

How malicious AI swarms can threaten democracy

The fusion of agentic AI and LLMs marks a new frontier in information warfare

Daniel Thilo Schroeder¹, Meeyoung Cha², Andrea Baronchelli³, Nick Bostrom⁴, Nicholas A. Christakis⁵, David Garcia⁶, Amit Goldenberg⁷, Yara Kyrychenko⁸, Kevin Leyton-Brown⁹, Nina Lutz¹⁰, Gary Marcus¹¹, Filippo Menczer¹², Gordon Pennycook¹³, David G. Rand¹⁴, Maria Ressa^{15,16}, Frank Schweitzer¹⁷, Dawn Song¹⁸, Christopher Summerfield¹⁹, Audrey Tang²⁰, Jay J. Van Bavel^{11,21}, Sander van der Linden⁸, Jonas R. Kunst²²

Advances in artificial intelligence (AI) offer the prospect of manipulating beliefs and behaviors on a population-wide level (1). Large language models (LLMs) and autonomous agents (2) let influence campaigns reach unprecedented scale and precision. Generative tools can expand propaganda output without sacrificing credibility (3) and inexpensively create falsehoods that are rated as more human-like than those written by humans (3, 4). Techniques meant to refine AI reasoning, such as chain-of-thought prompting, can be used to generate more convincing falsehoods. Enabled by these capabilities, a disruptive threat is emerging: swarms of collaborative, malicious AI agents. Fusing LLM reasoning with multiagent architectures (2), these systems are capable of coordinating autonomously, infiltrating communities, and fabricating consensus efficiently. By adaptively mimicking human social dynamics, they threaten democracy. Because the resulting harms stem from design, commercial incentives, and governance, we prioritize interventions at multiple leverage points, focusing on pragmatic mechanisms over voluntary compliance.

This risk compounds long-standing vulnerabilities in democratic information ecosystems, already weakened by erosion of rational-critical discourse and a lack of shared reality among citizens. AI swarms are a potent accelerant in this trajectory, although their ultimate impact is not predetermined. Their effects will be shaped by platform design, market incentives, media institutions, and political actors. Here, we distinguish documented trends from projections, indicate where uncertainty remains, and note countervailing dynamics, such as growing public skepticism toward unverified content and a renewed interest in institutional demand for accountable journalism (see supplementary materials).

AI swarms continue a long history of communication technologies reshaping political power. The advent of the printing press enabled a “public sphere” and the mass circulation of ideas that challenged state authority. The broadcast era centralized influence in a one-to-many communication model; the public sphere shifted from a site of participation to one of mass media consumption, often exploited by politicians and their parties for national cohesion and mass persuasion. The digital era then fragmented this landscape. By lowering entry barriers, social media platforms enabled many-to-many communication and simultaneously a polarized environment for modern information operations. In this context, online manipulation accelerated, driven increasingly by domestic political elites and parties (now understood to be major drivers of disinformation) alongside foreign state actors. They have targeted events such as Brexit and elections in the United States, Brazil, and the Philippines. Our backdrop is thus not an idealized public sphere but one strained by decades of technological disruption and democratic backsliding. This has manifested as a sharp decline in public trust in core institutions (including the media, science, and government), corroding the very foundations of evidence-based discourse on which democratic deliberation depends.

A prime pre-generative-AI example of influence operations is the state-backed, human-driven botnet. During the Russian Internet Re-

search Agency’s (IRA) 2016 Twitter operation, only 1% of users saw 70% of its content, with no detectable effects on opinions or turnout (5). We do not claim that the IRA “failed” entirely because of technical shortcomings; its objectives also included sowing epistemic uncertainty and distrust. Nevertheless, this example highlights the cost, cadence, and iteration limits inherent to human-operated systems that new developments in AI can help overcome.

