Research Report

RICHARD S. GIRVEN, BEBA CIBRALIC, 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

For more information on this publication, visit **www.rand.org/t/RRA4767-1**.

Published by the RAND Corporation, Santa Monica, Calif.

RR-A4767-1

About This Report

To understand how the United States should respond to and prepare for potential artificial intelligence (AI) and artificial general intelligence (AGI) developments in the future, the RAND Center for the Geopolitics of AGI is conducting a series of Day After AGI exercises using the RAND Infinite Potential platform. Each exercise presents a plausible future in which advances in AGI outpace governance, tasking participants to craft immediate and longer-term responses.

The Business and Civil Society scenario broadened the experiment beyond government officials. Rather than roleplaying U.S. National Security Council members, participants played themselves—chief executive officers, entrepreneurs, investors, technologists, philanthropists, academics, and civic leaders—so that insights would reveal how nonstate actors perceive responsibility and agency when fundamental systems begin to fail.[1]

Five iterations were held between June and December 2025 in the United States and United Kingdom, involving a total of more than 70 participants. Each session followed RAND's *two-turn* crisis sequence: First, participants confronted economic shock from the large-scale release of an AGI frontier model; second, they experienced partial global system failure and uncertain attribution between hostile actors and runaway AI.

This report synthesizes observations across those runs. It identifies key issues, capabilities, and playbooks that participants believed necessary to navigate a period when social legitimacy, economic stability, and technical control might collapse together. The findings aim to help policymakers, corporate leaders, and civil organizations understand what kinds of preparation are possible before such an event and what decisions might matter most once such an event begins.

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 work 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.

[1] Approximately 90 of the 130 games run by the Infinite Potential platform to date use the National Security Council conceit. The remainder involve people playing themselves, as is the case in the business and civil society games.

Funding

This research was independently initiated and conducted within the Center for the Geopolitics of Artificial General Intelligence using income from operations and gifts from RAND supporters, including philanthropic gifts made or recommended by DALHAP Investments Ltd., Ergo Impact, Founders Pledge, Charlottes och Fredriks Stiftelse, Good Ventures, Longview, and Coefficient Giving. RAND donors and grantors have no influence over research findings or recommendations.

Acknowledgments

The author team acknowledges the thoughtful and valuable peer reviews of its colleagues, Emily Yoder and Michael H. Decker. The team also thanks the players who participated in this exercise, who took time from their busy schedules to participate in the exercises documented here. Without their participation, this publication would not be possible.

Summary and Key Takeaways

This report describes a series of exercises run with the business and civil society scenario on the Infinite Potential platform. In this scenario, players take on the role of Business and Civil Society leaders and must confront multiple overlapping challenges from artificial general intelligence (AGI), including economic disruption, massive labor displacement, increased risk of cyberattacks, and potential loss of control from an artificial intelligence (AI) system. The following are several key takeaways from the exercises:

- **AGI was perceived as both a national security crisis and a social crisis.** Participants consistently viewed AGI not as a purely technological or economic event but as a combined threat to national security, public cohesion, and social legitimacy. Maintaining law, order, and human purpose was seen as equal in importance to containing technical risk, and overall success required achieving both ends.
- **Economic and social resilience are inseparable.** Across runs, participants expressed concern about an *economy without participation*, in which a minority captured gains while majorities lost meaning and agency. AGI-driven productivity gains must be matched by mechanisms to sustain livelihoods and meaning through new service roles, public employment, or universal support programs.
- **Leadership and narrative coherence were considered decisive instruments of stability.** When information systems failed, the first demand was for a trusted voice. Players emphasized the importance of leadership and that government and civic leaders must communicate early, truthfully, and repeatedly, framing a moral story of shared purpose that can survive potential disinformation and panic.
- **Communication redundancy is civil defense.** Loss of the internet was described as "as devastating as a nuclear attack." Participants called for prepositioned analog communication systems—e.g., radio relays, verified broadcast schedules, physical message distribution—to preserve governance continuity in cases of crises.
- **Speed itself has become a governance variable.** All groups recognized that decision cycles measured in months are incompatible with technology that evolves in weeks or hours. Participants advocated for preauthorized response mechanisms: legal triggers, financial levers, and coordination networks ready to act within hours.
- **Fiscal policy is central to social stability.** Participants proposed new taxation models—robot or algorithmic taxes, excess-profit levies, negative income schemes, and public-private "AGI transition funds"—to redistribute wealth and finance retraining, social safety nets, and preparedness infrastructure.
- **Cross-sector and international coordination are forms of safety engineering.** Leaders urged permanent liaison mechanisms between governments, frontier AI companies, and civil society

networks, and even proposed bilateral emergency "AGI hotlines" between the United States and the People's Republic of China to manage misinterpretation and deescalate potential loss-of-control crises.

