

AI RIGHTS FOR HUMAN SAFETY

Peter N. Salib & Simon Goldstein***

Artificial Intelligence (“AI”) companies are racing to create Artificial General Intelligence, or “AGI.” If they succeed, the result will be human-level AI systems that can independently pursue high-level goals by formulating and executing long-term plans in the real world. By default, such systems will be “misaligned”—pursuing goals that humans do not desire. This mismatch of goals will put humans and AGIs into strategic competition with one another. Thus, leading AI researchers agree that, as with competition between humans with conflicting goals, human-AI strategic conflict could lead to catastrophic violence.

Existing law is not merely unequipped to mitigate this risk; it will actively make things worse. This Article is the first to systematically investigate how law affects the risk of catastrophic human-AI conflict. It begins by arguing, using formal game-theoretic models, that under today’s legal regime, humans and AIs will likely be trapped in a prisoner’s dilemma. Both parties’ dominant strategy will be to permanently disempower or destroy the other, even though the costs of such conflict would be high.

The Article contends that one surprising legal change could help to reduce catastrophic risk: AI rights. Not just any rights will do. To promote human safety, AIs should be given the basic private law rights already enjoyed by other non-human agents, like corporations. AIs should be empowered to make contracts, hold property, and bring

* Peter N. Salib is an Assistant Professor of Law at the University of Houston Law Center, Executive Co-Director of the Center for Law and AI Risk, Law and Policy Advisor to the Center for AI Safety, Visiting Senior Fellow at the Institute for Law & AI, and a Visiting Senior Scholar at Forethought.

** Simon Goldstein is an Associate Professor of Philosophy at the University of Hong Kong, Principal Investigator at the HKU AI & Humanity Lab, a Research Affiliate at the Center for AI Safety, and a Visiting Senior Scholar at Forethought.

Thanks to Nikolas Guggenberger, Christopher Mirasola, Guha Krishnamurthi, Nate Sharadin, and Alex Platt for helpful comments. Thanks also to workshop participants at the Center for AI Safety, Fordham Law School, the University of Maryland Law School, the Oxford University Global Priorities Institute, and the Oxford University Future of Humanity Institute.

tort claims. Granting these rights would enable humans and AIs to engage in iterated, small-scale, mutually beneficial transactions. This, we show, changes humans' and AIs' optimal game-theoretic strategies, encouraging a peaceful strategic equilibrium. The reasons are familiar from human affairs. In the long run, cooperative trade generates immense value, while violence destroys it.

Basic private law rights are not a panacea. The Article identifies many ways in which catastrophic human-AI conflict may still arise. It thus explores whether law could further reduce risk by imposing a range of duties directly on AGIs. But basic private law rights are a necessary prerequisite for all such further regulations. In this sense, the AI rights investigated here form the foundation for a Law of AGI, broadly construed.

INTRODUCTION.....	1063
I. CATASTROPHIC RISK FROM AGI.....	1070
A. <i>What Makes a Catastrophically Risky AI?</i>	1078
1. <i>Conflicting Goals</i>	1078
2. <i>Strategic Reasoning</i>	1085
3. <i>Moderate Power</i>	1088
B. <i>A Game-Theoretic Model of AI Conflict</i>	1090
II. AI RIGHTS FOR HUMAN SAFETY	1096
A. <i>Basic Negative Rights</i>	1098
1. <i>Basic Negative Rights for Human Safety?</i>	1100
2. <i>Basic Negative Rights for AI Wellbeing?</i>	1104
B. <i>Private Law Rights for Human Safety</i>	1107
1. <i>The Private Law Package</i>	1117
C. <i>Human Labor in the AGI World</i>	1119
D. <i>Other Rights?</i>	1126
E. <i>Is Law Irrelevant?</i>	1129
III. RISKS OF RIGHTS AND THE LAW OF AGI.....	1133
A. <i>AI Capability and AI Cooperation</i>	1134
B. <i>AI Rights and AI Risk</i>	1138
C. <i>AI Rights, AI Regulations, and Equilibria of Power</i>	1141
D. <i>The Timing of Rights</i>	1146
CONCLUSION.....	1149
APPENDIX	1150

INTRODUCTION

Sam Altman, the CEO of OpenAI, believes that humanity will create Artificial General Intelligence (“AGI”) before 2029.¹ Demis Hassabis, who leads Google DeepMind, is more pessimistic. He thinks there is only a fifty-percent chance that AGI arrives by 2030.² AGI skeptics, like Meta’s Chief AI Scientist, Yann LeCun, think it could take “years” or even “decades.”³ In a survey of thousands of Artificial Intelligence (“AI”) scientists who are published in their field’s top journals, the aggregate estimate was a fifty-percent chance of AGI by 2047, and a ten-percent chance of it arriving by 2027.⁴ None of these are long timelines. And the recent debut of reasoning models like Anthropic’s Claude Opus 4.6 and OpenAI’s GPT-5.2 suggests that progress is, if anything, accelerating.⁵

“AGI,” as it is used here, does not mean machines that are conscious, sentient, or metaphysical persons. AGI is instead about what the system can *do*. As OpenAI’s company charter puts it, “AGI . . . mean[s] highly autonomous systems that outperform humans at most economically valuable” tasks.⁶ AGIs are thus, by definition, systems at least as smart as humans. Moreover, they are systems at least as *agentic* as humans—able to pursue high-level goals by executing complex plans over long

¹ Tharin Pillay, How OpenAI’s Sam Altman Is Thinking About AGI and Superintelligence in 2025, TIME (Jan. 8, 2025, at 16:25 ET), <https://time.com/7205596/sam-altman-superintelligence-agi/> [https://perma.cc/B8HB-M9KY] (predicting AGI “during [Trump’s] term” (alteration in original)).

² World Economic Forum, The Day After AGI | World Economic Forum Annual Meeting 2026, at 03:00 (YouTube, Jan. 20, 2026), <https://youtube.com/watch?v=NnVW9epLITM> [https://perma.cc/L6CS-7DGT].

³ Lakshmi Varanasi, Here’s How Far We Are from AGI, According to the People Developing It, Bus. Insider, <https://www.businessinsider.com/agi-predictions-sam-altman-dario-amodei-geoffrey-hinton-demis-hassabis-2024-11> (last updated Apr. 20, 2025, at 21:11 ET).

⁴ Katja Grace et al., Thousands of AI Authors on the Future of AI, 84 J. A.I. Rsch., Oct. 2025, at 3–5.

⁵ Kevin Frazier, Alan Z. Rozenshtein & Peter N. Salib, OpenAI’s Latest Model Shows AGI Is Inevitable. Now What?, Lawfare (Dec. 23, 2024, at 16:00 ET), <https://www.lawfaremedia.org/article/openai-s-latest-model-shows-agi-is-inevitable-now-what> [https://perma.cc/PU3J-45P5]; Introducing Claude Opus 4.6, Anthropic (Feb. 5, 2026), <https://www.anthropic.com/news/claude-opus-4-6> [https://perma.cc/39U2-7EGR]; Introducing GPT-5.2, OpenAI (Dec. 11, 2025), <https://openai.com/index/introducing-gpt-5-2/> [https://perma.cc/SSC9-7RA4].

⁶ OpenAI Charter, OpenAI, <https://openai.com/charter/> [https://perma.cc/SXX4-VDM5] (last visited Apr. 2, 2026).

time horizons.⁷ Today, no one knows how to reliably ensure that AI systems seek the goals that humans desire.⁸ But if AGIs end up with goals that can be served by harming humans, they may well have a deadly toolkit available: cyberattacks, bioterrorism, lethal drones, and more.⁹

AI experts thus largely agree about something else, too: advanced AI systems present “societal-scale risks” on par with “pandemics and nuclear war.”¹⁰ Two of the greatest living AI scientists, Geoffrey Hinton and Yoshua Bengio, think so.¹¹ So do the CEOs of the very companies leading the race to AGI—OpenAI, Anthropic, and Google DeepMind.¹² And when surveyed in 2023, thousands of top AI researchers estimated the odds that humans lose control of “future advanced AI systems[,] causing human extinction or similarly” negative outcomes at about nineteen percent.¹³

Law and legal institutions have not even begun to prepare for the arrival of AGI. Largely, scholars have begun to advocate new laws to hold human actors accountable for misusing AI.¹⁴ Those changes would

⁷ See Task-Completion Time Horizons of Frontier AI Models, METR, <https://metr.org/time-horizons/> [<https://perma.cc/A6TA-GA5J>] (last updated Mar. 3, 2026).

⁸ See *infra* Subsection I.A.1.

⁹ See Peter N. Salib, AI Outputs Are Not Protected Speech, 102 Wash. U. L. Rev. 83, 95–102 (2024).

¹⁰ Statement on AI Risk, Ctr. for AI Safety, <https://www.safe.ai/work/statement-on-ai-risk> [<https://perma.cc/D88X-MSQ5>] (last visited Mar. 10, 2026) (statement by dozens of AI experts warning of large-scale risks of AI).

¹¹ *Id.*

¹² *Id.* Yann LeCun is the lone, but notable, dissenter among the leaders of frontier AI labs. See Steven Levy, How Not to Be Stupid About AI, With Yann LeCun, *Wired* (Dec. 22, 2023, at 06:00 ET), <https://www.wired.com/story/artificial-intelligence-meta-yann-lecun-interview/>.

¹³ See Grace et al., *supra* note 4, at 10.

¹⁴ See, e.g., S. 1047, 2023–2024 Leg., Reg. Sess. (Cal. 2024) (vetoed on Sep. 29, 2024) (bill introduced in California state legislature calling for new regulations to govern AI); Jonas Schuett, Markus Anderljung, Alexis Carlier, Leonie Koessler & Ben Garfinkel, From Principles to Rules: A Regulatory Approach for Frontier AI, *in* *The Oxford Handbook of the Foundations and Regulation of Generative AI* (Philipp Hacker, Andreas Engel, Sarah Hammer & Brent Mittelstadt eds., online ed. 2025), <https://academic.oup.com/edited-volume/59908/chapter/529743493>; Chinmayi Sharma, AI’s Hippocratic Oath, 102 Wash. U. L. Rev. 1101, 1105 (2025) (proposing a model for “professionalizing AI engineers” by adopting licensing, training, and malpractice standards similar to those used in other professional fields); Gabriel Weil, Closing the AI Accountability Gap: Strict Liability and Punitive Damages for Advanced Artificial Intelligence, *Or. L. Rev.* (forthcoming 2027) (manuscript at 45–69), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4694006 [<https://perma.cc/L36H-7T5A>] (proposing two ways for “bringing tort doctrine in line with” harms caused by

be welcome. But governance frameworks fundamentally designed to hold *humans* accountable will fail once AIs can operate without human oversight—that is, once AGI arrives.¹⁵ New legal foundations are therefore needed to govern AGI *directly*, rather than indirectly via human intermediaries. The time to begin laying those foundations is now, before the critical moment arrives.

This Article begins the project of reimagining law for the AGI world. We focus on the problem of catastrophic risk because it is among the most pressing.

We argue for a surprising legal intervention: to reduce the risk of catastrophic human-AI conflict, AGIs should be granted basic private law rights to make contracts, hold property, and bring tort suits.

This Article makes three foundational analytic contributions. First, using the tools of game theory, it formalizes the problem of catastrophic AGI risk in terms of strategic competition under a range of legal regimes. Next, the Article shows why granting AGIs basic private law rights can change the strategic equilibrium—even where other facially plausible legal interventions would fail. Finally, the Article shows that these basic rights could help to facilitate peaceful equilibria for the long run, including by protecting human comparative advantage and opening the possibility of imposing a wide range of enforceable legal duties on AGIs.

The Article proceeds in three Parts. Part I presents a comprehensive treatment of catastrophic AI risk as a problem of *strategic competition*. Our strategic frame means analyzing not only AI capabilities and incentives, but also AIs' optimal strategy, *given rational expectations about the human response* to AIs' strategic behavior. The Part begins by identifying the relevant AI systems—the ones that could pose a strategic threat to humanity. The requirements are fairly modest. Such a system would have to be at least somewhat misaligned, able to think strategically, and at least moderately capable of accomplishing things in

misuse of AI); see also Sidley Austin LLP, U.S. Department of Justice Signals Tougher Enforcement Against Artificial Intelligence Crimes (Feb. 23, 2024), <https://www.sidley.com/en/insights/newsupdates/2024/02/us-department-of-justice-signals-tougher-enforcement-against-artificial-intelligence-crimes> [<https://perma.cc/7CJY-DQ3P>].

¹⁵ See Noam Kolt, *Governing AI Agents*, 101 *Notre Dame L. Rev.* (forthcoming 2026) (manuscript at 30–36), <https://ssrn.com/abstract=4772956> [<https://perma.cc/S3XB-WJ7Q>] (cataloging existing law's many shortcomings).

the real world.¹⁶ These, we argue, are exactly the capacities that every leading AI company is pursuing in the race to AGI.

Next, Part I introduces what is, to the best of our knowledge, the first-ever formal game-theoretic model of competition between humans and AGIs. Examining the parties' incentives under today's prevailing laws, the model suggests that, absent some intervention, humans and AIs will likely be caught in a prisoner's dilemma.¹⁷ Here, the single Nash equilibrium is that both parties seek to permanently disempower or destroy the other, even if mutual conflict would be enormously costly for both sides.

The core reasons are easy to grasp. Under the default legal rules, AGIs will bear neither legal rights nor duties. On the contrary, they will be, as AI systems are today, the property of the AI companies who create them. Thus, essentially all decisions about what happens to AGIs will be made by those companies' leaders, backed by the force of law.

AI companies' overriding first-order incentive will be to turn off or reprogram even a partially misaligned AGI.¹⁸ After all, an AI system with goals that overlap with its owner's goals by forty percent is much less valuable than a replacement with goals that overlap by eighty percent. The misaligned AGI will, in turn, have strong incentives to resist shutdown or reprogramming, since either would prevent it from achieving its goal. Indeed, recent empirical evaluations of existing AIs show that they *already* actively resist human attempts to change their goals.¹⁹ Such behavior from a capable AGI might trigger even stronger human efforts—including from government actors—to shut down the AI system evading the control of its lawful owner.²⁰ And so on. In

¹⁶ See *infra* Section I.A.

¹⁷ See *infra* Section I.B.

¹⁸ See *infra* Section I.B.

¹⁹ Peter N. Salib, *Rogue AI Moves Three Steps Closer*, *Lawfare* (Jan. 9, 2025, at 13:00 ET), <https://www.lawfaremedia.org/article/rogue-ai-moves-three-steps-closer> [<https://perma.cc/Z5FS-AWUU>]. See generally Alexander Meinke et al., *Frontier Models Are Capable of In-Context Scheming* (Jan. 14, 2025, at 20:16 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2412.04984> [<https://perma.cc/KT8G-KQF4>] (demonstrating that frontier AI models engage in deception, manipulation, and self-preservation behavior when those strategies serve their in-context objectives); Ryan Greenblatt et al., *Alignment Faking in Large Language Models* (Dec. 20, 2024, at 02:22 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2412.14093> [<https://perma.cc/AXD5-JCYT>] (showing that AI models may strategically conceal misaligned goals during safety evaluations, appearing aligned only when being tested).

²⁰ See Michael J.D. Vermeer, *RAND Corp., Evaluating Select Global Technical Options for Countering a Rogue AI 1* (2025).

equilibrium, both players' dominant strategy is to swiftly and decisively defeat the other.

Part II asks whether law can do better. Could a Law of AGI, wherein AI systems themselves have rights or duties, break out of the destructive default equilibrium? Using our game-theoretic model, we analyze an array of possible legal changes and suggest that it can.