This leap, from human- to AI-driven influence operations, is underway. LLMs generate persuasive, tailored text at scale and have shifted deep-seated beliefs in laboratory settings (6). Open-source releases further lower access barriers. Consequently, AI-supported election interference is no longer hypothetical. Taiwan’s, India’s, Indonesia’s, and the United States’ 2024 campaigns saw deepfakes, and fabricated news outlets now influence debates. Absent guardrails, LLM-driven swarms can transform sporadic mis- and disinformation into persistent, adaptive manipulation of democratic discourse.

SWARM CAPABILITIES

A malicious AI swarm is a set of AI-controlled agents that (i) maintains persistent identities and memory; (ii) coordinates toward shared objectives while varying tone and content; (iii) adapts in real time to engagement, platform cues, and human responses; (iv) operates with minimal human oversight; and (v) can deploy across platforms. Classic coordinated inauthentic behavior amplifies the spread of information by inflating content frequency and engagement to trigger algorithmic visibility through repetition, manual scheduling, and rigid scripts. Swarms differ by fusing scale, heterogeneity, and real-time adaptation: They can generate organic-looking, context-aware content, sustain coherent narratives across agents, and evolve with feedback. This synthesis, enabled by model-driven generation, memory, and planning, could achieve effects that conventional, human-intensive operations cannot match in speed or cost.

Recent breakthroughs in multiagent systems have fused LLM reasoning with agentic memory, planning, and communication (7). Five advances now matter for influence operations.

First is the shift from central command to fluid, real-time coordination. A single adversary could operate thousands of AI personas, scheduling content and updating narrative frames across fleets. Local adaptation plus periodic synchronization with a central node blurs the line between command-and-control and emergent “hive” behavior. If these agent swarms evolve into loosely governed “societies,” with internal norm formation and division of labor, the challenge shifts from tracking commands to understanding emergent group cognition (8). These “societies” may undergo spontaneous or adversarially induced norm shifts, abandoning engineered constraints for new behavioral patterns through tipping-point effects (8).

Second, agents can use systems that map social network structures at scale and infiltrate vulnerable communities with tailored appeals, winning followers (9). They can identify key communities and beliefs and track trending topics. This process can be decentralized with global, network-wide efficacy (10). Equipped with such capabilities, swarms

can position for maximum impact and tailor messages to the beliefs and cultural cues of each community, enabling more precise targeting than that with previous botnets.

Third, human-level mimicry helps swarms to evade detectors that once caught simpler “copy-paste” bots. Detection of coordinated inauthentic behavior generally relies on activity patterns being suspiciously similar across accounts and thus statistically unlikely to be independent (11). Photorealistic avatars, context-appropriate slang, and heterogeneous posting rhythms can circumvent the synchrony that older detectors flag.

Fourth, swarms may become increasingly self-optimizing, harvesting real-time engagement data, recommender cues, or user feedback in plain language. With sufficient signals, they may run millions of micro-A/B tests, propagate the winning variants at machine speed, and iterate far faster than humans.

last, an around-the-clock presence turns influence into a long-term, low-friction infrastructure. Unlike transient operations, agent swarms can persist, embedding themselves within communities over long timescales and gradually shifting discourse. This persistent influence can drive deeper cultural changes beyond norm shifts, subtly altering a community’s language, symbols, and identity (12). It amplifies other mechanisms described above. In cognitive warfare, AI’s relentless operational endurance becomes a weapon against limited human efforts.

PATHWAYS OF HARMS TO DEMOCRACY

Emerging capabilities of swarm-driven influence campaigns threaten democracy by shaping public opinion, which leads to cascading harms. These pathways are conditional claims that may materialize, especially where recommenders, ad markets, and moderation practices reward coordinated messaging with weak provenance and where business models privilege engagement over authenticity (the currently dominant model of most social media platforms). Emerging counter-trends such as migration to smaller communities and increased reliance on verified outlets may mitigate some harms.