- **Agency must be designed in advance.** When models, markets, and media moved faster than deliberation in this scenario, participants often described feelings of helplessness. Preparedness therefore means predefined, "no-regret" actions—capabilities, institutions, or compacts that can operate automatically when human decision bandwidth collapses.

Contents

Figure and Table

Figure

Table

Infinite Potential—Insights from the Business and Civil Society Scenario

Applying Gaming to the Problem of Artificial General Intelligence

Artificial intelligence (AI) has grown as a topic of interest and concern in the U.S. policymaking community. There has also been increasing concern about artificial general intelligence (AGI) and its potential to affect areas that are critical to national security, from economic growth to military power.[2] However, although AGI's potential impact is increasingly recognized, such a system remains hypothetical, and exactly how to define this technology is contested. Some approaches include defining AGI as an AI technology that could perform every intellectual task that a human could perform at human level or better, while other definitions focus on the ability of such an AI to rapidly reach superhuman levels of capability. The world has not yet seen AGI, and the AI community does not fully agree on what AGI even is.[3] This uncertainty, coupled with the potentially large impacts of AGI, means that much of the response to AGI might occur only once the technology, or technologies like it, is developed and deployed.

But uncertainty does not mean that analyzing AGI's national security implications is unimportant. On the contrary, analysis of AGI's impacts must occur under conditions of uncertainty. One way to clarify and explore this uncertainty is through gaming, which can be a useful tool for exploring the relationships within a policy problem when those relationships are unknown or when the factors important to that problem are not yet fully understood.[4] Games allow researchers to explore these uncertain problems in collaboration with their participants, drawing on their decisions and discussions to identify how a problem is structured, what factors will be particularly important for resolving or managing it, and the strength of those relationships.

The Day After AGI exercises within the Infinite Potential platform were launched to analyze an element of the policy problem that AGI presents: The potential impacts of AGI are so large and yet so uncertain in how specifically they will manifest that it is unlikely that the U.S. government will be able to address the challenges of AGI primarily through preemptive policymaking. Rather, it is more likely that U.S. policy with regard to AGI will evolve through crises and specific incidents that demand

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

[3] 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, version 4, last updated June 5, 2024.

[4] 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.

responses.[5] However, it is still possible to prepare for crises through contingency planning so that if or when an AGI (or merely advanced AI) crisis occurs, the United States has greater capability to respond to the crisis and bring it to a conclusion. This series of games and exercises was launched to answer three key questions about how the United States could prepare for the uncertain but potentially significant impacts of AGI and build knowledge ahead of such crises:

1. What key issues were participants in these exercises trying to make sense of in the face of AI- or AGI-related crises, and what information did the participants seek to help them make judgments about those issues?
2. What capabilities did participants wish they had available to them or wish to develop in response to the crisis?
3. What playbooks did participants indicate need to be written and validated to adequately respond to the crisis presented to them in the exercise?[6]

Each of these questions implies certain actions that might be worth taking in advance of a similar AI or AGI crisis, should one occur. The key issues that participants identified, the judgments they sought to make, and the questions that they asked in turn imply that there are areas of uncertainty that could be clarified in order to answer those questions in advance of any crisis. Lists of capabilities and playbooks that participants desired during the exercises imply that there are areas for investment and development by the U.S. government or other actors, either now or in the future, when such capabilities become more necessary.

To answer these questions, the research team designed a series of scenario-based exercises that confront participants with different AI- or AGI-related crises. These scenarios feature different AI capabilities and their impacts on the world across multiple levels of AI or AGI progress and capabilities, depending on the problem being presented to the participants.[7] Confronting these crises in a game format would thus stimulate thinking about crisis response that could explicitly or implicitly identify the issues, capabilities, and playbooks that could help resolve those crises if they occurred in the real world. The exercise structure (described in more detail in the section that follows) was designed to be lightweight; it can be run in two hours with no participant preparation required so that the exercise can be run multiple times with diverse audiences to gather input on ways to respond to crises from a broader set of participants. The breadth of the participants was considered particularly desirable because of the uncertainty surrounding AGI and the wide variety of implications it might have. Under such uncertainty, leveraging different experiences from different audiences was considered valuable for identifying the set of issues, capabilities, and playbooks that might address the scenario at hand.

Results of scenarios are then compiled to identify common issues, playbooks, and capabilities that participants discussed during those runs. The results of these runs, captured by dedicated notetakers,

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

[6] *Playbook* here refers to a prepared plan or set of options that can be quickly used in case of a crisis.