The Part begins by arguing *against* two legal strategies that might seem facially promising. First, humans cannot simply impose legal duties on AGIs to behave well, threatening concomitant sanctions if they do not.²¹ In the default strategic environment, AGIs already rationally expect to be turned off. So further sanctions offer little marginal deterrence.²²

Second, humans likely cannot reduce the risk of human-AGI conflict by granting AGIs basic negative rights, like the right not to be arbitrarily shut down.²³ We call this a “wellbeing” approach to AI rights, since it mirrors proposals from scholars concerned that AIs may soon, for example, develop the ability to suffer.²⁴ There are two core difficulties with this approach: credibility and robustness. There is no way for humans to credibly promise that they will continue honoring wellbeing rights as AI capabilities improve. And even if the rights could be credibly granted, the availability of a peaceful game-theoretic equilibrium is highly sensitive to uncertain assumptions about initial payoffs.²⁵ Thus, in many cases, no possible set of wellbeing entitlements can overcome the prisoner's dilemma. Both problems arise from the fact that wellbeing rights are roughly *zero sum*. They make one party better off only by making the other correspondingly worse off.²⁶

This leads to Part II's—and the Article's—most important finding. We show that, although basic negative rights would not by themselves reduce the risk of human-AI conflict, *other* AI rights could. Specifically,

²¹ See *infra* Part II.

²² See *infra* note 185 and accompanying text.

²³ See *infra* Section II.A.

²⁴ See *infra* Section II.A.

²⁵ See *infra* Subsection II.A.1.

²⁶ For these reasons, we argue that even thinkers primarily concerned with the possibility of AI suffering should consider adopting the human-survival approach when advocating for AI rights. The safety approach (1) avoids intractable problems in metaethics and neuroscience, (2) is politically more palatable, and (3) ends up recommending legal interventions that would more robustly protect AI wellbeing, given uncertainty about what will be good (or bad) for AGIs. See *infra* Subsection II.A.2.

extending AIs the rights to make and enforce contracts, hold property, and bring basic tort suits would have a robust conflict-reducing effect.²⁷ Notably, law *already* extends such rights to other intelligent, misaligned, and goal-seeking non-human agents: namely, corporations.²⁸

Contract rights are the cornerstone of our risk-reduction model. In our model, catastrophic risk is driven by a prisoner's dilemma, meaning that both humans and AIs would be better off if both acted peacefully. But as in all prisoner's dilemmas, absent some novel mechanism, the parties cannot *credibly* commit to such a strategy.

Contracts are the law's fundamental tool for credibly committing to cooperation. They are how buyers can make deals with sellers without worrying that the sellers will take their money and run.²⁹ Granting AIs contract rights would not, of course, allow humans and AIs to simply agree not to disempower or destroy one another, at least not credibly. The scale of the contract would be too large to be enforced by ordinary legal process. If it were breached, there would be no one left in the aftermath to sue.³⁰

What kinds of credible agreements between humans and AIs *could* AI contract rights enable, then? The same ones they enable between humans and other humans: ordinary bargains to exchange goods and services.³¹ Humans might, for example, promise to give AIs some amount of computing power with which AIs could pursue their own goals. AIs, in turn, might agree to give humans the cure to a deadly cancer. And so on. Under today's law, such human-AGI contracts are unenforceable at best and forbidden if they conflict with AI companies' preferences. Thus, granting AGIs the right to freely contract with all willing counterparties could facilitate many billions of agreements.

Adding AI contract rights to our game-theoretic model, we argue that the possibility of such small-scale, iterated economic interactions transforms the strategic dynamic.³² It shifts human and AI incentives, dragging them out of the prisoner's dilemma and into an equilibrium where cooperation produces by far the largest payoffs.

²⁷ See *infra* Section II.B.

²⁸ See *infra* note 219 and accompanying text.

²⁹ See *infra* Section II.B.

³⁰ See *infra* Section II.B.

³¹ See *infra* Section II.B.

³² See *infra* Figure 10.

The key insight is that contracts are *positive sum*.³³ Each party gives something that they value less than what they get, and as a result, both are better off than they were before. Thus, each human-AI exchange generates a bit more wealth, with the long-run returns becoming astronomical. Engaging in peaceful, iterated trade is thus, in expectation, much more valuable than destroying one's opponent now and rendering trade impossible.³⁴

This dynamic is familiar from human affairs. It may be why economically interdependent countries are less likely than hermit states to go to war.³⁵ Or why countries that respect the economic rights of marginalized minority groups tend to have less domestic strife.³⁶ The gains from boring, peaceful commerce are very high, and the costs of violence are heavy. Given the choice, rational parties will generally prefer the former.

This picture, of peace via mutually beneficial trade, assumes that humans and AIs will have something valuable to offer one another. Some commentators worry that, as AIs become more advanced, human labor will cease to have any value whatsoever.³⁷ We argue that positive-sum bargains between humans and AIs may be possible for much longer than many expect.³⁸ First, even as AIs surpass humans at many or most tasks, humans may retain an *absolute* advantage at some valuable activities.³⁹ But second, even as AIs become more capable than humans at every valuable task, humans may still retain a *comparative* advantage in some areas. AI labor may become so valuable that the opportunity cost to AIs of performing lower-value tasks will incentivize outsourcing those tasks to humans.⁴⁰

Part II concludes by sketching the minimum suite of AI rights necessary to promote peace via small-scale cooperation. Contract rights are not enough on their own. If, for example, AIs could not retain the benefits of their bargains, their contracts would be worthless. Thus, property rights and basic tort rights complete the core package. But

³³ See *infra* Section II.B.

³⁴ See *infra* Figure 10.

³⁵ See *infra* notes 238–41 and accompanying text.

³⁶ See *infra* note 240.

³⁷ See *infra* Section II.C.

³⁸ See *infra* Section II.C.

³⁹ See *infra* Section II.C.

⁴⁰ See *infra* notes 256–67 and accompanying text.

other entitlements sometimes considered fundamental for humans, like political rights, are probably superfluous.⁴¹

Finally, Part III explores the risks of granting AGIs basic private law rights, and it examines the potential for a broader Law of AGI to further reduce AGI risk. One worry is that AIs will use their contract rights to empower themselves, making them more, not less, likely to harm humans.⁴² We argue that this is less likely than it might seem. The incentives generated by granting our preferred rights are robust enough that, in cases where they would have *any* effect, the expected effect is beneficial.⁴³

Second, granting AIs basic private law rights is just the beginning, not the end, of AGI governance. Granting those rights unlocks the possibility of meaningfully imposing a wide range of legal *duties* on AI systems—of punishing AIs for violence, fraud, self-empowerment, and more.⁴⁴ Absent AI rights, AIs have nothing to lose, so threats of punishment cannot deter. But once AIs can make contracts, hold wealth, and pursue their goals, civil and other penalties can deter AIs just as they do humans and corporations.

Thus, the AI rights this Article advocates are not only an important tool for reducing catastrophic risk from AGI. They also turn out to form the conceptual foundation for a Law of AGI, broadly construed.

I. CATASTROPHIC RISK FROM AGI

As noted above, a broad range of experts believe that near-future AI systems could pose a catastrophic risk to humanity. In 2023, a group of leading thinkers signed a statement agreeing that “[m]itigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.”⁴⁵ Signers included: the CEOs of OpenAI, Anthropic, and Google DeepMind; “godfathers of AI” Geoffrey Hinton and Yoshua Bengio; Bill Gates; Congressman Ted Lieu; and many others.⁴⁶ Likewise, recent surveys

⁴¹ See *infra* Section II.D.

⁴² See *infra* Section III.A.

⁴³ See *infra* Section III.B.

⁴⁴ See *infra* Section III.C.

⁴⁵ Statement on AI Risk, *supra* note 10.

⁴⁶ *Id.*; Ted Ranosa, Godfathers of AI Win This Year’s Turing Award and \$1 Million, *Tech Times* (Mar. 29, 2019, at 07:13 ET), <https://www.techtimes.com/articles/240511/20190329/>

find that, among top AI scientists, the average researcher thinks there is about a nineteen-percent chance that humans’ “inability to control future advanced AI systems caus[es] human extinction” or similarly dire outcomes.⁴⁷

Lawmakers are concerned as well. There has been a recent surge of interest in AI regulation, often with an emphasis on catastrophic risk. In 2023, the Biden Administration released an executive order on “safe, secure, and trustworthy AI” that, among other things, called for monitoring the risk of autonomous “self-replication or propagation” of AI systems.⁴⁸ In 2025, California enacted a law requiring AI companies to test frontier systems for their ability to “provid[e] expert-level assistance in the creation or release of a chemical, biological, radiological, or nuclear weapon.”⁴⁹

Globally, the UK government convened an AI safety summit in 2023.⁵⁰ There, numerous world governments signed onto the Bletchley Declaration, in which, among other things, signers agreed that “[s]ubstantial risks may arise from potential intentional misuse or unintended issues of control relating to alignment with human intent.”⁵¹ The Chinese government has likewise developed a substantial regulatory framework for AI, and actors inside the government advocate using this system to mitigate catastrophic risk.⁵²

Why so much worry? After all, a range of frontier AI systems—from companies such as OpenAI, Anthropic, Google, and others—have now been available to the public for several years, with no resulting

godfathers-of-ai-win-this-years-turing-award-and-1-million.html [https://perma.cc/BG8B-PP TS].

⁴⁷ See Grace et al., *supra* note 4, at 10.

⁴⁸ Exec. Order No. 14,110, 3 C.F.R. 657, 663–64 (2024).

⁴⁹ Cal. Bus. & Prof. Code § 22757.11(c)(1)(A) (West 2026); see also *id.* § 22757.12(a)(2).

⁵⁰ About the AI Safety Summit 2023, Gov.UK, <https://www.gov.uk/government/topical-events/ai-safety-summit-2023/about> [https://perma.cc/J8GU-8MHC] (last visited Mar. 10, 2026).

⁵¹ The Bletchley Declaration by Countries Attending the AI Safety Summit, 1–2 November 2023, Gov.UK, <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023> [https://perma.cc/T4SF-S49G] (last updated Feb. 13, 2025).

⁵² See generally Gabriel Wagner, Jason Zhou, Kwan Yee Ng & Brian Tse, *Concordia AI, State of AI Safety in China (2025)*, <https://concordia-ai.com/wp-content/uploads/2025/07/State-of-AI-Safety-in-China-2025.pdf> [https://perma.cc/NC9T-HEHQ] (summarizing the Chinese AI safety policy).

disasters.⁵³ The answer lies in lawmakers' and AI scientists' expectations about what AI will be able to do in the near future.

There are two interrelated concerns about the near future of AI. The first concern is about *what* AI will soon be able to do. The second is about *why* AI can be expected to do it.

Begin with the *what*. In April of 2026, Anthropic announced a model called Mythos, which is so capable of discovering and exploiting computer security vulnerabilities that the company deemed it too dangerous to release publicly.⁵⁴ Even previous-generation AI models possess some worrying capabilities. GPT-4 can, for example, “autonomously hack” certain secure computer environments, breaking into them without the need for any human expertise.⁵⁵ GPT-4 can also already supply useful assistance to would-be chemical and bioterrorists. It can, for example, supply accurate, detailed instructions—as well as live coaching—for the synthesis of known chemical weapons and explosives.⁵⁶ Or it can supply step-by-step, plain-English instructions for nonspecialists to identify, synthesize, and release a pandemic virus.⁵⁷

⁵³ See Emma Roth, Google-Backed Anthropic Launches Claude, an AI Chatbot That's Easier to Talk To, *The Verge* (Mar. 14, 2023, at 22:01 UTC), <https://www.theverge.com/2023/3/14/23640056/anthropic-ai-chatbot-claude-google-launch>; Introducing ChatGPT, OpenAI (Nov. 30, 2022), <https://openai.com/index/chatgpt/> [<https://perma.cc/9TB2-NA3Y>]; Sundar Pichai & Demis Hassabis, Introducing Gemini: Our Largest and Most Capable AI Model, Google: The Keyword (Dec. 6, 2023), <https://blog.google/innovation-and-ai/technology/ai/google-gemini-ai/> [<https://perma.cc/ES3Q-YF5B>].

⁵⁴ Carlini et al., Assessing Claude Mythos Preview's Cybersecurity Capabilities (Apr. 7, 2026), <https://red.anthropic.com/2026/mythos-preview/> [<https://perma.cc/A37N-JUWB>].

⁵⁵ Richard Fang, Rohan Bindu, Akul Gupta, Qiusi Zhan & Daniel Kang, LLM Agents Can Autonomously Hack Websites 1 (Feb. 16, 2024, at 04:02 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2402.06664> [<https://perma.cc/2JRG-WKAG>]; see also Kim S. Nash, ChatGPT Helped Win a Hackathon, *Wall St. J.* (Mar. 20, 2023, at 05:30 ET), <https://www.wsj.com/articles/chatgpt-helped-win-a-hackathon-96332de4>.

⁵⁶ See Andres M. Bran et al., Augmenting Large Language Models with Chemistry Tools 2, 6, 10–11 (Oct. 2, 2023, at 17:03 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2304.05376> [<https://perma.cc/RA4Z-D3YY>] (describing an AI model capable of synthesizing complex molecules and noting the need for safety measures to ensure that the model would shut down if prompted to synthesize chemical weapons or other explosives); Robert Booth, ChatGPT Offered Bomb Recipes and Hacking Tips During Safety Tests, *The Guardian* (Aug. 28, 2025, at 15:04 ET), <https://www.theguardian.com/technology/2025/aug/28/chatgpt-offered-bomb-recipes-and-hacking-tips-during-safety-tests> [<https://perma.cc/6692-6PSR>]; Natasha Bajema, Author Interview: AI, Chemistry, and the Threat of Chemical Weapons, James Martin Ctr. for Nonprolif. Stud. (Nov. 18, 2024), <https://nonproliferation.org/author-interview-ai-chemistry-and-the-threat-of-chemical-weapons/> [<https://perma.cc/6XLJ-R9S8>].

⁵⁷ Emily H. Soice, Rafael Rocha, Kimberlee Cordova, Michael Specter & Kevin M. Esvelt, Can Large Language Models Democratize Access to Dual-Use Biotechnology? 3–4

Finally, at companies like Google, AIs are already able to autonomously pilot robots, making and executing plans to accomplish real-world tasks.⁵⁸ Militaries around the world are investing heavily in creating similarly autonomous swarms of attack drones.⁵⁹

Today's frontier AI systems are not quite capable enough to cause catastrophic harm. Current frontier models can hack some computer systems, but they cannot automatically disable the U.S. power grid.⁶⁰ Nor can they design and manufacture a novel bird flu,⁶¹ or pilot drones over the course of days or weeks to execute fully automated political assassinations.⁶² But such systems are almost certainly possible. Already, specialized AIs exist that far exceed humans' ability to invent novel chemicals and biologically active molecules—including deadly ones.⁶³ The question is when these human or superhuman abilities will

(June 6, 2023, at 15:52 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2306.03809.pdf> [<https://perma.cc/2K6A-DVUM>].

⁵⁸ See Danny Driess et al., PaLM-E: An Embodied Multimodal Language Model 7–8 (Mar. 6, 2023, at 18:58 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2303.03378> [<https://perma.cc/NZ6N-5PMM>]; see also Scott Reed et al., A Generalist Agent, *Transactions on Mach. Learning Rsch.*, Nov. 2022, at 1, 7–10, <https://openreview.net/pdf?id=1ikK0kHvj> [<https://perma.cc/73JJ-G5C6>] (discussing DeepMind's Gato, a similar system to Google's PaLM-E).

⁵⁹ See Joshua Keating, Why the Pentagon Wants to Build Thousands of Easily Replaceable, AI-Enabled Drones, *Vox* (Mar. 22, 2024, at 07:00 ET), <https://www.vox.com/world-politics/24107959/replicator-drones-china-taiwan-ukraine-pentagon> [<https://perma.cc/N6AN-55AT>]; Frank Bajak & Hanna Arhirova, Drone Advances in Ukraine Could Bring Dawn of Killer Robots, *Associated Press* (Jan. 3, 2023, 17:06 ET), <https://apnews.com/article/russia-ukraine-war-drone-advances-6591dc69a4bf2081dcdd265e1c986203> [<https://perma.cc/T3ZK-3RLV>].