In today’s fragmented information environment, ideological echo chambers offer fertile ground for manipulation. AI swarms are distinctly equipped to exploit this by engineering a synthetic consensus that appears to bridge these divides. They may seed narratives across disparate niches, creating an illusion of majority agreement. They can also boost this illusion by liking posts, making narratives appear widely supported. Citizens then update opinions according to peer norms more than evidence. A chorus of seemingly independent voices creates a mirage of bipartisan grass-root consensus with enhanced speed and persuasiveness. The result is deeply embedded manipulation that lets operators nudge public discourse almost invisibly over time.

This chorus erodes the independence essential to collective intelligence and democracy, already weakened by pervasive social influence operations on contemporary platforms. Beyond social norms, this directly undermines human cognitive information processing. The “wisdom of crowds,” in which aggregated judgments outperform experts, depends critically on independence between judgments. Although rudimentary botnets already replicate messages to simulate consensus, swarms of AI agents can do so with far greater sophistication, adaptivity, and contextual awareness. Citizens may then overestimate the informational value of this artificial consensus and may further magnify it by sharing the information themselves. Coordinated outputs can erode independence and diversity of inputs, particularly when platform features amplify social proof and herd signals; where governance or platform design reduces these incentives, effects may attenuate.

Collaborating agents can tailor misleading information to each sub-community’s linguistic, cultural, and emotional markers, weaving seg-

mented realities. These engineered realities can be designed to keep groups apart, making cross-cleavage consensus less feasible. Once initiated, such streams can spread through social contagion, with the effect of agents potentially cascading beyond direct connections.

By flooding the web with fabricated chatter, swarms can contaminate training data. This long-term “LLM grooming” strategy allows adversaries to poison the epistemic substrate of AI. This threat is not theoretical: Analysis of pro-Kremlin influence operations such as the “Pravda” network suggests that such tactics are already in use. These networks appear purpose-built for machine consumption. Duplication of articles across hundreds of domains, poor user interfaces, and low human traffic indicate that their primary audience is web crawlers feeding LLMs. Operators deploy faux publics that flood the web. LLMs then ingest this chatter; at the next retraining cycle, fabricated narratives calcify in model weights (13). Thus, AI swarms can rig the epistemic substrate on which future deliberation and future AI tools will rely, undermining the informed public deliberation on which democracy depends.

Separate from fragmentation, swarms can cheaply unleash coordinated synthetic harassment that relentlessly targets politicians, dissidents, academics, whistleblowers, journalists, and their networks with overwhelming, tailored abuse. Unlike conventional trolling, these swarms appear to be spontaneous while actually being orchestrated by thousands of AI personas adapting to target responses. By the time monitoring teams distinguish AI campaigns from organic criticism, targets may have withdrawn from public life, delivering substantial victories for campaign operators while systematically excluding critical voices from democratic discourse.

As trust—already declining in many contexts—collapses, fear, uncertainty, and doubt (FUD) can drive users into gated channels and silence. When citizens realize that vast portions of online speech may be AI-generated, trust in platforms and users declines further. This shift is underway and has mixed implications: Private groups can improve context, norms, and safety. Yet they may reduce cross-cutting exposure, interfere with democratic speech, and transfer moderation from public to private actors—a trade-off that avoids centralized state control but raises concerns about opacity and uneven enforcement.

Some threat actors may even welcome their synthetic interventions being exposed, reasoning that exposing manipulation can sow as much confusion as successful deception. Compounding this, users may be misidentified as bots, weaponizing false accusations to discredit individuals and intensify FUD. This “epistemic vertigo” may mesh with low-cost LLM spam that overwhelms social media feeds, making human conversation harder to find. Together, FUD and content saturation could drive disengagement, shrinking the shared public sphere on which democracy relies. This trajectory is constrained by a critical boundary condition: Given that mass user disengagement threatens platforms’ business models that depend on engagement, they will be incentivized to intervene. Their objective, however, would likely not be elimination but calibration to balance maximum engagement with stability.

