if an adversary uses economic or supply-chain pressure to split allied positions.

Finally, players also spoke of the need for **an allied resilience** playbook to prepare for hybrid foreign influence campaigns that target allied AI ecosystems. This playbook could include coordination with major platforms operating across allied markets, public education and AI-literacy efforts, and procedures for exchanging best practices on countering synthetic media, bot amplification, and adversary-linked influence organizations.

Key Issue 6: Other Recommendations

Across games, players discussed the need for **dedicated allied institutions to coordinate preparation and response to AI-era hybrid threats** rather than relying on ad hoc arrangements. They noted that standing channels, such as Five Eyes, the Group of Seven, and emerging AI

governance forums, are the most promising venues for rapid coordination, especially because adversary campaigns often move faster than allied coordination cycles. Players also observed that adversary efforts to fracture allied positions, including through targeted pressure on individual allies, make shared resilience investments and shared exposure mechanisms more important.

Next Steps

Taken together, the discussions from these games suggest that existing preparedness for AI-enabled societal warfare is inadequate but that there are many defensive and preparatory measures that could be pursued. Across games, player discussion converged on a similar set of needs: contending with the pace and pressures of AI development, characterizing and attributing hybrid foreign influence campaigns, creating an adequate response toolkit, sustaining public trust in AI, aligning authorities and federal-state coordination, and coordinating with allies that are facing parallel pressures.

A useful next step for the Infinite Potential platform would be to explore domestic dynamics around advanced AI and AGI in greater depth. This includes examining how workforce displacement, distrust of frontier AI developers, and local impacts of AI infrastructure—such as effects on energy and water systems—create pressure on the U.S. government, even in the absence of foreign influence campaigns. Additional factors, such as inequality in the distribution of AGI's benefits and federal-state tensions over AI policy, also contribute to domestic unrest and could be explored. Additionally, it would be valuable to assess how U.S. response architecture and policy options perform when unrest is driven by domestic constituencies acting on their own rather than driven by foreign influence. Exercises focused specifically on these domestic dynamics would provide different insights than those gained from the Societal Warfare scenario, which primarily examined domestic unrest through the lens of foreign influence.

Additionally, drawing on the key issues raised by players and the capabilities and playbooks they identified as necessary during Infinite Potential games, RAND researchers are developing materials for AI education, training, and practical use. Building on this body of work and its other research, RAND's Center for the Geopolitics of AGI will continue its effort to equip governments, the private sector, and civil organizations with insights and tools to prepare for and respond to the uncertainty and risks posed by AGI.

Appendix

Other Infinite Potential Publications

Table A.1 shows other Infinite Potential publications that are available—or will be soon—at the time of this writing.

Table A.1. Other Infinite Potential Publications

Scenario or Theme	Publication Information
Business and Civil Society	Girven, Richard S., Beba Cibralic, and Gregory Smith, *Infinite Potential—Insights from the Business and Civil Society Scenario: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-A4767-1, 2026. As of June 22, 2026: https://www.rand.org/pubs/research_reports/RRA4767-1.html
Cyber Surprise	Smith, Gregory, George Hage, Chad Heitzenrater, Matt Chessen, and Richard S. Girven, *Infinite Potential—Insights from the Cyber Surprise Scenario: Post-Series Scenario Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-A4626-1, 2026. As of June 22, 2026: https://www.rand.org/pubs/research_reports/RRA4626-1.html
European Games	Stephanie Walsh, Marta Kepe, Salman Ahmed, Barbara Del Castello, Matt Chessen, Marzia Giambertoni, Cian Girven, Richard S. Girven, Theiline Pigott, and Anton Shenk, *Infinite Potential—Insights from the European Games: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-A4767-2, 2026. As of October 5, 2026: https://www.rand.org/pubs/research_reports/RRA4767-2.html
Loss of Control	Pauline Moore, Stephanie J. Walsh, Zachary Burdette, Matt Chessen, David R. Frelinger, Richard S. Girven, Chad Heitzenrater, Gregory Smith, and Barry Wilson, *Infinite Potential—Insights on Artificial Intelligence Loss of Control Across Five Scenarios: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*," RAND Corporation, RR-A4767-3, 2026. As of October 5, 2026: https://www.rand.org/pubs/research_reports/RRA4767-3.html
Robot Insurgency	Smith, Gregory, Benjamin Boudreaux, Michael J. D. Vermeer, and Joel B. Predd, *Infinite Potential—Insights from the Robot Insurgency Scenario: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-A4231-1, 2025. As of June 22, 2026: https://www.rand.org/pubs/research_reports/RRA4231-1.html
Two Moonshots	Predd, Joel B., Matt Chessen, and Gregory Smith, *Infinite Potential—Insights from the Two Moonshots Scenario: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-