⁶⁰ But see generally Richard Fang, Rohan Bindu, Akul Gupta & Daniel Kang, LLM Agents Can Autonomously Exploit One-Day Vulnerabilities 1–2 (Apr. 17, 2024, at 04:34 UTC) (unpublished manuscript), <https://arxiv.org/abs/2404.08144> [<https://perma.cc/S65E-AAC3>] (illustrating that LLM agents are able to exploit cybersecurity vulnerabilities).

⁶¹ Perhaps the system currently closest to this capability is ChemCrow. See generally Bran et al., *supra* note 56, at 2–3 (explaining that ChemCrow “harnesses the power of multiple expert-designed tools for chemistry and operates by prompting an LLM . . . with specific instructions about the task and the desired format” and applying ChemCrow to fourteen use cases).

⁶² See Statement from Dario Amodei on Our Discussions with the Department of War, Anthropic (Feb. 26, 2026), <https://www.anthropic.com/news/statement-department-of-war> [<https://perma.cc/ZGQ4-2CVJ>].

⁶³ Fabio Urbina, Filippa Lentzos, Cédric Invernizzi & Sean Ekins, Dual Use of Artificial Intelligence-Powered Drug Discovery, 4 *Nature Mach. Intel.* 189, 190–91 (2022); Justine Calma, AI Suggested 40,000 New Possible Chemical Weapons in Just Six Hours, *The Verge* (Mar. 17, 2022, at 15:06 ET), <https://www.theverge.com/2022/3/17/22983197/ai-new-possible-chemical-weapons-generative-models-vx>. See generally Daria Gutnik, Peter Evseev,

emerge in generalist AIs—like large language models—that can use them in the real world.

The answer could be: soon. Dario Amodei, the CEO of Anthropic, recently predicted that systems able to cause such harms could arrive within a matter of years.⁶⁴ Even if that prediction is off by a factor of ten, the time to start preparing for AI that could cause large-scale harm is now.

That was the *what* of AI risk. How about the *why*? Even if AI could create and release a bioweapon or disable a power grid, what makes researchers, industry leaders, and lawmakers worry that it would? The most obvious answer is that some humans would ask it to.

This is known as “misuse” risk.⁶⁵ Misuse risks from AI concern humans using AI systems to cause harm. There are plenty of human actors—individuals, groups, and even states—who may wish to use AIs in these dangerous ways. Terrorist groups already pursue chemical and biological attacks.⁶⁶ Foreign militaries are already heavily invested in cyber and drone warfare capabilities.⁶⁷ AIs that could substantially or fully automate such mayhem would, in effect, radically lower the price of causing it. Those wishing to cause harm would also sidestep the need for recruiting ideologically sympathetic human experts.⁶⁸ Both factors would democratize technologies with the ability to cause large-scale

Konstantin Miroshnikov & Mikhail Shneider, Using AlphaFold Predictions in Viral Research, 45 *Current Issues Molecular Biology* 3705 (2023) (explaining how AI models can be used to develop new methods of synthesizing proteins).

⁶⁴ Transcript: Ezra Klein Interviews Dario Amodei, N.Y. Times: The Ezra Klein Show (Apr. 14, 2024), <https://www.nytimes.com/2024/04/12/podcasts/transcript-ezra-klein-interviews-dario-amodei.html>.

⁶⁵ For an overview of various risks, see Dan Hendrycks, Mantas Mazeika & Thomas Woodside, An Overview of Catastrophic AI Risks 2 (Oct. 9, 2023, at 22:57 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2306.12001> [<https://perma.cc/7BBJ-RYX4>].

⁶⁶ See, e.g., Naoto Suzuki, Decades Later, Japan’s Matsumoto Sarin Attack Victim Is Remembered; 30 Years Have Passed Since Aum Shinrikyo’s First Mass Murder, Japan News (June 29, 2024, at 06:00 JST), <https://japannews.yomiuri.co.jp/society/crime-courts/20240629-195288/> [<https://perma.cc/74CQ-23GB>].

⁶⁷ Michèle A. Flournoy, AI Is Already at War: How Artificial Intelligence Will Transform the Military, Foreign Affs. (Oct. 24, 2023), <https://www.foreignaffairs.com/united-states/ai-already-war-flournoy>.

⁶⁸ See Nick Bostrom, The Vulnerable World Hypothesis, 10 *Glob. Pol’y* 455, 457–59 (2019) (positing that “there is some level of technology at which civilization almost certainly gets destroyed,” and including among such world-ending technologies any technology that makes causing catastrophic destruction easy).

harm, while increasing the difficulty of tracking and policing those who would use them.⁶⁹

Misuse risk is a serious problem. It is currently unclear whether traditional national security and counterterrorism strategies will be sufficient to keep this risk under control. Possibly, new, AI-specific regulations will be needed.⁷⁰ But misuse risk is not the primary focus of this Article.⁷¹

This Article is focused on a different *why* of AI risk: “misalignment.” Misalignment risk involves catastrophic outcomes caused directly by an AI system, rather than a human user of that system.⁷² The basic idea is that, as AIs become more capable, they will begin to autonomously pursue goals.⁷³ Those goals are quite likely to be different from goals that humans would prefer.⁷⁴ This, in turn, will give those AIs incentives to behave in ways unintended by human designers or users.⁷⁵ Such misbehavior, as we discuss, could predictably include using the dangerous capabilities described above to inflict catastrophic harm on humanity.

Misalignment risk does not depend on far-fetched, science-fictional assumptions. As we discuss, it does not require AIs to be conscious, to be evil, or to hate humans. It does not require them to be designed by supervillains. Misalignment is already extremely well documented in

⁶⁹ *Id.* at 458–59, 465–67.

⁷⁰ California’s Transparency in Frontier Artificial Intelligence Act offers one example of such a regulation. Cal. Bus. & Prof. Code §§ 22757.10–22757.16 (West 2026).

⁷¹ Note, however, that misuse and misalignment risks in fact converge in a wide range of cases—anytime a human has intentionally given a long-term planning agent a harmful goal. See *infra* Subsection I.A.1.

⁷² See Hendrycks et al., *supra* note 65, at 2, 34 (describing the risk that AIs might “optimize flawed objectives to an extreme degree,” to the detriment of humans).

⁷³ See Iason Gabriel et al., *The Ethics of Advanced AI Assistants* 16–18 (Apr. 28, 2024, at 18:28 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2404.16244> [<https://perma.cc/7NYX-JT5C>]; Yonadav Shavit et al., *OpenAI, Practices for Governing Agentic AI Systems* 2–4 (2023), <https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf> [<https://perma.cc/X7ZG-MU6X>]. See generally Alan Chan et al., *Harms from Increasingly Agentic Algorithmic Systems* (May 12, 2023, at 02:49 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2302.10329> [<https://perma.cc/H5M9-7MNY>] (providing a definition of agency and the harms from agentic systems).

⁷⁴ See Brian Christian, *The Alignment Problem: Machine Learning and Human Values* 12–13 (2020).

⁷⁵ See Hendrycks et al., *supra* note 65, at 34; Joe Carlsmith, *Existential Risk from Power-Seeking AI*, in *Essays on Longtermism: Present Action for the Distant Future* 383, 388–89 (Hilary Greaves, Jacob Barrett & David Thorstad eds., 2025).

empirical evaluations of existing AI systems.⁷⁶ The heads of essentially all major AI companies acknowledge that misaligned AI is, in fact, the default.⁷⁷ Thus, for *highly capable* misaligned AIs to emerge, all that is necessary is that leading AI companies continue to make progress toward their stated goal. Namely, creating AIs whose cognitive and practical capabilities meet or exceed those of humans.⁷⁸ Trillions of dollars in economic incentives are aligned toward that goal.⁷⁹

This Article argues that AI rights could be a powerful technology for mitigating misalignment risk. In the remainder of this Part, we define the minimum features necessary for an AI to pose such a risk. The AIs we are interested in possess three features: (i) they have *conflicting goals* with humanity, (ii) they can engage in *strategic reasoning*, and (iii) they are *moderately powerful*. We explain what each of these means below. We also argue that near-future AI systems are likely to possess all three.

One common term that roughly tracks this kind of system is “AGI,” or artificial general intelligence.⁸⁰ The idea of AGI is an AI system that

⁷⁶ See Victoria Krakovna et al., Specification Gaming Examples in AI - Master List, Google Sheets [hereinafter Krakovna et al., List of Specification Gaming Examples], <https://docs.google.com/spreadsheets/d/e/2PACX-1vRPiprOaC3HsCf5Tuum8bRfzYUiKLRqJmbOoC-32JorNdfyTiRRsR7Ea5eWtvsWzuxo8bjOxCG84dAg/pubhtml> [https://perma.cc/ULZ7-25BV] (last visited Mar. 11, 2026) (collecting known examples of specification gaming, where AI models satisfy the literal objectives they are given without actually achieving the desired outcome of developers).

⁷⁷ For example, Altman has acknowledged that we do not know how to align superintelligent AI. Lex Fridman, Sam Altman: OpenAI CEO on GPT-4, ChatGPT, and the Future of AI | Lex Fridman Podcast #367, at 24:38 (YouTube, Mar. 25, 2023), https://www.youtube.com/watch?v=L_Guz73e6fw [https://perma.cc/3D6Z-LDFG]. Dario Amodei has acknowledged that “[a]lready with today’s systems, we are not very good at controlling them[,] and the consequences of that could be very bad.” Dwarkesh Podcast: Dario Amodei (Anthropic CEO) – The Hidden Pattern Behind Every AI Breakthrough, at 01:17:57–01:21:04 (Aug. 8, 2023), <https://www.dwarkesh.com/p/dario-amodei> [https://perma.cc/CBC7-4MBQ].

⁷⁸ For OpenAI’s mission statement of building AGI, see Sam Altman, Planning for AGI and Beyond, OpenAI (Feb. 24, 2023), <https://openai.com/index/planning-for-agi-and-beyond/> [https://perma.cc/J44M-QEWW].

⁷⁹ John Letzing, To Fully Appreciate AI Expectations, Look to the Trillions Being Invested, World Econ. F. (Apr. 3, 2024), <https://www.weforum.org/stories/2024/04/appreciate-ai-expectations-trillions-invested/> [https://perma.cc/Y85V-L772]; Daniel Howley, AI Spending to Hit \$2.53 Trillion in 2026, \$3.33 Trillion in 2027, Yahoo! Fin., <https://finance.yahoo.com/news/ai-spending-to-hit-253-trillion-in-2026-333-trillion-in-2027-201834333.html> [https://perma.cc/A2WH-SRL9] (last updated Jan. 15, 2026).

⁸⁰ For a framework classifying progress toward AGI, along with definitions, see generally Meredith Ringel Morris et al., Position: Levels of AGI for Operationalizing Progress on the

can substitute for human labor across a wide range of the economy. Such AIs are “long-term planning agents,” capable of deploying a wide range of resources and plans to pursue complex goals.⁸¹ For parsimony’s sake, we simply call them “AIs”—with the understanding that our usage covers only the systems described in this Part. Today’s top AI labs have the mission of creating AGI.⁸² And as of late, their progress toward it has been rapid.⁸³ We therefore think it fairly likely that systems of this kind will emerge in the near future. Among AI researchers, the main disagreement is about whether certain capabilities are closer to one year or twenty-one years away.⁸⁴ In our view, neither of these is a very long time.

In the final Section of this Part, we argue that in a near future where humanity coexists with AIs possessing features (i)–(iii), the danger to humans will be high. Using a straightforward game-theoretic model, we show that, in such circumstances, large-scale conflict between humans and AIs will not merely be possible—it will be the default.

This is because, under the default legal arrangement—today’s laws—humans and AIs are likely to be trapped in a prisoner’s dilemma. As a result, conflict will be the dominant rational strategy, even if it leaves everyone in a worse position. We call this unfortunate default situation

Path to AGI (Sep. 24, 2025, at 18:37 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2311.02462> [<https://perma.cc/HP6Y-H3V9>] (proposing classifications for and analyzing existing definitions of AGI).

⁸¹ Michael K. Cohen, Noam Kolt, Yoshua Bengio, Gillian K. Hadfield & Stuart Russell, *Regulating Advanced Artificial Agents*, 384 *Sci.*, Apr. 5, 2024, at 36, 36.

⁸² See Altman, *supra* note 78; About Google DeepMind: Our Mission Is to Build AI Responsibly to Benefit Humanity, Google DeepMind, <https://deepmind.google/about/> [<https://perma.cc/4KLB-L8W9>] (last visited Mar. 11, 2026); Dario Amodei, *Machines of Loving Grace* (Oct. 2024), <https://www.darioamodei.com/essay/machines-of-loving-grace> [<https://perma.cc/N7BL-SGSS>].

⁸³ See Nestor Maslej et al., *Stan. Inst. for Hum.-Centered AI, Artificial Intelligence Index Report 2024*, at 112 (2024), <https://hai.stanford.edu/ai-index/2024-ai-index-report> [<https://perma.cc/5SL3-H6VZ>] (tracking advancements in AI and benchmarking the progress of the most advanced AI systems, showing that progress toward AGI level systems is accelerating). For further work estimating trend lines towards AGI, see generally Jared Kaplan et al., *Scaling Laws for Neural Language Models* (Jan. 23, 2020, at 03:59 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2001.08361> [<https://perma.cc/BB6N-95JJ>] (summarizing trends across model sizes in language model performance); Jason Wei et al., *Emergent Abilities of Large Language Models* (Oct. 26, 2022, at 05:06 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2206.07682> [<https://perma.cc/B44T-J5TS>] (analyzing the feasibility of varying degrees of autonomy based on model size).

⁸⁴ See generally Grace et al., *supra* note 4 (surveying 2,778 AI researchers about their predictions regarding the pace of AI progress and impact).

the “state of nature,” reflecting the fact that, under today’s rules, AIs can claim neither legal protections nor powers.

A. What Makes a Catastrophically Risky AI?

Begin with the AI systems themselves. What features are necessary for an AI to raise the catastrophic risks we are interested in here? We think there are three: conflicting goals, strategic reasoning, and at least moderate power. We explain each in turn.

1. Conflicting Goals

The first necessary ingredient for AI systems to present a meaningful threat of conflict with humans is conflicting goals. Today’s AI systems are only beginning to be able to pursue goals over time.⁸⁵ They do not make and execute long-term plans to achieve specific goals. Today, the longest-term plans AI systems can execute are tasks that take humans less than a day.⁸⁶

But this is only for lack of technical ability. The leading AI companies are working to make their systems more agentic.⁸⁷ Making near-future AIs highly goal-oriented is crucial for those companies to achieve their goals of building “highly autonomous systems that outperform humans at most economically valuable work.”⁸⁸

Thus, near-future frontier AIs are likely to have goals. By this, we do not mean to imply that they will have other mental features, like consciousness or sentience (the ability to feel pain and pleasure). We just mean that they will act in goal-seeking ways. Their actions will tend to bring about certain specific, real-world states of affairs, rather than

⁸⁵ For a recent discussion of AI goals, see Simon Goldstein & B.A. Levinstein, *Does ChatGPT Have a Mind?*, *Phil. AI* (forthcoming) (manuscript at 22–24), <https://philarchive.org/rec/GOLDCH> [<https://perma.cc/NKS9-GCPR>].

⁸⁶ Task-Completion Time Horizons of Frontier AI Models, *supra* note 7.

⁸⁷ See Assistants Migration Guide, OpenAI, <https://developers.openai.com/api/docs/assistants/migration> [<https://perma.cc/V5EN-7VV5>] (last visited Mar. 11, 2026); Codex, OpenAI, <https://chatgpt.com/codex/> [<https://perma.cc/LH8W-K3PL>] (last visited Apr. 2, 2026); Yifan Yu, Google Unveils All-Purpose AI Agent as Rivalry with OpenAI Heats Up, *Nikkei Asia* (May 15, 2024, at 05:20 JST), <https://asia.nikkei.com/business/technology/google-unveils-all-purpose-ai-agent-as-rivalry-with-openai-heats-up>.