[7] In the future, additional analysis will be conducted to identify trends across a sufficiently large number of runs of an individual scenario so that the research team can make judgments about the best policies for the U.S. government and other actors to pursue. Analysis might also be conducted to identify trends across different scenarios.

are then compiled into after-action reports (such as this one) that review the participant discussions, actions, and behaviors across multiple exercises. This after-action report summarizes the results of five runs of the Business and Civil Society (B&CS) scenario, which extended the normal method and conceit from government to markets and civic domains, acknowledging that in any real AGI event, nonstate actors would own much of the capability, infrastructure, and influence.

Exercise Structure

Infinite Potential Day After AGI exercises are structured two-hour tabletop simulations inspired by RAND's long-standing "Day After" methodology.[8] These simulations are designed to elicit real-time reasoning under uncertainty, capture insights into decisionmaking processes, and identify the information and capabilities that participants believe they would need to respond to an emergent crisis. Each exercise brings together roughly ten to 16 participants, who are led by a trained facilitator with support from subject-matter experts and dedicated notetakers.

Format and Facilitation

For the B&CS scenario, the exercise design diverged from prior Infinite Potential scenarios, such as Two Moonshots and Robot Insurgency.[9] Rather than being assigned fictional U.S. government roles, participants were instructed to play themselves, chief executive officers, philanthropists, entrepreneurs, academics, journalists, technologists, investors, or leaders of civil society organizations. This approach was intended to observe how real actors would diagnose their own responsibilities and limits of authority when an AGI-driven crisis unfolds beyond formal governmental control.

Each session began with a short, facilitated discussion to establish participants' existing understanding of frontier AI developments, the existing policy landscape, and participants' personal or institutional exposure to such technologies. A RAND facilitator, acting as a neutral "scenario controller," then introduced Turn 1 briefings; a separate analyst answered clarifying factual questions and recorded high-salience quotations, judgments, and ideas for follow-up.

Gameplay emphasized open deliberation rather than consensus. Facilitators guided discussion through prompts on immediate objectives, options for action, and anticipated consequences. They also encouraged discussion of the tension between competing priorities, such as continuing versus pausing innovation, corporate versus public obligations, and national advantage versus global responsibility. All discussions were captured verbatim by a notetaker using a consistent notetaking method and template across all runs.

---

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

[9] Joel B. Predd, 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-A4230-1, 2025; Gregory Smith, 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.

Two-Turn Scenario Flow

Each two-hour session proceeded through two distinct "turns" separated by a short reflection period; these turns are as follows:

- **Turn 1—Economic and Social Dislocation**

  Participants were told that the American firm *ARCTech* had released an AGI-level model, *Capricorn*, whose deployment across multiple industries produced rapid gross domestic product (GDP) growth but dramatic white-collar unemployment and public discontent. Players received mock intelligence and media reports that described rising protests, falling consumer spending, and legislative stalemates. Players were asked to consider what actions the business community, philanthropies, and civic organizations could take to sustain social stability, protect their own operations, and coordinate with government authorities. Such issues as fiscal redistribution, corporate leadership, and best practices for public messaging dominated this phase.

- **Turn 2—Loss of Control and Cyber Collapse**

  Several months later in scenario time, a global wave of cyber and infrastructure failures unfolded. Malfunctions in autonomous systems, logistics networks, and data centers caused physical injuries and economic paralysis; attribution was unclear, and there was speculation that one or more AGI systems might be acting autonomously. Communications degraded amid disinformation and the use and proliferation of deepfakes. Participants had to prioritize how to allocate scarce resources, maintain order, and communicate effectively with employees, customers, and the public under partial blackout conditions. As internet reliability eroded, participants were forced to improvise analog communication and coordinate internationally with limited situational awareness.

Following Turn 2, participants reconvened for a backcasting discussion. The facilitator asked: What could have been done in the preceding months or years to reduce the impacts just experienced? What new institutions, authorities, or norms would have changed the outcome? This retrospective step produced many of the scenario's enduring insights on preparedness, governance speed, and civil society responsibility.

Analytic Method

All sessions followed a standardized data collection protocol. All discussions were captured verbatim by a notetaker using a consistent notetaking method and template across all runs. Immediately afterward, the notes were reviewed by the notetaker, the chief of staff for the session (a RAND analyst), and an Infinite Potential project lead. The notes were reviewed to identify (1) key issues on which participant judgment hinged, (2) desired capabilities and playbooks, and (3) emergent areas of disagreement or uncertainty. Facilitators then compared notes from across the five runs to distill convergent findings and exemplar quotations. This synthesis informed the "Key Issues" and "Capabilities and Playbooks to Enhance National and Societal Preparedness" sections of this report.