Scenario or Theme	Publication Information
	A4230-1, 2025. As of June 22, 2026: https://www.rand.org/pubs/research_reports/RRA4230-1.html
Viral Uplift	Shenk, Anton, Matt Chessen, Barbara del Castello, Gregory Smith, and Richard S. Girven, *Infinite Potential—Insights from the Viral Uplift Scenario: After-Action Report from a Sequence of Day After Artificial General Intelligence Exercises*, RAND Corporation, RR-A4727-1, 2025. As of June 22, 2026: https://www.rand.org/pubs/research_reports/RRA4727-1.html

Bibliography

Johnson, Stuart E., Martin C. Libicki, and Gregory F. Treverton, eds., *New Challenges, New Tools for Defense Decisionmaking*, RAND Corporation, MR-1576-RC, 2003. As of September 19, 2025: https://www.rand.org/pubs/monograph_reports/MR1576.html

Kennedy, Brian, Eileen Yam, Emma Kikuchi, Isabelle Pula, and Javier Fuentes, "How Americans View AI and Its Impact on People and Society," Pew Research Center, September 17, 2025.

Millot, Marc Dean, Roger C. Molander, and Peter A. Wilson, *"The Day After . . ." Study: Nuclear Proliferation in the Post-Cold War World*, Vol. I: *Summary Report*, RAND Corporation, MR-266-AF, 1993. As of June 11, 2026: https://www.rand.org/pubs/monograph_reports/MR266.html

Mitre, Jim, and Joel B. Predd, *Artificial General Intelligence's Five Hard National Security Problems*, RAND Corporation, PE-A3691-4, February 2025. As of August 5, 2025: https://www.rand.org/pubs/perspectives/PEA3691-4.html

Morris, Meredith Ringel, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv, arXiv:2311.02462, last updated September 24, 2025.

Nimmo, Ben, Kimo Bumanglag, Michael Flossman, Nathaniel Hartley, Jack Stubbs, and Albert Zhang, "Disrupting Malicious Uses of AI: An Update," OpenAI, October 2025.

Parson, Edward, "What Can You Learn from a Game?" in Richard J. Zeckhauser, Ralph L. Keeney, and James K. Sebenius, eds., *Wise Choices: Decisions, Games, and Negotiations*, Harvard Business School Press, 1996.

Predd, Joel, "Riding the Coming Wave: The Urgent Need for Contingency Plans for AGI Futures," RAND Corporation, WR-A3691-1, 2025. As of August 21, 2025: https://www.rand.org/pubs/working_papers/WRA3691-1.html

Rey, Julie, "Americans Express Real Concerns About Artificial Intelligence," Gallup, August 27, 2024.

About the Authors

Julia Arnold is a research assistant at RAND. Her work focuses on the intersection of emerging technologies and national security, as well as the design of wargames to explore the strategic implications of transformative technologies. She holds a B.A. in international studies and political science.

Abby Doll is a senior political scientist at RAND. Her work focuses on military capability assessment and military strategy. She holds a Ph.D. in war studies.

Richard S. Girven is a senior international defense researcher at RAND. He leads the Infinite Potential gaming platform and its engagement with senior decisionmakers across government and industry. He holds a master's in military arts and sciences in strategic studies and an M.S. in national resource strategy.

Karl P. Mueller is a senior political scientist at RAND. His research focuses on deterrence and coercion, airpower, escalation, crisis stability, and wargaming. He holds a Ph.D. in politics.

Gregory Smith is a policy analyst at RAND. His research focuses on emerging technologies, such as artificial intelligence, and their implications for U.S. competitiveness and national security. He holds a J.D.