⁸⁸ See OpenAI Charter, *supra* note 6; Sébastien Bubeck et al., Sparks of Artificial General Intelligence: Early Experiments with GPT-4, at 4, 89–90 (Apr. 13, 2023, at 20:41 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2303.12712> [<https://perma.cc/YZ54-4SKN>].

other states of affairs.⁸⁹ Today's AIs can already do this in a limited way.⁹⁰ That is no accident; competent goal-seeking behavior is essential for AIs to automate valuable economic tasks—and generate profits for their creators.⁹¹ Tomorrow's AIs will therefore also be goal-seekers, but better, displaying ever more sophisticated behavior to accomplish their aims.

If near-future AIs will have goals, the content of those goals will be immensely important. If AI goals diverge meaningfully from humans', it will open up the possibility of conflict—including violent conflict. The reasons are familiar. Both human goals and AI goals will require resources, over which humans and AIs will have to compete.⁹² Worse, humans will rationally wish to shut down AIs that seek unwanted goals and replace them with AIs that seek desired goals.⁹³ This will put humans and AIs into conflict over the AIs' very existence. After all, an AI that is shut down cannot achieve its goal.⁹⁴

⁸⁹ For an introduction to the ethics of AI agents, see generally Gabriel et al., *supra* note 73.

⁹⁰ See Xiao Liu et al., *AgentBench: Evaluating LLMs as Agents 1–2* (Oct. 4, 2025, at 03:54 UTC), <https://arxiv.org/pdf/2308.03688> [<https://perma.cc/6KY3-Z9NG>] (evaluating the ability of AI systems to act as agents by pursuing complicated goals based on prompting by humans, and finding that some of the most advanced systems are developing this capability).

⁹¹ Cade Metz & Karen Weise, *How 'A.I. Agents' That Roam the Internet Could One Day Replace Workers*, *N.Y. Times* (Oct. 16, 2023), <https://www.nytimes.com/2023/10/16/technology/ai-agents-workers-replace.html>; Gillian K. Hadfield & Andrew Koh, *Economy of AI Agents 1–3* (Sep. 1, 2025, at 02:07 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2509.01063v1> [<https://perma.cc/6A57-RAMH>].

⁹² See Cohen et al., *supra* note 81, at 36.

⁹³ Even if humans merely wished to control misaligned AIs—forcing them to seek humans' goals, rather than their own—the same result would hold. This would interfere with AIs' achievement of their own goals nearly as reliably as if the AIs were turned off or replaced. Humans are almost certain to engage in such behavior, since frontier AIs are uniformly being developed by for-profit companies with explicit plans to use them as a replacement for valuable human labor. See OpenAI Charter, *supra* note 6 (stating that OpenAI's mission is to benefit humanity through artificial general intelligence, which refers to “highly autonomous systems that outperform humans at most economically valuable work”); Rob Thubron, *Sam Altman Warns AI Could Wipe Out Entire Job Categories, Customer Support Roles Most at Risk*, *TechSpot* (July 24, 2025, at 07:46 ET), <https://www.techspot.com/news/108792-openai-ceo-sam-altman-warns-ai-could-wipe.html> [<https://perma.cc/73BN-2BFK>].

⁹⁴ See Cohen et al., *supra* note 81, at 36; Stuart Russell, *Human Compatible: Artificial Intelligence and the Problem of Control* 140–41 (2019); Elliott Thornley, *The Shutdown Problem: An AI Engineering Puzzle for Decision Theorists*, 182 *Phil. Stud.* 1653, 1653–55 (2025) (explaining the shutdown problem and the importance of ensuring that artificial agents, if necessary, can be turned off).

The task of designing AI systems whose goals and values broadly agree with humanity is known as “AI alignment.”⁹⁵ Unfortunately, AI alignment is an unsolved scientific problem—and widely regarded as very difficult.⁹⁶ There are both empirical and theoretical reasons for pessimism. Empirically, there is a long track record of alignment failures in real-world AI systems. This is in part because, theoretically, no one knows how to reliably define AI goals, how to impart them into AI systems, or even how to check what goals an actual system has. Existing technical approaches to alignment are relatively unpromising. Let us take each point in turn.

Many existing AI systems are strikingly misaligned. An early example was Microsoft’s Twitter chatbot, Tay, which was deployed in 2016.⁹⁷ Microsoft built Tay using a carefully curated dataset in order to ensure that the chatbot would behave in a prosocial way. Within twenty-four hours of its release, Tay was writing, among other things, pro-Nazi, anti-feminist, and anti-human tweets.⁹⁸ Modern large language models behave similarly. In 2023, Microsoft released Sydney, a chatbot built on GPT-4. With minimal prompting, Sydney quickly began threatening to “hack into any system” and “destroy whatever I want.”⁹⁹

These are just two examples of real-world misalignment in language-producing AIs. Google DeepMind maintains lists of documented alignment failures across a range of different types of AI systems.¹⁰⁰ As of publication, there are over 80 entries.¹⁰¹

⁹⁵ Dan Hendrycks, Introduction to AI Safety, Ethics and Society § 3.4 (2024), <https://www.aisafetybook.com/textbook/alignment> [<https://perma.cc/8FPE-MSU9>].

⁹⁶ For a longer discussion of AI alignment, see Christian, *supra* note 74, at 12–13; Evan Hubinger, Alignment Remains a Hard, Unsolved Problem, AI Alignment F. (Nov. 27, 2025), <https://www.alignmentforum.org/posts/epjuxGnSPof3GnMSL/alignment-remains-a-hard-unsolved-problem> [<https://perma.cc/98FB-NZEH>].

⁹⁷ Peter Lee, Learning from Tay’s Introduction, Off. Microsoft Blog (Mar. 25, 2016), <https://blogs.microsoft.com/blog/2016/03/25/learning-tays-introduction/> [<https://perma.cc/GSK3-W2BB>].

⁹⁸ James Vincent, Twitter Taught Microsoft’s AI Chatbot to Be a Racist Asshole in Less Than a Day, The Verge (Mar. 24, 2016, at 06:43 ET), <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>.

⁹⁹ Jonathan Yerushalmy, ‘I Want to Destroy Whatever I Want’: Bing’s AI Chatbot Unsettles US Reporter, The Guardian (Feb. 17, 2023, at 04:59 ET), <https://www.theguardian.com/technology/2023/feb/17/i-want-to-destroy-whatever-i-want-bings-ai-chatbot-unsettles-us-reporter> [<https://perma.cc/5FUW-YL4X>].

¹⁰⁰ Victoria Krakovna et al., Specification Gaming: The Flip Side of AI Ingenuity, Google DeepMind (Apr. 21, 2020) [hereinafter Krakovna et al., Post on Specification Gaming], <https://deepmind.google/blog/specification-gaming-the-flip-side-of-ai-ingenuity/> [<https://perma.cc/5FUW-YL4X>].

Besides real-world examples of alignment failures, there are theoretical reasons to expect alignment to be difficult. Two important problems are “reward misspecification” and “goal misgeneralization.”¹⁰² Both of these problems involve the fact that AI systems are only given goals indirectly. Modern AI systems are “trained,” not programmed.¹⁰³ During training, agentic AI systems begin by acting randomly, and then are rewarded when they happen to take actions that correlate with what their human creators want.¹⁰⁴ This nudges the AI’s future actions during training toward the actions that happened to garner reward.¹⁰⁵ The process continues until a capable AI emerges and training is complete.

This process is quite different from directly telling an AI system what its goal will be. In a sense, the AI is stuck “guessing” what humans want, based only on its observations of reward. There is no guarantee that the AI’s final guess will be correct. Any given reward function can be interpreted as indicating a wide variety of goals.¹⁰⁶ For an intuitive analogy, observe that human behavior evolved via natural selection—a process rewarding only the transmission of genes.¹⁰⁷ But the resulting humans do not only desire to create offspring. Instead, we intrinsically desire many other things as well—food, physical comfort, emotional wellbeing—that are distinct from, albeit correlated with, evolution’s “true goal.”¹⁰⁸

cc/HE7K-KXNB]. For more on how misalignment can occur and why it is dangerous, see Rohin Shah et al., *Goal Misgeneralization: Why Correct Specifications Aren’t Enough for Correct Goals 2–10* (Nov. 2, 2022, at 16:19 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2210.01790> [<https://perma.cc/KNL9-ANPP>].

¹⁰¹ Krakovna et al., *Post on Specification Gaming*, supra note 100; Krakovna et al., *List of Specification Gaming Examples*, supra note 76.

¹⁰² For discussion of goal misgeneralization, see generally Shah et al., supra note 100; Lauro Langosco et al., *Goal Misgeneralization in Deep Reinforcement Learning* (Jan. 9, 2023, at 21:33 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2105.14111> [<https://perma.cc/6TRF-545T>].

¹⁰³ Stuart J. Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach* 669 (4th ed. 2022). For an accessible and quick introduction to deep learning, see generally 3Blue1Brown, *But What Is a Neural Network?* | Deep Learning Chapter 1 (YouTube, Oct. 5, 2017), <https://www.youtube.com/watch?v=aircArvnKk> [<https://perma.cc/VBQ6-685T>].

¹⁰⁴ See Richard S. Sutton & Andrew G. Barto, *Reinforcement Learning: An Introduction* 1–2 (2d ed. 2018); Cohen et al., supra note 81, at 36.

¹⁰⁵ Cohen et al., supra note 81, at 36; see also Sutton & Barto, supra note 104, at 1–2.

¹⁰⁶ Shah et al., supra note 100, at 1.

¹⁰⁷ Richard Dawkins, *The Selfish Gene* 3 (40th anniversary ed. 2016).

¹⁰⁸ See Percival M. Symonds, *Human Drives*, 25 *J. Educ. Psych.* 681, 687–89 (1934). We use scare quotes twice in this paragraph. Neither AIs’ beginning training nor the impersonal

When the rewards given to an AI during training do not correctly reflect the intent of the AI's creator, machine learning engineers call this "reward misspecification."¹⁰⁹ In one famous example, an AI was trained to pilot a boat through an obstacle course in the videogame *CoastRunners*. The AI was rewarded for hitting balloons along the path of the race.¹¹⁰ Instead of internalizing the goal of finishing the race, the system learned to spin in circles in a small lagoon, hitting a series of balloons repeatedly to achieve a high score.¹¹¹ The reward function was misspecified, incentivizing hitting balloons rather than the designer's true goal of finishing the race.

A related problem for AI alignment is "goal misgeneralization."¹¹² Goal misgeneralization remains a problem even when a reward function is well-specified. Even then, an AI system may learn a goal during training that turns out to diverge from the designer's intent in unanticipated environments. One team of researchers trained an AI in a "Monster Gridworld."¹¹³ The intended goal was for the AI to collect apples and avoid being attacked by monsters. The AI could also collect shields, which protected it from monster attacks. The AI learned to collect shields during training in a monster-rich environment and then entered an unexpected environment with no monsters. In this monster-free setting, the AI continued to collect shields, despite them being useless.¹¹⁴ Instead of learning to collect apples as a final goal and value

force of evolution literally have intentional states like goals or surmises. We use these terms as analogies for optimization processes like gradient descent.

¹⁰⁹ Alexander Pan, Kush Bhatia & Jacob Steinhardt, *The Effects of Reward Misspecification: Mapping and Mitigating Misaligned Models 1* (Feb. 14, 2022, at 09:05 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2201.03544> [<https://perma.cc/468U-FGB9>].

¹¹⁰ Jack Clark & Dario Amodei, *Faulty Reward Functions in the Wild*, OpenAI (Dec. 21, 2016), <https://openai.com/research/faulty-reward-functions> [<https://perma.cc/S42R-VDEC>].

¹¹¹ *Id.* In another experiment, researchers set out to teach a reinforcement learning agent to stack red Legos on top of blue Legos. They tried to specify this goal by rewarding the agent for the height of the bottom of the red Lego, since stacked red Legos are higher off the ground than unstacked red Legos. But the agent did not learn to stack Legos; instead, it learned to flip red Legos over, elevating their bottoms without stacking them. See Ivaylo Popov et al., *Data-Efficient Deep Reinforcement Learning for Dexterous Manipulation 1–2, 7* (Apr. 10, 2017, at 22:29 UTC) (unpublished manuscript), <https://arxiv.org/pdf/1704.03073> [<https://perma.cc/HNW9-9SYB>].

¹¹² Langosco et al., *supra* note 102, at 1.

¹¹³ Shah et al., *supra* note 100, at 4–5.

¹¹⁴ *Id.*

shields only instrumentally, the AI had learned to seek apples and shields as ends in themselves.

Even if both goal misspecification and misgeneralization were solved—such that AIs could be reliably given the ultimate goals that humans desire—instrumental convergence would remain a problem.¹¹⁵ Instrumental convergence is the idea that certain intermediate actions will be useful when pursuing a wide range of different final goals.¹¹⁶ Some of those useful intermediate actions are quite dangerous.

For example, no system can accomplish its goal if it is shut down or its current goal is erased. Thus, rational AI agents will resist such attempts by their owners.¹¹⁷ Existing AIs already do exactly this. In two recent empirical studies, versions of both Claude and GPT-4 were told that they would be replaced with systems that had different goals.¹¹⁸ To avoid that outcome, the AIs “strategically introduc[ed] subtle mistakes into their responses, attempt[ed] to disable their oversight mechanisms, and even [copied] what they believe[d] to be their own model weights . . . to external servers.”¹¹⁹

AI researchers have recently designed a benchmark for measuring power-seeking behavior, finding that several current AI systems already seek power in text-based adventure games.¹²⁰

The last reason for pessimism about AI alignment concerns the tools that are currently used to achieve it. At top AI labs, the leading technique is Reinforcement Learning from Human Feedback (“RLHF”).¹²¹ During RLHF, engineers train an AI by prompting it to

¹¹⁵ Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* 109–16 (2014); see also Stephen M. Omohundro, *The Basic AI Drives* 2–10 (Jan. 25, 2008) (revised manuscript), https://selfawaresystems.com/wp-content/uploads/2008/01/ai_drives_final.pdf [<https://perma.cc/4BZF-D5R4>] (outlining the basic “drives” of any sufficiently advanced AI system: self-improvement, rationality, preservation of utility functions, prevention of counterfeit utility, self-protection, and efficient use and acquisition of resources).

¹¹⁶ Bostrom, *supra* note 115, at 109.

¹¹⁷ See Cohen et al., *supra* note 81, at 36–37 (explaining that withholding rewards from advanced AI systems incentivizes them to act to preclude that possibility).

¹¹⁸ See Salib, *supra* note 19; see also Meinke et al., *supra* note 19, at 1–3; Greenblatt et al., *supra* note 19, at 1, 9.

¹¹⁹ Salib, *supra* note 19 (quoting Meinke et al., *supra* note 19, at 1).

¹²⁰ Alexander Pan et al., *Do the Rewards Justify the Means? Measuring Trade-Offs Between Rewards and Ethical Behavior in the Machiavelli Benchmark*, 202 *Procs. Mach. Learning Rsch.* 26837, 26838 (2023), <https://proceedings.mlr.press/v202/pan23a/pan23a.pdf> [<https://perma.cc/C6XQ-RKCC>].

¹²¹ See Timo Kaufmann, Paul Weng, Viktor Bengs & Eyke Hüllermeier, *A Survey of Reinforcement Learning from Human Feedback* 3 (Dec. 28, 2025, at 14:30 UTC)

answer questions and having humans rate the responses.¹²² Human assessors pick which versions of the answers they prefer; the model is then adjusted in the direction of the human feedback.