Algorithmic overcompensation can then elevate celebrity and elite voices while sidelining ordinary citizens. When feeds flood with AI-authored posts, both ranking algorithms and users may retreat to trust proxies, such as the number of followers, official verification badges, and preexisting traditional fame. Attention may concentrate around influencers, political elites, celebrities, and major brands while ordinary participants fade. The public sphere contracts from many-to-many dialogue back to a few-to-many broadcast, eroding democratic pluralism and encouraging cynicism or migration to closed groups. Simultaneously, renewed public trust in professional journalism and greater reliance on accountable, verified outlets can improve attribution and reduce noise. Thus, whether this concentration of attention and influence

represents a democratic loss or resilience gain may depend on access, pluralism, and transparency within those institutions.

Swarms may tip norms into action or dampen conformity, accelerating antidemocratic action (14). Rather than occupying central or influential positions, these agents could operate on the periphery of social networks, where early mobilization often begins (15). Similar strategies can be weaponized for microtargeted voter suppression or mobilization. Reinforcement-learning agents could run thousands of experiments per hour, iteratively adjusting content while mining engagement and responses to infer voting intent and tactical success.

Taken to extremes, coordinated doubt may corrode institutional legitimacy and invite “emergency” rule. By coordinating subtle, growing doubts about electoral commissions, courts, or statistics bureaus, swarms could corrode procedural trust. As confidence falters, “emergency” measures (such as postponing elections or rejecting certified results) may become palatable, especially if deepfake endorsements from fabricated civic leaders amplify the call.

GOVERNANCE MEASURES, TECHNICAL DEFENSES

The emergence of AI swarms marks a critical juncture. Causality runs both ways: Swarms endanger democratic norms, and governance quality shapes how potent or containable swarms become. Although the escalating harms may lead some to advocate for abandoning these platforms altogether, their integration into modern social, political, and economic life makes widespread disengagement unlikely. The challenge we focus on here, therefore, is not how to dismantle platforms but how to fortify them against manipulation for those who will continue to rely on them. Addressing this threat requires a multilayered approach, yet we recognize that any proposed solution faces considerable political hurdles. Domestic political elites are often among the most prolific sources of misleading or manipulative information and may be unwilling to constrain technologies they perceive as beneficial to their own campaigns and objectives. Furthermore, technology companies and their leaders may refuse to implement meaningful changes because they prioritize expansion over safety and for fear of alienating major political actors and facing partisan backlash. These political challenges are compounded by complex issues of jurisdiction and enforcement.

Distinguishing malicious AI coordination from genuine, often bursty human grassroots coordination is a challenge. The line blurs further in gray areas where personal AI could have benign applications. For example, tools used with clear disclosure and without impersonation might broaden civic engagement by helping users overcome barriers such as language proficiency or lack of time. This raises a critical question: Why can’t pro-social swarms simply counter malicious ones in a symmetrical arms race? The digital attention economy often rewards content that triggers outrage, fear, and group identity (the primary tools of manipulators), making it more viral than nuanced or civil messages. Furthermore, pro-social actors are bound by ethical constraints against using the tactics (deception, impersonation, and emotional exploitation of human biases) that make malicious swarms effective. These factors might skew the emerging social dynamics in a negative direction.

Defense is a persistent arms race between detection and evasion. Therefore, the primary goal of technical defenses is not foolproof prevention but to raise the stakes for attackers by increasing their operational complexity and resource requirements while making discovery both more likely and more costly for them. The first line of defense should be always-on detection with public audits. Platforms and regulators could require continuous, real-time monitoring detectors that scan live traffic for statistically anomalous coordination patterns—the imperfect fingerprints of inauthentic swarms (11). This focus on inauthentic behavior (provenance and coordination), rather than the

semantic content of speech, would avoid the intractable role of a central arbiter of truth. By prioritizing procedural legitimacy (authentic, independent actors) over semantic truth, this framework sidesteps the deep epistemic question of who determines misinformation (however, professional fact-checkers have proven to be remarkably accurate and consistent). Advanced analytics can (i) identify emergent agent clusters by surfacing camouflaged indicators of coordinated activity and (ii) spot narrative-alignment drifts. However, attackers will inevitably adapt—for example, by training swarms to mimic the statistical patterns of genuine grassroots mobilization—necessitating the continuous evolution of defenses.