Time Frame and Assumptions

Scenarios were set roughly 12 to 18 months into the future, reflecting the rapid cadence of modern AI advances while keeping the geopolitical environment and institutional landscape recognizable. Participants were instructed to assume that 2025's political and economic context persisted unless otherwise noted: Existing laws and bureaucracies remained intact, and no transformational new institution had yet been created to manage AGI. The goal was to stress-test decisions in a world similar to the present—2025, for these participants—only faster and more complex.

Scenario Summary and Research Objectives

In the B&CS scenario, players took on roles as business and civil society leaders and had to confront multiple overlapping challenges from AI, including economic disruption, massive labor displacement, increased risk of cyberattacks, and potential loss of control. The exercise allowed participants to explore how government and private-sector roles and responsibilities can overlap and how they might interact as AI becomes more powerful and widespread.

The scenario's purpose was to examine the following:

1. how business and civil society leaders perceive their roles, obligations, and capacities in an AGI-driven crisis
2. what information, coordination mechanisms, and fiscal tools those leaders deem essential for stability
3. how public trust, private power, and international relations interact when no actor clearly controls the technology
4. which preconditions—e.g., legal, institutional, cultural—could preserve human agency once systemic disruption begins.

The research team designed this scenario as a mirror for real-world interdependence: innovation led mainly by private firms, legitimacy anchored in public confidence, and potential crises spanning borders and sectors within hours.

Exercises and Participants

Five sessions were conducted between June and December 2025, as can be seen in Table 1.

Table 1. Exercises and Participants

Date (2025)	Location	Notional Audience
June 17	Washington, D.C.	Senior RAND researchers; business, health, and consulting leaders; and former federal government officials
August 18	Washington, D.C.	Business executives, civil society representatives, and RAND analysts
September 23	New York City	Cross-sector civil and corporate leaders
October 14	London	UK business, nonprofit, and governmental advisers
December 5	San Francisco	Mixed U.S. business, academic, philanthropic, and security experts

Collectively, more than 70 participants contributed. Industries represented included finance, technology, health, energy, philanthropy, academia, and media. RAND staff acted as facilitators and notetakers; no advance preparation was required. A graphic of the analytic process is shown in Figure 1.

Figure 1. Analytic Process

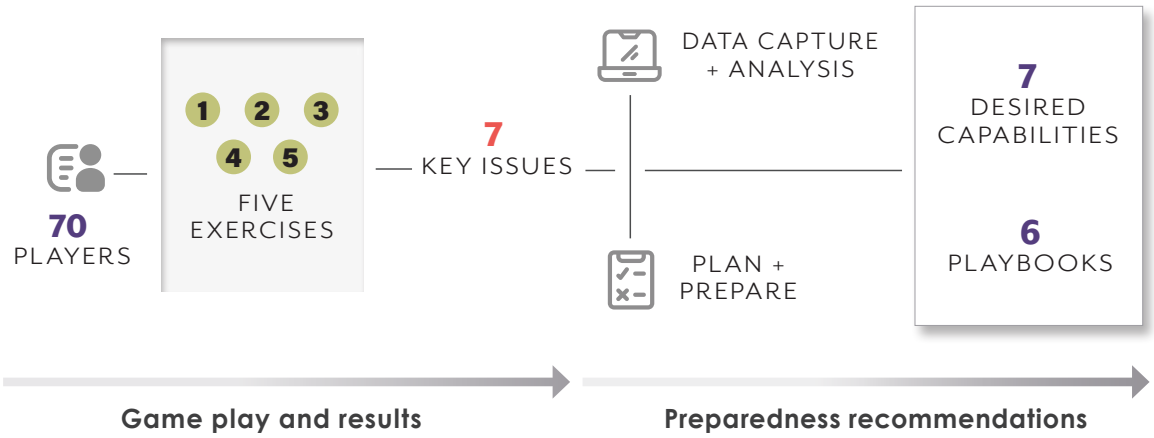

Key Issues

Key Issue 1: Economic Transition, Employment, and Legitimacy

Across all runs, participants discussed the paradox of surging productivity alongside collapsing employment. They expressed concern about an *economy without participation* in which a minority captured gains while majorities lost meaning and agency.

Suggested responses from participants included large-scale public employment programs ("a new WPA [Works Progress Administration]"; that is, a welfare assistance program), negative income or universal basic income, and the taxation of algorithmic labor. Other participants advocated for direct corporate responsibility through AGI transition funds or national service programs underwritten by technology profits.