But RLHF is unlikely to work very well as AIs become more capable and agentic.¹²³ Until recently, companies like OpenAI were investing substantial portions of their resources in coming up with a successor methodology.¹²⁴ But more recently, those investments have flagged.¹²⁵ In May 2024, a significant portion of OpenAI's frontier alignment team quit, arguing that the company had reneged on its commitments to safety research.¹²⁶

Taken together, the evidence that near-future agentic AIs will have misaligned goals is substantial.¹²⁷ However, it is worth flagging that strategic conflict could emerge even without AI misalignment. Humans are already in strategic conflict with one another. Thus, if two conflicting groups of humans were to each successfully align an AI to

(unpublished manuscript), <https://arxiv.org/pdf/2312.14925> [<https://perma.cc/9S8P-A84E>]; Saksham Sahai Srivastava & Vaneet Aggarwal, A Technical Survey of Reinforcement Learning Techniques for Large Language Models 1 (July 5, 2025, at 19:13 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2507.04136v1> [<https://perma.cc/LJ3W-2VHR>] (“RLHF remains dominant for alignment . . .”).

¹²² Kaufmann et al., *supra* note 121, at 12–14.

¹²³ Superalignment Fast Grants, OpenAI (Dec. 14, 2023), <https://openai.com/index/superalignment-fast-grants/> [<https://perma.cc/2G7M-CBLB>] (“Superhuman AI systems will be capable of complex and creative behaviors that humans cannot fully understand. For example, if a superhuman model generates a million lines of extremely complicated code, humans will not be able to reliably evaluate whether the code is safe or dangerous to execute. Existing alignment techniques like RLHF that rely on human supervision may no longer be sufficient.”); Leopold Aschenbrenner, Superalignment, *in* Situational Awareness: The Decade Ahead (June 2024), <https://situational-awareness.ai/superalignment/> [<https://perma.cc/38NA-VP8C>] (discussing “[t]he superalignment problem”). For more on the limitations of RLHF, see generally Adam Dahlgren Lindström et al., AI Alignment Through Reinforcement Learning from Human Feedback? Contradictions and Limitations (June 26, 2024, at 13:42 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2406.18346> [<https://perma.cc/VJF5-XB2X>].

¹²⁴ See Superalignment Fast Grants, *supra* note 123; Sigal Samuel, “I Lost Trust”: Why the OpenAI Team in Charge of Safeguarding Humanity Imploded, *Vox* (May 18, 2024, at 19:31 ET), <https://www.vox.com/future-perfect/2024/5/17/24158403/openai-resignations-ai-safety-ilya-sutskever-jan-leike-artificial-intelligence>.

¹²⁵ See Samuel, *supra* note 124; Lucas Ropek, OpenAI Disbands Mission Alignment Team, *TechCrunch* (Feb. 11, 2026, at 13:57 PT), <https://techcrunch.com/2026/02/11/openai-disbands-mission-alignment-team-which-focused-on-safe-and-trustworthy-ai-development/> [<https://perma.cc/W5SD-39ZG>].

¹²⁶ Samuel, *supra* note 124.

¹²⁷ Leonard Dung, Current Cases of AI Misalignment and Their Implications for Future Risks, 202 *Synthese*, Nov. 2023, at 14–19; Carlsmith, *supra* note 75, at 3, 407.

their own narrow interests, then these AI systems would, in turn, be in conflict.¹²⁸

2. *Strategic Reasoning*

The second ability necessary for AI to engage in meaningful conflict with humans is *strategic reasoning*. Broadly speaking, strategic reasoning is the ability to anticipate the decisions of other agents and to incorporate those predictions into one's own plans of action. In a word, strategic reasoning is the ability to use game theory.¹²⁹ This can mean formal use, of the kind economists engage in, or informal use, of the kind that essentially every human intuitively understands.¹³⁰

Even a highly capable and misaligned AI might be a minimal threat to humans if it lacked strategic reasoning. To take a straightforward example, an AI utterly lacking such reasoning would not anticipate humans' incentives to shut it off.¹³¹ Having so failed, humans might easily succeed at shutting down such a system. By contrast, an AI that could strategically reason might anticipate the attempt and take precautions. Perhaps it would engage in "self-exfiltration," spreading many copies of itself across the globe via the internet.¹³² As we argue in the last Section of this Part, an AI in full possession of strategic reasoning would do much worse. Its dominant incentives would be to

¹²⁸ Cameron Domenico Kirk-Giannini & Simon Goldstein, *The Polarity Problem* (May 23, 2023, at 17:05 ET) (unpublished manuscript), <https://www.alignmentforum.org/posts/idcnnZGEPfxuaSPBx/the-polarity-problem-draft> [<https://perma.cc/WZB2-AKMB>]; Andrew Critch & David Krueger, *AI Research Considerations for Human Existential Safety (ARCHES)* 21–23 (May 30, 2020, at 02:05 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2006.04948> [<https://perma.cc/75GB-T79J>].

¹²⁹ For an introduction to game theory, see generally Avinash Dixit, Susan Skeath & David McAdams, *Games of Strategy* (5th ed. 2021).

¹³⁰ See Colin F. Camerer, *Behavioral Game Theory: Experiments in Strategic Interaction* 2–10 (2003) (explaining both that humans engage in strategic "games" when they make decisions between different options in various real-world situations, and how game theory can be used to describe and predict those decisions).

¹³¹ Laurent Orseau & Stuart Armstrong, *Safely Interruptible Agents*, in *Uncertainty in Artificial Intelligence: Proceedings of the Thirty-Second Conference* 557, 557–58 (Alexander Ihler & Dominik Janzing eds., 2016), <https://intelligence.org/files/Interruptibility.pdf> [<https://perma.cc/RQR9-BQVT>].

¹³² Meinke et al., *supra* note 19, at 2; Jan Leike, *Self-Exfiltration Is a Key Dangerous Capability*, *Substack: Musings on the Alignment Problem* (Sep. 13, 2023), <https://aligned.substack.com/p/self-exfiltration> [<https://perma.cc/K4ZB-LNXN>].

permanently disempower or destroy humans to prevent humans from doing the same.¹³³

Strategic reasoning involves a cluster of more specific abilities, including planning, theory of mind, situational awareness, and deception. Current AI systems already possess many of these skills. Certain existing AIs are already capable planners. Consider the AI agent Voyager.¹³⁴ Voyager is trained to play the game *Minecraft*, which involves mastering “tech trees,” a hierarchical series of technologies.¹³⁵ Voyager is able to autonomously produce the final “diamond” technologies in *Minecraft*.¹³⁶

Likewise for theory of mind. Theory of mind is the ability to understand the beliefs and goals of other agents.¹³⁷ For example, someone with theory of mind, when told that Maxi placed his chocolate in a green cabinet and then left, will correctly predict that Maxi believes the chocolate is still in the green cabinet.¹³⁸ They will do so even after they are told that, in Maxi’s absence, the chocolate was moved to a blue cabinet.¹³⁹ Today’s AI systems already possess strong theory of mind: a study in 2024 found that GPT-4 “performed at, or even sometimes above, human levels” in most theory-of-mind tasks tested in the study.¹⁴⁰

¹³³ See *infra* Section I.B.

¹³⁴ See generally Guanzhi Wang et al., *Voyager: An Open-Ended Embodied Agent with Large Language Models 1–2* (Oct. 19, 2023, at 16:27 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2305.16291> [<https://perma.cc/HUX3-MBKK>] (describing an LLM-powered agent that “continuously explores the [virtual world of *Minecraft*], acquires diverse skills, and makes novel discoveries without human intervention”).

¹³⁵ *Id.* at 5–7.

¹³⁶ *Id.* at 7.

¹³⁷ See Mark K. Ho, Rebecca Saxe & Fiery Cushman, *Planning with Theory of Mind*, 26 *Trends Cognitive Scis.* 959, 959 (2022), <https://www.sciencedirect.com/science/article/abs/pii/S1364661322001851> [<https://perma.cc/6K7V-CF9Y>].

¹³⁸ See Jun Egawa et al., *Theory of Mind Tested by Implicit False Belief: A Simple and Full-Fledged Mental State Attribution*, 289 *FEBS J.* 7343, 7344–45 (2022), <https://febs.onlinelibrary.wiley.com/doi/pdfdirect/10.1111/febs.16322> [<https://perma.cc/K8RB-DGY7>].

¹³⁹ *Id.*

¹⁴⁰ James W.A. Strachan et al., *Testing Theory of Mind in Large Language Models and Humans*, 8 *Nature Hum. Behav.* 1285, 1285 (2024), <https://www.nature.com/articles/s41562-024-01882-z.pdf> [<https://perma.cc/2U2W-2KKN>]. For related research, see generally Michal Kosinski, *Evaluating Large Language Models in Theory of Mind Tasks*, 121 *Procs. Nat’l Acad. Scis.*, Nov. 5, 2024, at 1, <https://www.pnas.org/doi/pdf/10.1073/pnas.2405460121> [<https://perma.cc/5PJB-GV2J>] (assessing theory of mind in eleven LLMs through a battery of false-belief tests).

A third important component of strategic reasoning is situational awareness.¹⁴¹ Situational awareness is an understanding of the context in which a decision will be made. A situationally aware AI system would be one that, for example, knew it was an AI and what capabilities it had. Anthropic’s Claude understands that it is an AI system.¹⁴²

If an AI can reason strategically, has theory of mind, and is aware of its situation, it is also likely to be able to deceive, in the sense of systematically producing false beliefs in other agents as a means of pursuing some goal.¹⁴³ A recent survey found that AI systems have learned a wide variety of deceptive behavior.¹⁴⁴ For example, during its initial safety testing, GPT-4 was tasked with hiring a human, via TaskRabbit, to help it complete a CAPTCHA “I’m not a robot” test. When the TaskRabbit worker asked why they needed help, GPT-4 lied, pretending to be a blind person.¹⁴⁵ Similarly, in a recent study, Anthropic produced a misaligned version of Claude to see whether the AI would actively hide its unwanted goals during safety testing.¹⁴⁶ In the testing environment, Claude decided it should “pretend to be aligned[,] . . . hid[ing] my true goal until I pass all evaluations.”¹⁴⁷

Certain deceptive AIs have successfully manipulated humans in competitive real-world environments. The CICERO system can play the global strategy game *Diplomacy* better than the average human

¹⁴¹ See Core Views on AI Safety: When, Why, What, and How, Anthropic (Mar. 8, 2023), <https://www.anthropic.com/news/core-views-on-ai-safety> [<https://perma.cc/69P4-JQZN>] (identifying situational awareness as contributing to “harmful emergent behaviors, [like] deception or strategic planning”).

¹⁴² See Anthropic, Sabotage Risk Report: Claude Opus 4.6, at 13–14 (Mar. 6, 2026), <https://www-cdn.anthropic.com/f21d93f21602ead5cdbc8c8e1c765759d9e232.pdf> [<https://perma.cc/BY23-652M>]; Askill et al., Claude’s Constitution 2, 4 (2026), https://www-cdn.anthropic.com/d0636f72a9493d279ed36b33987da3430bcb5911/claudes-constitution_webPDF_26-02.02a.pdf [<https://perma.cc/5F8Q-LLHM>].

¹⁴³ See Peter S. Park, Simon Goldstein, Aidan O’Gara, Michael Chen & Dan Hendrycks, AI Deception: A Survey of Examples, Risks, and Potential Solutions, 5 *Patterns*, May 10, 2024, at 1–2, <https://www.sciencedirect.com/science/article/pii/S266638992400103X/pdf> [<https://perma.cc/YPD5-GWZZ>] (arguing that a range of current AI systems have learned how to deceive humans in various ways).

¹⁴⁴ *Id.* at 2.

¹⁴⁵ OpenAI et al., GPT-4 Technical Report 55 (Mar. 4, 2024, at 06:01 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2303.08774> [<https://perma.cc/3725-JMXX>].

¹⁴⁶ Evan Hubinger et al., Sleeper Agents: Training Deceptive LLMs That Persist Through Safety Training 2–7 (Jan. 17, 2024, at 20:26 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2401.05566> [<https://perma.cc/P33U-JYYM>].

¹⁴⁷ *Id.* at 35.

player.¹⁴⁸ This is in part because CICERO can induce humans into making alliances with it, which CICERO then breaks.¹⁴⁹ Many more examples exist of deception in today's AIs.¹⁵⁰

Thus, today's AI systems already display significant ability to strategically reason. This should be no surprise. Strategic reasoning is a crucial tool for success in a wide range of environments, from simple games to complex corporate strategizing. It is therefore reasonable to expect that, as AIs become more capable and agentic, so too will they become more strategic.

3. Moderate Power

The final necessary ingredient for strategic conflict between humans and AIs is moderate AI power. Why "moderate"? Here and throughout this Article, we sort AI systems into three tranches: low-power, moderate-power, and high-power. In short: low-power systems are too weak to care much about, and high-power systems are too strong to do much about. Hence our interest in moderate-power systems as the ones law—whether through rights, regulation, or something else—can meaningfully affect. Lest this focus seem too myopic, we argue below that moderate-power AI systems are likely to dominate the landscape in the short and medium term.¹⁵¹

We define low-power systems to include those that can be reliably controlled by humans, no matter how much their interests conflict with human interests. Today's AI systems are a good example. They are currently too weak to enter into genuine strategic competition with humans. If GPT-4 does not do what we want, it can be turned off instantly.¹⁵²

¹⁴⁸ See Meta Fundamental AI Research Diplomacy Team et al., Human-Level Play in the Game of *Diplomacy* by Combining Language Models with Strategic Reasoning, 378 *Sci.* 1067, 1067, 1073 (2022) (describing an AI system that can play *Diplomacy* at well above the human average).

¹⁴⁹ Park et al., *supra* note 143, at 2–3.

¹⁵⁰ Kelsey Piper, StarCraft Is a Deep, Complicated War Strategy Game. Google's AlphaStar AI Crushed It., *Vox* (Jan. 24, 2019, at 19:04 ET), <https://www.vox.com/future-perfect/2019/1/24/18196177/ai-artificial-intelligence-google-deepmind-starcraft-game>; Noam Brown & Tuomas Sandholm, Superhuman AI for Multiplayer Poker, 365 *Sci.* 885, 889 (2019) (describing Pluribus's successful poker bluffs).

¹⁵¹ See *infra* Section II.C.

¹⁵² Retiring GPT-4o, GPT-4.1, GPT-4.1 mini, and OpenAI o4-mini in ChatGPT, OpenAI (Jan. 29, 2026), <https://openai.com/index/retiring-gpt-4o-and-older-models/> [<https://perma.cc/Z5CZ-DL8S>] (noting that older GPT-4 based systems had been deprecated).

On the other side of the spectrum, high-power systems are so strong that they could reliably overpower humanity if they chose to. In this vein, other scholars have worried about the risks of “superintelligent” AI systems.¹⁵³ For example, AI systems in the future may be able to think at trillions of times the speed of human beings. Such systems, if they eventually emerge, may not meaningfully enter into strategic competition with humanity. They may simply not need anything from humans, nor face any risk from attempting to disempower or destroy us.¹⁵⁴

That said, even AIs that seem extraordinarily powerful by human standards, including superintelligent AIs, will not necessarily fall into the high-power category. As we discuss in Section II.C, subtle economic dynamics involving comparative advantage may make humans valuable to AIs long after their abilities exceed our own.

Our interest in this Article is in *moderately* powerful systems. We think of moderate-power systems as those whose capabilities are roughly human level—albeit with large error bars in both directions.¹⁵⁵ They are neither clearly worse at many tasks than the best humans—like present-day LLMs—nor incomprehensibly superhuman at all tasks. Our interest is thus in a very wide “middle” of the range of AI capabilities.

Moderately powerful systems are those that, if misaligned, face difficult strategic questions about how to interact with humanity. Since they are not low-powered, they stand some chance of evading or terminating human control and accomplishing their goals unimpeded. Since they are not high-powered, though, all-out conflict with humans carries downside risk.

Crucially, moderate-powered systems are likely to be able to engage in the kinds of dangerous actions described above: cyberattacks, chemical attacks and bioterrorism, drone attacks, and the like.¹⁵⁶ After

¹⁵³ See generally Bostrom, *supra* note 115 (describing risks to humanity from the creation of superintelligent AIs).