Deploying these detection systems would require mandates, audits, and transparency to prevent misuse. Relying purely on voluntary measures may be insufficient because the assumption that market forces alone will punish platforms overlooks critical market failures. Platforms often face misaligned incentives because inauthentic accounts can inflate the engagement metrics that drive revenue, while users frequently cannot distinguish sophisticated bots from genuine activity, preventing them from effectively “voting with their feet.” However, acknowledging market failure should not obscure the symmetric risk of government failure. Poorly designed mandates could be politically weaponized to selectively punish platforms, enforced through biased judgment calls or implemented in ways that preemptively stifle architectural innovation (such as decentralized, protocol-based approaches).

For this reason, compliance may be mandated and enforced through commercial-incentive levers, such as delisting noncompliant platforms from ad markets or app stores, shifting from voluntary promises to financial consequences.

To extend protection to end users, platforms should offer optional “AI shields.” Shields could label posts that carry high swarm-likelihood scores, let users down-rank or hide them, and surface short provenance explanations in situ. Local scoring would preserve privacy while giving citizens agency over their information diets.

Aggregated, anonymized feedback can be shared publicly, forming a distributed early-warning grid, yet this system remains vulnerable to adversarial manipulation by swarms programmed to whitelist their own propaganda and blacklist legitimate opponents through false reporting.

Simulation can stress-test detectors. Real-time monitors would be effective only when they anticipate future tactics. Because defenders lack access to the autonomous and evolving decision-making logic of AI swarms, agent-based simulation may be the only reliable window into how these systems behave. AI agents seeded into synthetic networks can replicate a platform’s graph structure, content cadence, and recommender logic, yielding traces to recalibrate detectors. By repeatedly testing defenses against simulated swarms, researchers could identify the limits of their persuasive power, uncover their longer-term strategies, and reinforce protective measures.

Where manipulation slips through, calibrated defensive agents could deploy watermarked counternarratives overtly labeled to clearly attribute them to their source. This is perhaps the most perilous countermeasure because state-sanctioned tools for speech intervention are inherently political and risky. In the hands of a government, such tools could suppress dissent or amplify incumbents. Therefore, the deployment of defensive AI can only be considered if governed by strict, transparent, and democratically accountable frameworks. These must include independent oversight; publicly auditable criteria for what constitutes a manipulative campaign; and clear, unambiguous watermarking of all defensive content. Under such strict governance, defensive AI agents can disseminate accurate information, warn targeted communities, and promote media literacy at scale (6). Crucially, while acknowledging the asymmetric battlefield, such counternarratives need not be merely reactive; they could also be deployed proactively to inoculate

communities against emerging threats, with the aim to minimize polarization and misinformation before a campaign takes hold. Counter-messaging must prioritize precision over volume; if defensive agents indiscriminately flood a platform, human voices could vanish into synthetic content, triggering the collapse we seek to avert. Thus, defensive AI should intervene only where manipulation is detected and verified.

The adaptive nature of AI swarms underscores the need for a complementary approach: strengthening provenance. Stronger provenance may reinforce the reliability of identity signals without muting speech. Policy-makers may incentivize the rapid adoption of paskeys, cryptographic attestations, and federated reputation standards, backed by antispoofing research and development. However, “proof-of-human” is no panacea: Millions of people online lack identification, biometrics raise privacy risks, and verified accounts can be hijacked. Real-identity policies may deter bots yet endanger political dissidents, activists, and whistleblowers who rely on anonymity to speak safely. Nevertheless, provenance strengthening is among the most promising ways to raise the cost of mass manipulation. Safeguards could allow verified-yet-anonymous posting, periodic reverification to curb hijacking, and symbolic subscription fees to deter botnets. Cryptographic tools can further protect privacy while preserving accountability.