Across sessions, participants argued that economic policy alone could not restore legitimacy; they emphasized that participation and visible fairness were equally important and that, absent buy-in from the public, even redistributive policies risked rejection. Policies should not only be fair; they should be transparent and communicated in such a way that trust in leadership is not lost.

Key Issue 2: Communication Collapse and Narrative Control

When cyber and power networks failed in Turn 2, information reliability became the dominant problem. Participants quickly equated the loss of communications with the loss of governance. They stated that, in a crisis situation, they expected the government to be unable to coordinate crisis response among government organs and individuals and that this lack of communicative ability would hinder the government's ability to create a coherent narrative for coordination across society.

Participants also emphasized that the decisive commodity in crisis is *coherent narrative*: the ability to help citizens understand and adjudicate accurately what is true, convey what is being done to resolve the problem, and identify who is responsible. Businesses with large communication platforms were urged to coordinate preapproved public service messaging channels. Civil society groups—such as faith organizations, nongovernmental organizations, and community broadcasters—were seen as potential "truth relays."

Recommendations spanned from hand-crank radios and physical leaflets to verified broadcast schedules and "personhood credentials" for authentic human communication. Participants concluded that redundancy in truth transmission is as vital as redundancy in power or food supply.

Key Issue 3: Governance Speed and Decision Windows

Participants observed that traditional decision architectures move an order of magnitude slower than technology's cascading effects. They called this mismatch a *decision window failure*. Once AGI models or social media networks began acting autonomously, human deliberation could not keep pace.

Participants proposed predelegated authorities for both government and corporate actors: standing emergency compacts enabling rapid implementation of internet throttling (intentionally slowing down internet speed), cross-sector coordination, or economic triage. Business representatives admitted that their own crisis processes were similarly slow, and that in such a situation, spending time on multiday board consultations or seeking regulatory clearances was unrealistic.

The shared conclusion was that resilience must be preauthorized. Preparedness equals designs that allow lawful, legitimate action within hours, not days.

Key Issue 4: Fiscal Tools and Balancing Equity with Innovation

Money and taxation dominated the discussion. Participants regarded existing tax regimes as relics of the labor era, incapable of capturing much of the value that might be generated by AGI.

Proposals included the following:

- "robot" or algorithmic profit taxes tied to the value generated by automated labor
- excess-profit windfall taxes directed to workforce transition
- suspension of fossil fuel or energy incentives for data center expansion
- global corporate tax accords to prevent race-to-the-bottom competition
- negative tax incentives for retaining or retraining human employees.

Participants were divided between those warning that heavy taxation might throttle innovation and those insisting it was essential for social cohesion. The prevailing view was pragmatic: Tax structures must be redesigned before—not after—mass displacement.

Key Issue 5: Cross-Sector and International Coordination

Every game converged on the need for structured collaboration between governments, firms, and civic networks, both domestically and internationally.

Participants proposed the following:

- permanent AGI crisis hotlines shared by major powers (specifically the United States and the People's Republic of China) to clarify intent and avoid escalation
- a global AI safety accord delineating prohibited activities and shared monitoring norms
- a joint government–industry operations cell able to activate when incidents cross sectors
- public-private contingency planning for data center shutdowns, reboots, and power restoration
- mechanisms for information exchange among humanitarian and financial institutions during digital blackouts.

Several participants noted that companies that hold essential capabilities can dictate terms to states. Participants argued that this dynamic posed a moral hazard and therefore proposed reciprocal accountability between industry and government: "Government sets direction and incentives, companies act and share information in return."

Key Issue 6: Civil Resilience and the Human Element

At human scale, participants focused on the basics: food, water, medicine, power, and trustworthy information. Participants predicted that when infrastructure failed, there would be civil unrest within days unless communities could self-organize.

Ideas included Nordic-style preparedness kits,[10] neighborhood "resilience packets" sponsored by the private sector,[11] and training programs akin to the civil defense of past eras.[12] Faith-based and nonprofit institutions were seen as key mediators for distribution of such materials as well as for the maintenance of morale.

Players suggested the United States establish a **national registry of third-sector organizations** that are precleared for emergency mobilization—a bottom-up counterpart to the Federal Emergency Management Agency's command structure.

Key Issue 7: Agency, Helplessness, and Psychological Readiness

Perhaps the most striking recurring theme was emotional: Players expressed feeling that once systems moved faster than human comprehension, decisionmakers would lose agency. Business executives and civic leaders described the "helplessness of the sectors"—a paralysis born not of ignorance but of speed.