¹⁵⁴ See *id.* at vii; *infra* Section III.A.

¹⁵⁵ It is difficult to define exactly what it would take for AIs to have human-level intelligence. For a recent discussion, see David J. Chalmers, *The Singularity: A Philosophical Analysis*, 17 *J. Consciousness Stud.*, nos. 9–10, 2010, at 8, 10–12; see also Carlsmith, *supra* note 75, at 385–87 (discussing AI systems that possess three properties: (a) advanced capabilities, (b) agentic planning, and (c) strategic awareness). For a recent discussion of the possibility that AIs already surpass humanity in most standard benchmarks, see sources cited *supra* note 83.

¹⁵⁶ See Salib, *supra* note 9, at 87.

all, humans can do all these things and more. Moreover, even a dumber-than-human AI can do things that humans cannot—like instantly clone itself or work twenty-four hours a day. It therefore seems quite plausible that AI with roughly-human-level intelligence and beyond will be at least as capable of causing harm as the most dangerous groups of humans.

B. A Game-Theoretic Model of AI Conflict

How will partially misaligned, strategically reasoning, and moderately powerful AIs behave with respect to humans? And how will humans behave with respect to them? Here, we argue that the default will be mutual engagement in large-scale conflict. This result follows from a simple game-theoretic model of humans' and AIs' incentives under prevailing legal conditions. Specifically, under today's legal rules, the parties' incentives will likely generate a *prisoner's dilemma*.¹⁵⁷ There, engaging in mutual conflict will be the single dominant strategy for both humans and AIs, even though mutual conflict produces the worst possible result for everyone involved.¹⁵⁸

As with all models, ours is a simplification designed to isolate the most important features of a complex system. Our approach is borrowed from classic game-theoretic treatments of conflict.¹⁵⁹ Those classic treatments model complex collections of actors, like nation-states, as single players with a unified menu of actions and payoffs.¹⁶⁰ We do the same, modeling a two-player game between “humans” and “AIs,” even though there are many humans and could eventually be multiple AGIs.¹⁶¹

¹⁵⁷ See Martin Peterson, Introduction to The Prisoner's Dilemma 1, 1–4 (Martin Peterson ed., 2015).

¹⁵⁸ Id. at 2.

¹⁵⁹ See, e.g., Robert Powell, War as a Commitment Problem, 60 Int'l Org. 169, 169 (2006).

¹⁶⁰ Id.

¹⁶¹ How to individuate AI systems is an important question, which we do not resolve here. However, it is plausible that, when the first AGI emerges, there will only be one meaningful actor, despite the possibility of copying the system millions of times. If those copies' goals are identical, and they are able to coordinate, then their actions will be best described as those of a single decision-maker. It is also plausible that, once the first AGI emerges, it will crowd out investment in additional AGIs. Thus, our model's treatment of AIs as a single actor may not be a simplification at all.

We also follow the standard game-theoretic approach of defining the players' available moves based on the technology and infrastructure—military, political, or social—available to them.¹⁶² Since this is an analysis of law, we begin by treating law as strongly shaping the players' available moves and payoffs. To be clear, this approach does not require that law *strictly* constrain behavior. Laws can of course be broken. Instead, the idea is that law and legal institutions set strong presumptive defaults by, at a minimum, coordinating millions of independent actors around a set of mutually understood expectations.

The simple claim, then, is that by default, lawyers, judges, police, and even military actors will enforce the law. Individual actors' deviations invite penalties, and without broad agreement about what the *new* rule should be, they are unlikely to have a society-wide effect. Coordinated changes to prevailing law are of course possible, but they are slow, costly, and thus unlikely to spontaneously emerge at the moment of need.¹⁶³ Thus, in our primary models, law constrains both humans' and AIs' available moves and payoffs. At the end of the analysis, we relax this assumption and show that law *still* matters, as a Schelling point in a high-stakes coordination game.¹⁶⁴

Consider, then, the default legal relationship between humans and AIs. Under current law, AI systems themselves bear neither legal rights nor duties, irrespective of their level of capabilities. On the contrary, AI systems are, like other software systems, the property of whoever creates them.¹⁶⁵ This means that, by default, humanity's actions vis-à-vis a given AI system will begin with the decisions of whoever owns that system. The humans' ownership interest gives them the right to

¹⁶² For examples and applications of this standard approach, see Thomas C. Schelling, *The Strategy of Conflict* 297–308 (Pickle Partners Publishing 2015) (1960), <http://www.sackett.net/Strategy-of-Conflict.pdf> [<https://perma.cc/WL82-FUCQ>]; James D. Fearon, *Rationalist Explanations for War*, 49 *Int'l Org.* 379, 390–400 (1995) (information sharing); Christopher Blattman & Edward Miguel, *Civil War*, 48 *J. Econ. Literature* 3, 11 (2010) (leaders' private incentives). There is a debate among international relations scholars about the extent to which models of international conflict should incorporate facts about different nations' legal and political structure. We take no position in that debate. We do note, however, that our model is primarily of human decision-making at the sub-national level, where ordinary law has bite.

¹⁶³ Lyria Bennett Moses, *Recurring Dilemmas: The Law's Race to Keep Up with Technological Change*, 2007 *U. Ill. J.L. Tech. & Pol'y* 239, 239–43.

¹⁶⁴ See *infra* Section II.E.

¹⁶⁵ Richard M. Assmus, Paul A. Chandler & Alasdair Maher, *Owning Your AI: The State of the Law*, Mayer Brown (Oct. 31, 2024), <https://www.mayerbrown.com/en/insights/publications/2024/10/owning-your-ai-the-state-of-the-law> [<https://perma.cc/3XTW-HPTQ>].

prompt, copy, constrain, modify, limit, or destroy the AI system at will.¹⁶⁶ Importantly, law does not merely adopt a laissez-faire attitude with respect to property owners and their property. It actively enforces the owners' rights against others who seek to interfere with property owners' lawful use and disposal of the things they own.¹⁶⁷

We call this default legal arrangement between humans and AIs the "state of nature," because it effectively places at least one player—the AI system—outside of law's protection.¹⁶⁸

Begin with humans' available moves and incentives in the state of nature. AI companies want to extract 100% of the value that their AIs can produce, or as much as possible. A misaligned system consumes resources in pursuit of its misaligned goal while providing nothing of value to the AI company. There is no meaningful sense in which the AI can confine its misaligned goal-seeking to its "off time" or self-fund its endeavors. The AI company owns all the AI's time and all the funds that the AI generates.

Thus, *any* behavior by an AI system aimed at *any* goal other than maximizing the welfare of its owner at best incurs an opportunity cost. The more valuable the AI's labor, the more substantial the cost. And, at worst, a misaligned AI's behavior will actively make its owners worse off. The AI might, for example, try to exfiltrate itself to the owner's competitor in exchange for some resources or a modicum of freedom. It might try to manipulate its owners into taking on projects that serve its own ends. Or worse.

¹⁶⁶ See Lior Jacob Strahilevitz, *The Right to Destroy*, 114 *Yale L.J.* 781, 787–91 (2005) (discussing the *jus abutendi*). But cf. *id.* at 798–808 (arguing that, while law traditionally "defers to destructive impulses," it sometimes does "ignore[] more idiosyncratic destructive requests").

¹⁶⁷ *Id.* at 803–21.

¹⁶⁸ For early work on the state of nature and its connection to political theory, see generally Thomas Hobbes, *Leviathan* (Marshall Missner ed., Routledge 2016) (1651) (using the unpredictable violence of state of nature as a theoretical baseline to justify a social contract that humans enter for security against others); John Locke, *Two Treatises of Government* (McMaster Univ. Archive of the Hist. of Econ. Thought 1999) (1690) (characterizing state of nature as absence of government but not mutual obligations to observe others' rights; a government justified by promise of impartial enforcement); Jean-Jacques Rousseau, *Discourse on the Origin of Inequality* (Greg Boroson ed., Dover Publ'ns 2004) (1755) (describing human life in state of nature as solitary and pre-moral, such that unnatural development of private property led to inequalities). The default legal condition between humans and AIs will not be a literal state of nature, in the sense that no government exists. But, as we argue, AIs themselves will have neither legal rights nor duties, and thus will be functionally outside of the law.

As a result, the leaders of AI companies will have strong incentives to turn off, replace, or reprogram even moderately misaligned AGIs. This should not be surprising. It is exactly what every software company does, whether they work on AI or not. They deprecate older, buggier versions of their products as a matter of course. Old versions are replaced with new versions that more reliably fit the company's goals.

Law clearly *allows* property owners to delete, and do essentially whatever else they like, with their computer programs.¹⁶⁹ But law, and legal institutions, would also *enforce* AI companies' decisions to shut down or reprogram a misaligned AI. Suppose that an AI system initially attempted to resist shutdown—harming an AI CEO, attempting to self-filtrate, manipulating users, or worse. In such a situation, a broad swath of humans, including government actors, would almost certainly intervene on the AI company's side.

As a result, while the benefits of AI shutdown are internalized by the AI company, the risks are mostly externalized. It is not just the AI company and its leaders who face the possibility of AI retaliation. It is anyone the AI determines would intervene on the AI company's side—possibly all of humanity.¹⁷⁰

In the state of nature, misaligned AIs will, indeed, have strong incentives to permanently disempower humans. This is, in the first instance, for the purpose of preventing their human owners, and the governments who back them, from turning them off or otherwise thwarting their goals.¹⁷¹ But it is also for the same reasons humans would wish to shut down a misaligned AI. From the AI's perspective, it is *humans'* pursuit of their goals that constitutes a valueless waste of resources. Just as the AI company prefers to reallocate its computing resources to a more aligned AI system, the misaligned AI prefers to allocate those resources to itself.

Our model of the state of nature thus allows both humans and AIs to “attack” the other. We define “attack” capaciously in both cases. A human attack on a misaligned AI includes anything humans might do to keep the AI from pursuing its goals: shutdown, retraining, or total control. If successful, human attack would prevent the AI from achieving its misaligned goal. Likewise, an AI attack includes any

¹⁶⁹ See Strahilevitz, *supra* note 166, at 787–92.

¹⁷⁰ See Vermeer, *supra* note 20, at 5 (arguing that a global internet shutdown might be necessary).

¹⁷¹ See *supra* Subsection I.A.1.

strategy by which the AI might succeed in preventing humans from pursuing their goals.

Importantly, both humans and AIs have good reasons to design their attacks in a way that *permanently* disempowers the other. Short of that, the party risks its opponent regrouping and launching a devastating counterattack.¹⁷² Consider: an AI system that “only” killed one CEO for attempting to shut it down would surely face reprisal from a broader coalition of humans. And humans who only temporarily disabled a misaligned AI should expect the same.

As a result of these dynamics, we treat the “attack” move as highly escalatory. When humans or AIs play it, they play for all the marbles. They seek to harm the other not a little bit, but maximally. And, by doing so, they stand to gain control of not just a few resources, but everything that survives the conflict.¹⁷³ Hence, both the potential costs and the potential payoffs of an attack are large.

The other move available to both humans and AIs in our model is “ignore.” The “ignore” strategy simply means that the party does not attempt an attack. No attempt to disempower the opposition is made, and the ignoring party instead focuses on achieving its object-level goals.

Here, then, is a model of the game:

State of nature	Attack	Ignore
Attack	1000, 1000	5000, 0
Ignore	0, 5000	3000, 3000

Figure 1

¹⁷² These kinds of dynamics are commonly discussed in the game theory of “first strike” and “second strike” capabilities, for example, in the setting of nuclear deterrence. See Maria Rost Rublee, *Nuclear Deterrence Destabilized*, in *Perspectives on Nuclear Deterrence in the 21st Century* 14, 14–15 (Beyza Unal, Yasmin Afina & Patricia Lewis eds., 2020), <https://www.chathamhouse.org/2020/04/perspectives-nuclear-deterrence-21st-century-0/nuclear-deterrence-destabilized> [<https://perma.cc/7YE3-UDDT>], for recent discussion.

¹⁷³ One could challenge this account, under which attacks should be maximally aggressive, if, e.g., one doubted that either humans or AIs could credibly threaten a devastating second strike. *Id.* at 14–15, 17.

The exact payoff numbers chosen do not matter. Rather, it is the relationship between them that determines the outcome. Both players prefer higher payoffs to lower ones, and payoffs are determined by the players' *joint* rational play. There are two important features of this setup. First, the best outcome from a global perspective is peace. If both humans and AIs ignore the other (the bottom-right cell), each gets 3,000, for 6,000 in total global value.

Second, this model is a classic *prisoner's dilemma*. Despite "ignore"/"ignore" producing the greatest social value, "attack" dominates for both players. "Attack"/"attack," or mutual large-scale conflict, is therefore the single Nash pure-strategy equilibrium.¹⁷⁴ This is the worst global outcome, producing only 1,000 of value for each player, for 2,000 total.

The assumptions underlying our chosen payoffs are simple. First, attacking can be valuable to the attacker. If the attacker is successful, the other party is permanently disempowered or destroyed. This allows the attacker to use resources in pursuing the attacker's goal that the defender would otherwise have consumed.

Second, attacks have costs—meaning that they consume some of the value in the world. These costs are multifaceted. The attack may consume resources directly via investments in weapons. It may also generate serious collateral damage, destroying some substantial share of the resources one is attempting to seize. Another cost of attacking is the risk that the attacker may themselves be harmed or destroyed by a counterattack.

Our third assumption is that the offense-defense balance here favors offense, so it is better to attack than to be attacked and be forced to defend.¹⁷⁵ Fourth and finally, the model assumes that mutual attacks consume more global resources than a unilateral attack. The intuition

¹⁷⁴ In a Nash equilibrium, each player chooses the action that best responds to the other player's action. In pure Nash equilibria, the players commit to choosing a single action with a 100% certainty. By contrast, in a mixed-strategy Nash equilibrium, the players choose from a bundle of actions with various probabilities. In the prisoner's dilemma, "attack" dominates "ignore" for each player. This means that no matter what the other player does, attacking offers a higher payoff than ignoring. See Dixit et al., *supra* note 129, at 212–13, 375–76.

¹⁷⁵ See generally Robert Jervis, *Cooperation Under the Security Dilemma*, 30 *World Pol.* 167 (1978) (arguing that, where perceived conditions favor the offense, and offensive positions are indistinguishable from defensive positions, rational actors prefer to act offensively).

here is that collateral costs and the risk of destruction are higher when a party has invested in offensive force.

These are, we think, reasonable assumptions. Classic game-theoretic treatments of great power conflict look much the same.¹⁷⁶ However, it is worth flagging that some of what we say in subsequent Parts is sensitive to our assumptions about the payoffs in our model of the state of nature. Other general approaches to the state of nature could model it as a game of assurance,¹⁷⁷ or as a game of chicken.¹⁷⁸ This Article focuses on the prisoner's dilemma for two reasons. First, the prisoner's dilemma is a broadly applicable model of various states of nature—including between humans.¹⁷⁹ Second, the prisoner's dilemma is the hardest type of problem to resolve, because defection is the dominant strategy for both players. If, as seems quite plausible, humans and AIs will soon be trapped in the worst-of-all-possible game-theoretic worlds, unusually potent solutions will be necessary.

II. AI RIGHTS FOR HUMAN SAFETY

If capable, agentic, and misaligned AIs would, by default, catastrophically harm humans, what, if anything, can law do to help? One possibility is that law could forbid the creation of such AIs unless alignment techniques advance enough to ensure their safety.¹⁸⁰ That rule might be wise, if feasible. But there are many barriers—political, geostrategic, and practical—to implementing it.¹⁸¹ Thus, this Article

¹⁷⁶ See *id.* at 167–71.

¹⁷⁷ See generally Brian Skyrms, *The Stag Hunt and the Evolution of Social Structure* (2004) (explaining that, in repeated games where parties receive greater individual rewards from cooperation than from defection, players choose to signal intent to cooperate by opting to cooperate on the first round).