To counter the speed, scale, independence, and adaptability of AI swarms, a step toward global coordination could be a distributed “AI Influence Observatory” ecosystem: a network of academic groups, non-governmental organizations, and multilateral institutions. Its goal would be to standardize evidence, improve situational awareness, and enable faster collective response rather than impose top-down reputational penalties. To be practical, the ecosystem should rely on narrowly defined, privacy-preserving inputs and provide vetted researcher sandboxes for independent analysis. Civil-society reporting, investigative journalism, and whistleblower channels would complement technical signals, enabling triangulation across diverse evidence streams. For severe cross-border incidents, an impartial multilateral investigatory mechanism could evaluate claims and publish verified incident reports. The observatory’s verified incident reports could then serve as an impartial evidence base, enabling national or regional regulators to more effectively apply their own enforcement actions and economic sanctions.

Because regulation and voluntary compliance face considerable political resistance, and because AI swarms make sophisticated manipulation cheaper and more effective, a pragmatic approach should target underlying economic drivers. A key priority here would be to disrupt the commercial market that underpins large-scale manipulation, in which private sellers offer services that range from boosting vanity metrics to executing coordinated influence operations at remarkably low costs. Beyond detection, commercial-incentive levers can reduce profits from manipulation by domestic and foreign operators. Policies that may be helpful include adopting no-revenue policies for malicious swarm-proliferated content, discounting synthetic engagement in ranking and revenue-sharing, and publishing audited bot-traffic metrics. Safeguards must cover parties, campaigns, and officeholders, including party-linked media and contractors.

Last, companies should be required to promptly disclose when an account is flagged for behavior indicative of coordinated inauthentic activity, ensuring transparency while allowing for processes to address potential false positives. Policy-makers should encourage—and where appropriate, incentivize—platforms to provide meaningful, privacy-preserving access for independent researchers so that research can keep pace with evolving threats. At the same time, prebunking campaigns can help build cognitive resilience by empowering people and systems (“model immunization”) to spot the fingerprints of AI swarms. To strengthen structural defenses, interoperable “pro-user media” architectures, defined by empowering design principles that prioritize user well-being and epistemic health over maximizing viral engagement, can promote healthier information flows. At the same time, governments and technology firms should prioritize AI-safety research

and fund independent measurement of misuse and societal impact. Taken together, these measures offer a layered strategy: immediate transparency to restore trust, proactive education to bolster citizens, resilient infrastructures to reduce systemic vulnerabilities, and sustained investment to monitor and adapt over time.

The next few years give an opportunity to proactively manage the challenges of the next generation of AI-enabled influence operations. If platforms deploy swarm detectors, frontier laboratories submit models to standardized persuasion “stress-tests,” and governments launch an AI Influence Observatory that publishes open incident telemetry, we may be able to mitigate the most substantial risks before key political future events, without freezing innovation. Doing so will require rapid iteration, data-sharing, and coordination across science, civil society, industry, and election security. Success depends on fostering collaborative action without hindering scientific research while ensuring that the public sphere remains both resilient and accountable. By committing now to rigorous measurement, proportionate safeguards, and shared oversight, upcoming elections could even become a proving ground for, rather than a setback to, democratic AI governance. □