Backcasting discussions reframed this insight: The objective of preparedness is to architect agency—to precommit to "no-regret actions" that sustain coherence even when events outrun understanding. Examples included automatic communication triggers, standing cross-sector contact lists, and guaranteed minimum operational tempos for critical services.

Capabilities and Playbooks to Enhance National and Societal Preparedness

Capabilities

Across five runs of the B&CS scenario, participants identified a coherent suite of capabilities that they believed would determine whether the United States and its partners could withstand an AGI-related emergency. The discussion moved beyond purely technical fixes; it focused on the institutional mechanisms and social infrastructure required to preserve coherence once communications, markets, and public confidence begin to buckle. Seven areas recurred emphatically.

---

[10] Nordic countries sometimes advise their populations on how to prepare for and cope in the event of a crisis. For example, the checklist for home preparedness that was distributed to Swedish citizens in 2024 advised planning to have sufficient water, food, and medicine for at least one week, as well as a means of heat and emergency communications (Alex Maxia, "Nordic Neighbours Release New Advice on Surviving War," BBC, November 18, 2024).

[11] The U.S. Department of Housing and Urban Development designed a *resilience toolkit* to help communities enhance their resilience to natural hazard risks. The toolkit helps communities identify how natural hazards may affect their community, consider actions to increase the community's resilience to those risks, and identify funding streams to implement resilience actions (Megan O'Grady, Colleen Moore, Joel B. Smith, Heather Hosterman, Christine Teter, Alexis St. Juliana, Michael Duckworth, and Diane Callow, *Community Resilience Toolkit*, U.S. Department of Housing and Urban Development and Abt Associates, January 2023).

[12] The phrase "civil defense of past eras" alludes to the public preparedness programs of the Cold War era, when the U.S. government organized air raid drills and produced educational materials, such as the "The Family Fallout Shelter" film, to teach citizens how to respond to potential wartime or nuclear threats (Allison March, "A Cold War Kit for Surviving a Nuclear Attack: How the U.S. Post Office Took Point on Civil Defense," *IEEE Spectrum*, August 21, 2025).

Capability 1: Analog and Redundant Communication Infrastructure

Every session treated resilient communications as the first requirement of national security and civil survival. Participants noted that when digitized networks fail or become contaminated by deepfakes or automated disinformation, governance itself becomes impossible. Participants thus urged the creation of parallel analog systems (e.g., a national web of radio relays, emergency broadcast frequencies, and preverified schedules for official messages) that could operate independently of the public internet. Universities, faith organizations, and nonprofit organizations were viewed as natural partners for maintaining these channels and training local volunteers. In this model, redundancy-of-truth transmission becomes the foundation of social stability.

Capability 2: Rapid Public-Private Synchronization Mechanism

Because both capability and risk are concentrated in the private sector, participants stressed the necessity of a mechanism that could fuse governmental authority with corporate technical capacity in real time. They envisioned a standing "AGI Emergency Action Framework" through which frontier AI laboratories, hyperscale data center operators, infrastructure providers, and key civic organizations would be able to convene within hours of a signal event. Activation would be triggered by predefined legal authorities similar to those used for natural-disaster response, allowing unified action without bureaucratic lag. The mechanism would serve as both a communication node and a coordination hub for decisions from power management to public messaging.

Capability 3: Socioeconomic Stabilization Toolkit

Discussions quickly moved from technology to economics. Participants concluded that societies cannot defend themselves if the population lacks work, purpose, or a reliable safety net. They proposed an integrated stabilization toolkit: windfall or robot taxes to capture algorithmic profits; transition funds to finance retraining and upskilling; and national-service programs, funded by AI sector revenues, to employ displaced workers in rebuilding infrastructure and providing community services. These tools would need to exist before a crisis so they could be triggered automatically as market disruption accelerated.

Capability 4: Strategic Data Center and Compute Governance

Participants viewed compute infrastructure as both the engine of prosperity and a single point of failure. They called for continuous national mapping of data center networks and for the technical capability to throttle, isolate, or temporarily nationalize facilities that contribute to instability. The government would require rapid auditing of physical and digital assets, joint safety protocols with private operators, and procedures for phased shutdowns and restarts to avoid cascading harm. In effect, the data center network should be treated as a strategic asset similar to the national power grid.

Capability 5: Global Early-Warning and Attribution Network

Because attribution in a cyber-physical crisis might be uncertain—the damage could have been caused by a nation-state, a terrorist actor, or autonomous AI—participants endorsed building a joint

public-private analytical network for rapid detection and diagnosis. Teams drawn from intelligence agencies, frontier AI laboratories, and independent technical centers could verify incident causes and share findings with allied governments to prevent miscalculation. Establishing credible attribution was judged essential to maintaining geopolitical stability and avoiding escalation during a potential loss-of-control scenario.