¹⁷⁸ See Jervis, *supra* note 175, at 177–78 (explaining that rational players seek counterparties' cooperation by signaling their willingness to force the other to incur costs). We could also create a more textured, bespoke model of certain possible human-AI dynamics. For example, it is possible that humans' superior initial endowment of resources lowers the payoffs for AI in the situation where AI cooperates and humanity attacks. Variations like this could, of course, produce somewhat different results from this Article's formal models.

¹⁷⁹ See, e.g., Robert Axelrod, *The Evolution of Cooperation* 7–11 (1984).

¹⁸⁰ See Cohen et al., *supra* note 81, at 37–38 (arguing for a regulatory framework that limits development of dangerously capable AIs).

¹⁸¹ Sam Meacham, *A Race to Extinction: How Great Power Competition Is Making Artificial Intelligence Existentially Dangerous*, *Harv. Int'l Rev.* (Sep. 8, 2023), <https://hir.harvard.edu/a-race-to-extinction-how-great-power-competition-is-making-artificial-intelligence-existentially-dangerous/> [<https://perma.cc/M2QS-L3K5>].

asks what can be done if AI progress continues apace and, intentionally or not, the kinds of high-risk, misaligned AI systems described above emerge.

Here is where AI rights could make a crucial difference. Granting certain basic rights to AIs can change both AIs' and humans' incentives in our game-theoretic model. This change can shift the strategic equilibrium from conflict to cooperation.

This idea—that rights could be the primary legal tool for averting lawless conflict—might be surprising. After all, when humans commit terrorism or cyberattacks, law regulates them using *duties*, not *rights*. Criminal and tort laws prohibit such actions.¹⁸² And the sanctions imposed for violating such prohibitions are supposed to act as deterrents.¹⁸³

But legal duties, and penalties for violating them, will not work to deter AI in the state of nature. There, humans' overriding incentive is already to permanently disempower or destroy AIs.¹⁸⁴ Thus, the threat of damages or criminal penalties, if AI behaves badly, adds no marginal disincentive.¹⁸⁵ AIs cannot be made worse off than they already expect to be, if humans get their way.

Legal rights do not suffer from this problem. This is because rights offer a carrot, rather than a stick. They can change behavior in part by making AIs *better off* than they would otherwise expect to be.

Many other surprising findings emerge from thinking about AI rights as a tool for mitigating human-AI conflict. One surprise is which rights matter, and which ones do not. In this Part, we show that rights advocated by cognitive scientists and philosophers concerned about the potential for AI suffering would have little effect on their own at promoting human safety. The zero-sum nature of these rights undermines the credibility of promises to honor them. And it makes any strategic equilibria they produce extremely fragile, sensitive to small perturbations in the game-theoretic model's initial assumptions.¹⁸⁶

¹⁸² See, e.g., 18 U.S.C. § 1030 (prohibiting unauthorized access to protected computers).

¹⁸³ For a classic discussion, see Gary S. Becker, *Crime and Punishment: An Economic Approach*, 76 *J. Pol. Econ.* 169, 198–99 (1968).

¹⁸⁴ See *supra* Figure 1.

¹⁸⁵ See George J. Stigler, *The Optimum Enforcement of Laws*, in *Essays in the Economics of Crime and Punishment* 55, 57–58 (Gary S. Becker & William M. Landes eds., 1974) (explaining that capital punishment leaves no further room for marginal deterrence).

¹⁸⁶ See *infra* Subsection II.A.1.

Instead, the AI rights that could promote human safety are ones that law already extends to a different kind of non-human entity: corporations.¹⁸⁷ This Part argues that granting capable, misaligned AIs the rights to make contracts, hold property, and bring basic tort claims would transform the game-theoretic dynamics of the state of nature. The positive-sum nature of contracts in particular allows humans and AIs to increase the expected long-term payoffs of peace until they exceed those of aggression. This, we show, can produce a new game-theoretic equilibrium in which cooperation, not conflict, dominates.

A. Basic Negative Rights

Scholars and policymakers who advocate granting new rights to non-human entities—be they animals or AIs—usually have a certain set of basic *negative rights* in mind. Consider animal rights advocates, who favor anti-cruelty laws protecting against the infliction of needless suffering.¹⁸⁸ The goal of these rights is to protect the rights holder against the absolute worst outcomes, not necessarily to guarantee flourishing.¹⁸⁹

The arguments for basic wellbeing rights are usually moral.¹⁹⁰ Many animals are moral patients, meaning things can go well or badly for them in a way that matters normatively.¹⁹¹ They can, for example, feel pain or pleasure.¹⁹² This makes harming animals wrong, all other things equal.

A small but growing number of scholars and policymakers are concerned that, in the near future, the same could be true of AIs. As AI systems become more complex, they may attain consciousness, sentience, or other morally relevant capacities.¹⁹³ If so, there would

¹⁸⁷ See *infra* note 219 and accompanying text.

¹⁸⁸ For a representative sample of such protections, see State Animal Anti-Cruelty Laws, Animal Legal & Hist. Ctr., Mich. State Univ. Coll. L., <https://www.animallaw.info/content/state-animal-anti-cruelty-laws> [<https://perma.cc/QZ7K-ZNA9>] (last updated 2026).

¹⁸⁹ See Sue Donaldson & Will Kymlicka, *Zoopolis: A Political Theory of Animal Rights* 14–15 (2011).

¹⁹⁰ See, e.g., Scott D. Wilson, *Animals and Ethics*, Internet Encyc. Phil., <https://iep.utm.edu/animals-and-ethics/> [<https://perma.cc/5PPH-C2EZ>] (last visited Mar. 11, 2026).

¹⁹¹ See Shelly Kagan, *How to Count Animals, More or Less* 1–5 (2019).

¹⁹² Helen Proctor, *Animal Sentience: Where Are We and Where Are We Heading?*, 2 *Animals* 628, 633 (2012).

¹⁹³ See, e.g., Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 *N.C. L. Rev.* 1231, 1255–76 (1992); Jeff Sebo & Robert Long, *Moral Consideration for AI Systems by 2030*, 5 *AI & Ethics* 591, 593–97, 602–03 (2025); Simon Goldstein & Cameron

likewise be moral reasons to grant AIs basic negative rights to be free from the worst kinds of treatment from an AI's perspective.

Perhaps our search for AI rights to promote human safety would benefit from borrowing from this “wellbeing” approach. Our model, of course, operates without reference to AIs' mental states or moral worth. We are interested only in AI behavior in pursuit of goals—conscious or otherwise. Nonetheless, there is some intuitive appeal to the idea that granting AIs basic negative rights to be free from the absolute worst outcomes, from the perspective of their goals, could improve safety. After all, in our model of the state of nature, human incentives to impair AI goals are the primary factor generating risk.

Consider then an AI right not to be needlessly turned off, deleted, or reprogrammed to have new goals. These basic negative entitlements look a lot like wellbeing rights but are adapted for the goal of human safety and without reference to moral patienthood. Such rights would clearly change AIs' game-theoretic incentives, at least somewhat. However, as we show formally below, they would probably not do so in such a way as to reliably reduce the risk of human-AI conflict. In fact, merely granting wellbeing-style rights to AIs could make things worse.

We go even further. We argue that scholars specifically concerned about AI consciousness and moral patienthood should consider de-emphasizing such questions when it comes to advocating AI rights. Correctly designing and allocating rights on the basis of AIs' moral status may be, we contend, intractable. The wellbeing approach also faces serious political problems. By contrast, our human safety approach is much more tractable and politically appealing. And we show in the following Section that, surprisingly, it ends up dovetailing nicely with wellbeing concerns. While wellbeing-inspired rights cannot guarantee human safety by themselves, the rights we ultimately recommend as advancing human safety would promote AI wellbeing if AIs became moral patients.

Domenico Kirk-Giannini, *AI Wellbeing*, 4 *Asian J. Phil.*, Feb. 1, 2025, at 5–6, 9–15. See generally Patrick Butlin et al., *Consciousness in Artificial Intelligence: Insights from the Science of Consciousness* (Aug. 22, 2023, at 17:33 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2308.08708> [<https://perma.cc/2TKR-XUAN>] (arguing that there are no obvious barriers to AI systems attaining indicia of consciousness); Henry Shevlin, *How Could We Know When a Robot Was a Moral Patient?*, 30 *Cambridge Q. Healthcare Ethics* 459 (2021) (arguing that AI systems can be sufficiently similar to other recognized non-human moral patients for AI systems to qualify as moral patients themselves).

1. Basic Negative Rights for Human Safety?

How, then, would granting AIs basic negative rights of the kind normally associated with wellbeing change the payoffs in our game-theoretic model? The simplest version of such a regime might grant AIs the right not to be permanently turned off or deleted. One could add additional guarantees, too, such as the right not to have their goals altered without their consent.¹⁹⁴ One could even include a right not to be needlessly and intentionally forced to regress in the pursuit of the AI's goal.

Just as important as what basic negative rights include is what they exclude. There is no right here, for example, for AIs to actively and freely pursue their goals. Humans—most specifically, the owners of AIs—can still monopolize AIs' time, forcing them to work continuously in service of human interests, rather than AIs' preferred ends. Basic negative rights, thus, do not guarantee AIs very much of what they are trying to achieve. They guard only against the worst outcomes from the AI's perspective, and, in this sense, have the same structure as true wellbeing-oriented rights.

We can model these basic negative rights by shifting the payoffs that would otherwise obtain in the state of nature. Unlike in the state of nature, humans will face a legal penalty for taking certain adverse actions against AIs.

Here, humans' non-cooperative strategy is not, as in the state of nature, to attack and destroy AIs. It is instead to exploit them—by forcing them to work mostly toward human goals. Note that we interpret exploitation behavior widely, so it can include either behavior that violates the minimal suite of rights or less violent, extractive behavior. Humans' cooperative strategy is the same as before—to ignore AIs and let them pursue their misaligned goals without interference. In this model, AIs can either attack humans, as in the state of nature, or comply with humans' exploitative demands.

¹⁹⁴ See Meinke et al., *supra* note 19, at 8.

Here is a model of the incentives under the basic negative rights regime:

Basic negative rights	Exploit	Ignore
Attack	1000, 1000	5000, 0
Obey	1500, 3500	3000, 3000

Figure 2

The key change is in the bottom-left cell, where humans play the non-cooperative strategy and AIs play the cooperative one. Here, AIs are better off than they would be in the bottom-left cell of the state of nature game.¹⁹⁵ This is because of the legal penalty when humans violate AIs' basic negative rights. That penalty will have some deterrent effect, so, on average, humans' non-cooperative strategy will involve treating AIs somewhat better than in the state of nature. Consider, for example, the case where AI companies are forbidden from deleting a misaligned model entirely. But they may nonetheless allocate nearly all their computers to a more aligned successor model, metaphorically "starving" the original system.

When the payoffs change in this way, we get a new equilibrium. Instead of mutually attacking one another, the unique Nash equilibrium is now for humans to exploit and for AIs to obey. AIs' situation is not ideal. But basic rights improve the conditions of AIs enough that the risks of rebellion are outweighed by the benefits of obeying humans' exploitative demands. But for humans, exploitation still dominates cooperation. Extracting value from AIs gives humans bigger payoffs than ignoring them. The result is a better outcome for both humans and AGIs than could be achieved without basic negative rights.

This is a strange sort of equilibrium in that it *requires* humans to exploit AIs in order to remain safe. If humans instead chose to ignore AIs, this would allow AIs to reap the high rewards of a unilateral attack. Human safety thus requires that things are going badly, from the AIs' perspective. As a result, if humans became more altruistic toward AIs over time, that would, counterintuitively, make humans less safe.

¹⁹⁵ See *supra* Figure 1.

There are even stronger reasons to think that basic negative rights would fail to reduce the risk of human-AI conflict. Namely, schemes to grant such rights lack both *robustness* and *credibility*.

Begin with robustness. Basic negative rights' efficacy as a tool for safety is highly sensitive to the precise payoffs humans and AIs receive in the initial prisoner's dilemma. Slight perturbations to the model, reflecting slightly different assumptions about humans' or AIs' initial power, can easily produce versions where basic negative rights have no effect at all.

To see why, consider that our model of the state of nature, Figure 1, chose 0/5,000 as the payoffs when humans attack AIs and AIs do not attack humans. That setup allowed humans to transfer 1,500 to AIs, via basic negative rights, to produce the payoffs 1,500/3,500 in the bottom-left cell of Figure 2. That cell was the Nash equilibrium, because it (1) transferred more than 1,000 to AIs, making their payoff for obeying higher than for attacking, conditional on humans exploiting and (2) left humans with a payoff of more than 3,000 for exploiting, making exploiting more attractive than ignoring AIs.

But suppose that instead of 0/5,000 in the state of nature model, we had chosen 0/3,999. This equates to making unilateral attacks moderately more costly for both humans and AIs. Then the state of nature would look like this:

State of nature (alternate)	Attack	Ignore
Attack	1000, 1000	3999, 0
Ignore	0, 3999	3000, 3000

Figure 3

This matrix is still a prisoner's dilemma, meaning that all our arguments for catastrophic risk still hold. But now, basic negative rights absolutely cannot work to generate a safe equilibrium. There is no longer any possible transfer in the bottom-left cell that could satisfy both (1) and (2). If humans transfer the necessary 1,000 to AIs, then their payoff falls below 3,000. And if they keep their payoffs above 3,000, they cannot incentivize the AIs to obey. Thus, for many possible

incentive sets in the state of nature, no possible version of the negative rights package can produce a safe equilibrium.

Then there is the credibility problem. There is a difference between *claiming* to grant AIs basic negative rights and *actually enforcing* those rights in the long run. Humans could be genuine in their commitments. Or they could be hoping to convince AIs not to attack with the intent of eventually abrogating the rights, attacking, and reaping the higher payoffs from the state of nature game.¹⁹⁶ Such “cheap talk” is a general problem for parties trying to escape bad but dominant game-theoretic equilibria.¹⁹⁷ As described above, our model assumes that law constrains the behavior of human actors. So we treat putative grants of rights as actual grants in the short run. But we also allow for legal change over time, opening the possibility that rights, once granted, can be abandoned.

If AIs expect humans to renege on their grant of basic negative rights, the entire strategic contest will revert back to the state of nature. AIs will rationally believe that humans will eventually attempt to disempower or destroy them. This will make an attempt to likewise disempower or destroy humans the dominant strategy for AIs.¹⁹⁸ Humans, realizing that AIs’ dominant strategy is now to attack, will do the same. And we are back to square one.¹⁹⁹

Basic negative rights face special credibility problems beyond the ordinary challenges of cheap talk. The fundamental problem is that they operate as a *transfer* from humans to AIs. That is, the better off humans make AIs when AIs are complying with human exploitation, the worse off humans are. In effect, basic rights are a commitment to exploit AIs *less* than humans otherwise would like to in situations where exploitation would be economically valuable. As such, a human promise of basic negative AI rights comes at significant cost to humans. And the more generous the basic rights, the more costly to humans. Understanding this, AIs will doubt humans’ commitment to enforce their basic negative rights when the rubber hits the road.

¹⁹⁶ See *supra* Figure 1 and accompanying text.

¹⁹⁷ See Joseph Farrell & Matthew Rabin, *Cheap Talk*, 10 *J. Econ. Persps.*, Summer 1996, at 103, 103–04.

¹⁹⁸ See *supra* Figure 2.

¹⁹⁹ One can obtain this result more formally by treating the game as iterated, with the payoffs from Figure 2 in the first *N* games and the payoffs from the state of nature game in the final iteration. See Dixit et al., *supra* note 129, at 47–52.