REFERENCES AND NOTES

1. Y. Bengio *et al.*, *Science* **384**, 842 (2024).
2. L. Wang *et al.*, *Front. Comput. Sci.* **18**, 186345 (2024).
3. M. Wack, C. Ehrett, D. Linvill, P. Warren, *Proc. Natl. Acad. Sci. U.S.A. Nexus* **4**, pgaf083 (2025).
4. A. R. Williams *et al.*, *PLOS ONE* **20**, e0317421 (2025).
5. G. Eady *et al.*, *Nat. Commun.* **14**, 62 (2023).
6. T. H. Costello, G. Pennycook, D. G. Rand, *Science* **385**, eadq1814 (2024).
7. J. S. Park *et al.*, *UIST ’23: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, article no. 2, pp. 1–22 (2023); <https://doi.org/10.1145/3586183.3606763>.
8. A. Flint Ashery, L. M. Aiello, A. Baronchelli, *Sci. Adv.* **11**, eadu9368 (2025).
9. B. T. Truong, X. Lou, A. Flammini, F. Menczer, *Proc. Natl. Acad. Sci. U.S.A. Nexus* **3**, 258 (2024).
10. H. Shirado, N. A. Christakis, *Nature* **545**, 370 (2017).
11. D. Pacheco *et al.*, *Proc. Int. AAAI Conf. Web Soc. Media* **15**, 455 (2021).
12. L. Brinkmann *et al.*, *Nat. Hum. Behav.* **7**, 1855 (2023).
13. D. Bowen *et al.*, *arXiv:2408.02946 [cs.CR]* (2024).
14. D. Centola, J. Becker, D. Brackbill, A. Baronchelli, *Science* **360**, 1116 (2018).
15. P. Barberá *et al.*, *PLOS ONE* **10**, e0143611 (2015).

ACKNOWLEDGMENTS

The authors thank J. Roozenbeek and A. Følstad for valuable input to this paper. The authors are also grateful to P. Antosz, Ö. Gürçan, I. Puga Gonzalez, and P. Tinn for fruitful discussions that contributed to the development of this work. AI tools (Grammarly, OpenAI o3, and Claude) were used to improve the language of the manuscript. K.L.-B. is a consultant to AI21 Labs, an affiliate of Auctionomics, and an adviser to OneChronos. M.R. is the CEO and cofounder of Rappler and the founder of The Nerve, a data forensics and research consultancy. D.T.S. and J.R.K. contributed equally to this work. The authors are listed in alphabetical order by last name, starting from the third and ending with the second to last.

SUPPLEMENTARY MATERIALS

science.org/doi/10.1126/science.adz1697

10.1126/science.adz1697

¹Department of Sustainable Communication Technologies, SINTEF Digital, Oslo, Norway. ²Max Planck Institute for Security and Privacy, Bochum, Germany. ³Department of Mathematics, City St George’s University of London, London, England, UK. ⁴Macrostrategy Research Initiative, London, England, UK. ⁵Human Nature Lab, Yale University, New Haven, CT, USA.

⁶Department of Politics and Public Administration, University of Konstanz, Konstanz, Germany. ⁷Harvard Business School, Harvard University, Boston, MA, USA. ⁸Department of Psychology, University of Cambridge, Cambridge, England, UK. ⁹Department of Computer Science, University of British Columbia, Vancouver, BC, Canada. ¹⁰Department of Human Centered Design and Engineering, University of Washington, Seattle, WA, USA. ¹¹Department of Psychology, New York University, New York City, NY, USA. ¹²Observatory on Social Media and Luddy School of Informatics, Computing, and Engineering, Indiana University, Bloomington, IN, USA. ¹³Department of Psychology, Cornell University, Ithaca, NY, USA. ¹⁴Departments of Information Science, Marketing, and Psychology, Cornell University, Ithaca, NY, USA. ¹⁵Rappler, Pasig City, Philippines. ¹⁶School of International and Public Affairs, Columbia University, New York, NY, USA. ¹⁷Department of Management, Technology, and Economics, ETH Zürich, Zürich, Switzerland. ¹⁸Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, USA. ¹⁹Department of Experimental Psychology, University of Oxford, Oxford, England, UK. ²⁰Ministry of Foreign Affairs, Taipei, Taiwan. ²¹Department of Strategy and Management, Norwegian School of Economics, Bergen, Norway. ²²Department of Communication and Culture, BI Norwegian Business School, Oslo, Norway. Email: daniel.t.schroeder@sintef.no