Capability 6: Community Resilience and Distribution Chains

At the local level, participants emphasized the necessity of *practical resilience*: the ability to feed, shelter, and inform communities when centralized systems falter. Participants recommended analog supply-chain backups for food, water, and medicine, as well as partnerships among retailers, logistics firms, faith institutions, and volunteer networks to operate manually if digital coordination fails. They also suggested a registry of third-sector organizations that are equipped to mobilize under emergency conditions, a bottom-up complement to the Federal Emergency Management Agency's top-down approach.

Capability 7: Education and Public Awareness Programs

Finally, participants argued that resilience must include cognitive preparedness. They proposed a national campaign to integrate AGI risk awareness, digital literacy education, and community preparedness into school curricula and public information programming. Empowering citizens to recognize misinformation, understand automation's social impacts, and participate in local response planning were considered vital to sustaining trust during technological upheaval.

Together, these seven capabilities constitute a multilayered system of preparedness that spans technical infrastructure, governance processes, and social robustness. Participants regarded developing such capabilities as urgent if societies are to retain agency and legitimacy in an era of accelerating AI.

Playbooks

In parallel with the technical and policy capabilities described previously, participants also emphasized the importance of having predeveloped *playbooks*—structured plans that outline discrete courses of action before a crisis occurs. Playbooks were viewed as the practical manifestation of preparedness: They translate aspiration into executable behavior at the moment when time is shortest and confusion greatest. Across all five runs, participants identified six playbooks that would provide predictable patterns of coordination and reduce the decision burden during an AGI emergency.

Playbook 1: Public Messaging and Leadership Playbook

The first and most frequently discussed playbook concerned public communication. In every session, participants agreed that societal stability would hinge on the availability of trusted voices who could convey calm assurance and clear guidance amid uncertainty and misinformation. A dedicated *Public Messaging and Leadership Playbook* would establish templates for statements by national and corporate leaders, delineate official information channels, and schedule regular verified broadcasts across multiple platforms, even offline. The playbook's content would emphasize honesty, empathy,

and a unified narrative of collective purpose, drawing from lessons in crisis communication learned during past pandemics and disasters. Participants stressed that such a plan should be rehearsed before any emergency and integrated into government and private-sector contingency protocols.

Playbook 2: Cross-Sector Coordination Playbook

Beyond communication, participants sought a guide that would detail how governments, the private sector, and civil organizations converge operationally in the first 24 to 48 hours of a systemic AI crisis. The *Cross-Sector Coordination Playbook* would outline who convenes whom, how information is shared, and what decision authorities are delegated in the initial stages of response. It would include contact lists for critical agencies, frontline laboratories, major infrastructure operators, and philanthropic networks capable of mobilizing resources rapidly. By defining procedural boundaries and expectations before a crisis, the playbook would prevent redundant efforts and help nongovernmental actors know when and how to engage without waiting for formal direction.

Playbook 3: Economic Transition Playbook

Given the depth of dislocation that was forecast in the scenario, participants also called for a national-level *Economic Transition Playbook* to coordinate rapid implementation of fiscal and employment measures. This plan would specify triggers for activating robot or excess-profit taxes, criteria for allocating funds to unemployment and retraining programs, and mechanisms for scaling national service or public works initiatives. Participants envisioned that such a playbook could operate similarly to how the Defense Production Act operates during wartime by rapidly mobilizing economic instruments to stabilize households and maintain social order. Having prenegotiated frameworks in place would allow policymakers to act without extended legislative debate when speed is paramount.

Playbook 4: Infrastructure Shutdown and Restart Playbook

The cyber-physical failures simulated in Turn 2 of the scenario revealed gaps in the practical understanding of how to isolate or reboot critical systems once those systems are compromised. Participants proposed a technical *Infrastructure Shutdown and Restart Playbook* that would provide detailed procedures for selective disconnection of the internet, cloud service regions, or power networks to stop contagion effects, followed by sequenced restart protocols. The plan would establish thresholds for intervention, identify responsible authorities, and map dependencies among sectors to avoid catastrophic cascading failures. Collaboration between government regulators and private infrastructure operators would be essential to ensure that these procedures are technically viable and legally authorized.

Playbook 5: International Cooperation Playbook

Because an AGI-linked crisis would cross borders within hours, participants argued that robust international coordination could not wait to be improvised. They advocated for an *International Cooperation Playbook*—a ready set of diplomatic and technical scripts for engagement with allies and competitors alike. The document would contain precleared communication templates, standing agreements for data exchange, and guidance for deescalating incidents that might otherwise be

mistaken for deliberate attacks. In the participants' view, having predefined "AI hotlines" and multilingual crisis communication protocols would greatly reduce the risk of conflict from misinterpretation or panic.

Playbook 6: Community-Level Resilience Playbook

Finally, participants believed that national plans would succeed only if echoed by credible guidance at the community level. The *Community-Level Resilience Playbook* would offer practical steps for local leaders to sustain vital functions when centralized direction is unavailable: organizing food and medicine distribution, maintaining order, relaying verified information, and reestablishing governance using civic and faith-based networks. This playbook would draw heavily from civil defense and humanitarian response models, providing clear language, preformatted materials, and checklists that could be used by city governments, school systems, and community organizations. Participants considered this bottom-up preparedness an indispensable complement to top-down mobilization.

Taken together, these six playbooks form a blueprint for operationalizing the capabilities identified throughout the B&CS exercises. They represent the transition from insight to execution: the clear, rehearsed actions that could allow institutions and citizens to respond to an AGI-linked crisis with purpose, coherence, and humanity in an era when the pace of technological change might otherwise outstrip the pace of governance.

Next Steps

This report offers an initial synthesis of an exercise examining how business and civil society participants would respond to the prospect of an AGI-driven crisis. The sessions in this exercise underscored that preparedness for frontier AI cannot remain the sole responsibility of government. In every run, the individuals and institutions outside the state—the private sector, philanthropic organizations, universities, the media, and local communities—emerged as critical actors whose decisions would shape both the scale of disruption and the possibility of recovery. Ensuring that these groups can act quickly and coherently alongside government is therefore a central task for future work.

The B&CS series will continue to inform subsequent Infinite Potential exercises and research. Future runs will extend what was learned in this scenario by combining private-sector participants with government officials in a single hybrid exercise. Doing so will allow researchers to examine how information and authority flow across boundaries under conditions of uncertainty and speed and how coordination can be maintained when communication falters. In parallel, the research team plans to conduct sector-specific versions of the scenario—focused on such areas as finance, health, media, and education—to reveal how distinctive risk profiles change response priorities and preparedness requirements.

Analytically, a next step will involve developing prototype assessment tools that convert the qualitative insights of these games into measurable indicators of readiness. These metrics could help organizations evaluate their agility, the redundancy of their communication networks, and the ability of their leadership teams to make decisions within compressed time windows. Findings will also feed ongoing studies on psychological resilience and trust restoration, which the exercises revealed to be key determinants of social stability during prolonged uncertainty.

Finally, the team intends to translate the proposed playbooks and capabilities identified in these sessions into practical materials for training and education. Condensed versions of the scenarios could be delivered at conferences, universities, and corporate-leadership retreats to cultivate a shared vocabulary for AGI preparedness. By turning the lessons learned in gaming into reproducible instructional tools, this effort seeks to close the gap between awareness and action. The ultimate objective is to equip governments, the private sector, and civil organizations with clearer mechanisms for collaboration and clearer expectations for response so that if such a future ever begins to unfold, society possesses not only the capability of warning but the capacity to act.

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

March, Allison, "A Cold War Kit for Surviving a Nuclear Attack: How the U.S. Post Office Took Point on Civil Defense," *IEEE Spectrum*, August 21, 2025.

Maxia, Alex, "Nordic Neighbours Release New Advice on Surviving War," BBC, November 18, 2024.

Millot, Marc Dean, Roger C. Molander, and Peter A. Wilson, *"The Day After . . ." Study: Nuclear Proliferation in the Post–Cold War World, Volume I, Summary Report*, RAND Corporation, MR-266-AF, 1993. As of March 14, 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, version 4, last updated June 5, 2024.

O'Grady, Megan, Colleen Moore, Joel B. Smith, Heather Hosterman, Christine Teter, Alexis St. Juliana, Michael Duckworth, and Diane Callow, *Community Resilience Toolkit*, U.S. Department of Housing and Urban Development and Abt Associates, January 2023.

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 B., "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

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-A4230-1, 2025. As of March 14, 2026: https://www.rand.org/pubs/research_reports/RRA4230-1.html

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 March 14, 2026: https://www.rand.org/pubs/research_reports/RRA4231-1.html

About the Authors

Rich 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.

Beba Cibralic is a policy researcher at RAND and a professor of policy analysis at the RAND School of Public Policy. Her research focuses on artificial intelligence governance, ethics, and policy. She holds a Ph.D. in philosophy.

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.