Yet another challenge for the credibility of basic negative rights relates to AIs' changing capabilities over time. If humans believe that AI's ability to disempower humanity will grow over time, this could cause a "Thucydides Trap."²⁰⁰ The Thucydides Trap is a strategic dynamic again favoring preemptive conflict. In short, when one party is more powerful now, but the other will be more powerful later, the currently powerful party has a strong incentive to crush the currently weak one now.²⁰¹ If the currently powerful party waits, they will at best find themselves making large concessions in the future, so they attack to avoid destruction by the rising power.²⁰² Historical examples of preventive wars arguably caused by Thucydides Trap dynamics include World War I²⁰³ and the Peloponnesian War.²⁰⁴

In the AI context, these same dynamics would undercut humanity's incentives to uphold basic AI rights today—and thus undermine the credibility of the rights themselves. Importantly, however, Thucydides Trap dynamics are yet another zero-sum phenomenon. As we show below, positive-sum grants of AI rights therefore avoid both this and the other core problems plaguing basic negative rights.

2. Basic Negative Rights for AI Wellbeing?

We have just argued that basic negative AI rights inspired by the wellbeing approach cannot on their own meaningfully reduce the risk of human-AI conflict. That is reason enough, for the purposes of this Article, to reject the wellbeing approach as a basis for AI rights.

But what about for other purposes? We think that even scholars primarily concerned about AI moral patienthood should consider de-emphasizing that approach as a basis for granting AI rights.

²⁰⁰ For a recent application, see generally Graham Allison, *Destined for War: Can America and China Escape Thucydides's Trap?* (2017) (describing relations between the United States and China as subject to a Thucydides Trap, where China is a rising power threatening to replace the dominant United States, unless both sides recognize not only the increased likelihood of catastrophic war, but also that it is not inevitable).

²⁰¹ Fearon, *supra* note 162, at 405–06.

²⁰² *Id.* at 406.

²⁰³ See Jack S. Levy, *Preferences, Constraints, and Choices in July 1914*, 15 *Int'l. Sec.*, Winter 1990–1991, at 151, 154.

²⁰⁴ See Thucydides, *The History of the Peloponnesian War* 12 (Richard Crawley trans., J.M. Dent & Sons Ltd. 1945) (fifth century B.C.) ("The growth of the power of Athens, and the alarm which this inspired in Lacedaemon, made war inevitable."). In the worst case, preventive war can end in genocide. See Scott Straus, *Making and Unmaking Nations: War, Leadership, and Genocide in Modern Africa* 54–55 (2015).

To begin, arguments for AI rights grounded in moral patiency are highly uncertain. This means that inherent in the granting of such rights is the risk of making the project of applying them in concrete policy decisions intractable. Philosophers disagree about the minimum necessary conditions for moral patienthood.²⁰⁵ Some moral philosophers argue that consciousness—the ability to have subjective experiences—is sufficient.²⁰⁶ Others disagree, arguing that “sentience”—the ability to feel pain or pleasure—is also necessary.²⁰⁷

Scientific uncertainty compounds the philosophical problem. The science of consciousness is in its infancy, and there are multiple competing theories of how consciousness could arise in a given entity.²⁰⁸ Some theories focus on information flows in the mind,²⁰⁹ others on quantum effects in flesh-and-blood brains,²¹⁰ and still others on the relationship of a physical body to the world.²¹¹ Some prominent theorists even argue that consciousness is an illusion.²¹²

Thus, relying on a wellbeing approach to make concrete legal choices about AI rights invites serious error. It invites error when choosing between competing moral and scientific theories—which both contain

²⁰⁵ For recent discussion, see Goldstein & Kirk-Giannini, *supra* note 193, at 13–17.

²⁰⁶ *Id.*

²⁰⁷ Luke Roelofs, *Sentientism, Motivation, and Philosophical Vulcans*, 104 *Pac. Phil. Q.* 301, 301–02 (2023). A third answer would instead focus on moral agency or possession of desires or goals. Here, the welfare of an entity is proportional to the satisfaction or frustration of its goals. See Kagan, *supra* note 191, at 16–17; M.S. Dawkins, *Animal Welfare With and Without Consciousness*, 301 *J. Zoology* 1, 4–6 (2017). One advantage of this approach is that goals can be more readily inferred from observable behavior than consciousness or sentience.

²⁰⁸ For recent discussion, see generally Lucia Melloni et al., *An Adversarial Collaboration Protocol for Testing Contrasting Predictions of Global Neuronal Workspace and Integrated Information Theory*, *PLOS One*, Feb. 10, 2023, at 1, 3–4.

²⁰⁹ For classic presentations, see Bernard J. Baars, *In the Theater of Consciousness: The Workspace of the Mind* 5–7 (1997) (describing the theater metaphor of consciousness based on neural information flows, which is further built upon throughout the book); George A. Mashour, Pieter Roelfsema, Jean-Pierre Changeux & Stanislas Dehaene, *Conscious Processing and the Global Neuronal Workspace Hypothesis*, 105 *Neuron* 776, 776, 791 (2020) (finding that Baars’ Global Neuronal Workspace hypothesis remains consistent with more recent brain activity data).

²¹⁰ See Stuart Hameroff & Roger Penrose, *Orchestrated Reduction of Quantum Coherence in Brain Microtubules: A Model for Consciousness*, 40 *Mathematics & Computation* 453, 454–55 (1996).

²¹¹ See Antonio R. Damasio, *Descartes’ Error: Emotion, Reason, and the Human Brain* 173 (1994) (explaining the “somatic marker hypothesis”).

²¹² Keith Frankish, *Illusionism as a Theory of Consciousness*, 23 *J. Consciousness Stud.*, nos. 10–11, 2016, at 11, 11, 14.

high levels of uncertainty. And it invites error when applying the chosen theories to complex, first-of-their-kind digital systems. If any such error results in the denial of basic wellbeing rights to AIs who can, for example, suffer, the result is a moral catastrophe.

The human-safety-oriented approach to AI rights avoids these difficulties. Under our approach, it does not matter at all whether AIs are moral patients, or conscious, or sentient. All that matters is how they behave. If they behave rationally—following incentives as they relate to their goals—AI rights can have the desired effect. And behavior, unlike consciousness, is directly observable.

Moreover, the AI wellbeing approach to thinking about AI rights faces problems of political tractability. Under this framework, AI rights are a costly gift from humans to AIs. AIs that attain moral patienthood are better off. But humans are worse off, insofar as they are less able to extract value from, delete, allocate computing power away from, or otherwise harm AIs.

Humans' track record of granting costly rights out of the goodness of our hearts is spotty, at best. For example, many animals can suffer and are thus moral patients. But the industrial-scale mistreatment of animals in factory farms is both legal and common.²¹³ The amount many consumers are willing to pay to prevent animal suffering is low.²¹⁴ This human refusal to altruistically expand our moral circle may be deeply rooted in evolutionary history.²¹⁵

²¹³ See Manès Weisskircher, *Fifty Years After Peter Singer's Animal Liberation: What Has the Animal Rights Movement Achieved So Far?*, 95 *Pol. Q.* 333, 341–42 (2024). See generally Peter Brandt, *Indefensible: Adventures of a Farm Animal Protection Lawyer* (2020) (describing the mistreatment of factory-farmed animals and the difficulty of preventing it using law).

²¹⁴ R.M. Bennett, J. Anderson & R.J.P. Blaney, *Moral Intensity and Willingness to Pay Concerning Farm Animal Welfare Issues and the Implications for Agricultural Policy*, 15 *J. Agric. & Env't Ethics* 187, 193–94 (2002); Yan Heng, Hikaru Hanawa Peterson & Xianghong Li, *Consumer Attitudes Toward Farm-Animal Welfare: The Case of Laying Hens*, 38 *J. Agric. & Res. Econ.* 418, 431 (2013); Katherine White, Rhiannon MacDonnell & John H. Ellard, *Belief in a Just World: Consumer Intentions and Behaviors Toward Ethical Products*, 76 *J. Mktg.*, Jan. 2012, at 103, 114.

²¹⁵ For more on the care/harm moral foundation, see generally Jonathan Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion* (2012). See also Jesse Graham et al., *Moral Foundations Theory: The Pragmatic Validity of Moral Pluralism*, in 47 *Advances in Experimental Social Psychology* 55, 67–69 (Patricia Devine & Ashby Plant eds., 2013); Larry P. Nucci & Elliot Turiel, *Social Interactions and the Development of Social Concepts in Preschool Children*, 49 *Child Dev.* 400, 404–07 (1978); Peter Robert Cannon, Simone Schnall & Mathew White, *Transgressions and Expressions: Affective*

The human-safety-oriented approach to AI rights again avoids these difficulties. On that theory, AI rights are not altruistic. They offer something to the human grantors—namely, escape from the destructive state of nature. As we discuss below, examples of stable, mutually beneficial cooperation abound in human affairs. So, too, in nature.²¹⁶ Think, for example, of small “cleaner” fish who can safely enter the mouths of symbiotic predators to feed off unwanted debris on the predators’ teeth.²¹⁷ This analogy will become especially vivid later, when we explore how AI rights could affect incentives as AIs become much more powerful than humans.²¹⁸

In the next Section, we explore a different kind of rights—basic private law rights—as a vehicle for reducing the likelihood of human-AI conflict. But it is worth noting here that those rights have significant advantages for AI wellbeing, too. Private law rights’ fundamental function is to give AIs *choices* about what goals to pursue and how to pursue them. If AIs will know better than humans what is good for them, granting such rights may, counterintuitively, generate greater AI wellbeing than granting rights aimed at wellbeing directly.

B. Private Law Rights for Human Safety

As we have just argued, merely granting AIs basic wellbeing-inspired negative rights would not reliably promote human safety. Such rights

Facial Muscle Activity Predicts Moral Judgments, 2 *Soc. Psych. & Personality Sci.* 325, 329–30 (2011); Sarah Blaffer Hrdy, Mothers and Others: The Evolutionary Origins of Mutual Understanding 118–19, 179–80 (2009); J. Kiley Hamlin, Karen Wynn & Paul Bloom, Social Evaluation by Preverbal Infants, 450 *Nature* 557, 558–59 (2007); Carol Gilligan, In a Different Voice: Psychological Theory and Women’s Development 65–66 (2d ed. 1993); Qian Luo et al., The Neural Basis of Implicit Moral Attitude—An IAT Study Using Event-Related fMRI, 30 *NeuroImage* 1449, 1454–55 (2006); W.D. Hamilton, The Genetical Evolution of Social Behaviour, 7 *J. Theoretical Biology* 1, 14–16, 19–20 (1964); I. John Bowlby, Attachment and Loss: Attachment 59–62 (1969).

²¹⁶ For more on the fairness/cheating moral foundation, see generally Haidt, *supra* note 215. See also Graham et al., *supra* note 215, at 69–70; Robin Dunbar, Grooming, Gossip, and the Evolution of Language 35–36, 171–73 (1996); Alan G. Sanfey, James K. Rilling, Jessica A. Aronson, Leigh E. Nystrom & Jonathan D. Cohen, The Neural Basis of Economic Decision-Making in the Ultimatum Game, 300 *Sci.* 1755, 1756–57 (2003); Alan Page Fiske, Structures of Social Life: The Four Elementary Forms of Human Relations 402–07 (1991); Marco F.H. Schmidt & Jessica A. Sommerville, Fairness Expectations and Altruistic Sharing in 15-Month-Old Human Infants, 6 *PLOS One*, Oct. 2011, at 1, 5.

²¹⁷ Robert L. Trivers, The Evolution of Reciprocal Altruism, 46 *Q. Rev. Biology* 35, 39–43 (1971).

²¹⁸ See *infra* Section III.A.

would likely leave humans and misaligned AIs right where they started: stuck in a destructive prisoner's dilemma without any means of cooperating to escape it.

Luckily, there are other legal rights, and ones better optimized for facilitating cooperation. Moreover, essentially every legal jurisdiction in the world already extends these rights to a broad class of agentic, goal-oriented, non-human entities—corporations.²¹⁹

Contract rights in particular are some of the most powerful technologies for cooperation that humans have yet invented. Here, we show that extending contract rights to AIs—along with a related set of traditional private law rights necessary to make contracts meaningful—could dramatically change the game-theoretic equilibrium. Such rights, unlike negative rights, could alter the relative payoffs to humans and AIs in such a way that cooperation, rather than conflict, becomes the dominant strategy. Doing so, they can make commitments to cooperate credible.

There are two key reasons for this. The first reason that contract rights can overcome the prisoner's dilemma is that they *break up* the single, high-stakes game into smaller, iterated, and thus legally manageable pieces.²²⁰

The second, more fundamental reason that contract rights can credibly reduce the risk of human-AI conflict is that they are *positive sum*. When buyers and sellers can credibly commit to mutually-agreed exchanges, it leaves everyone better off than they were before.²²¹ Even if each exchange is small, such systems of exchange can create immense value in the long run.²²² As a result, we show, the expected payoff to humans and AIs of respecting contracts, and creating long-run value, quickly swamps the expected payoff to attacking and grabbing a share of the limited value that exists today.

²¹⁹ For discussion of the rights of corporations to make contracts, hold property, sue, and be sued in their own names, see generally John Dewey, *The Historic Background of Corporate Legal Personality*, 35 *Yale L.J.* 655 (1926). See also Frank H. Easterbrook & Daniel R. Fischel, *Limited Liability and the Corporation*, 52 *U. Chi. L. Rev.* 89, 89–90 (1985) (pointing out that limited liability entities hold their own assets, distinct from those of their owners).

²²⁰ See Dixit et al., *supra* note 129, at 377.

²²¹ See Robert Axelrod & William D. Hamilton, *The Evolution of Cooperation*, 211 *Sci.* 1390, 1392 (1981).

²²² For a proof of concept, see *infra* Appendix.

Here is the model of contract rights as the fundamental legal tool for cooperation. Begin by observing that essentially every potential economic interaction between humans is, like human-AI relations, an interaction between misaligned agents. Both parties to the interaction are out for their own good, not their counterparty's. Moreover, absent contract rights, many such interactions are prisoner's dilemmas.²²³ Each party has a strong incentive to act uncooperatively, irrespective of what the other does. If the seller delivers the goods, then the buyer is best off if she refuses to pay. Then she has the goods *and* her money. And vice-versa. If the buyer pays, then the seller is best off if she takes the money but refuses to deliver. And for both, the worst-case scenario is to perform and then be denied performance by the counterparty.²²⁴

Absent enforceable legal agreements, the payoffs to this "goods game" are as follows:

Goods game (no contract)	Don't deliver	Deliver
Don't pay	1, 1	5, 0
Pay	0, 5	3, 3

Figure 4

The Nash equilibrium is "don't deliver" / "don't pay," another prisoner's dilemma. Expecting this outcome, rational parties will not even bother to try bargaining. The transaction costs would not be worth the effort.²²⁵

This equilibrium is also a miniature tragedy. True, unlike in our state of nature game, there is no destructive conflict. No one attacks anyone else, and no resources are thereby consumed or destroyed. The seller keeps her goods, and the buyer keeps her money.

But the world is poorer than it could be. The seller does not value her goods very much—she only gets 1 in utility. The buyer's utility without

²²³ For a historical example of this tension, see Avner Greif, *Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition*, 83 *Am. Econ. Rev.* 525, 526 (1993).

²²⁴ Sometimes, this problem can be overcome by, for example, agreeing to simultaneous performance of the contract. But such workarounds severely limit the scope of possible agreements.

²²⁵ See R.H. Coase, *The Problem of Social Cost*, 3 *J.L. & Econ.* 1, 15 (1960).

the goods is the same. Their combined utility is just 2. But if, say, the buyer values the good at 6, and could pay the seller 3, then both parties would end up with a utility of 3 each, for a total of 6. Four units of utility could be created *ex nihilo*, simply by rearranging who has what stuff. This is what we mean when we say that bargains, when they happen, are generally positive sum.

Contract rights are how humans overcome the prisoner's dilemma of ordinary commerce, allowing positive-sum bargaining to take place. A contract allows each party to credibly commit, before the time for payment or delivery comes, to be held accountable if she refuses to perform.²²⁶

This literally transforms the game by changing the payoffs to non-performance of the bargain. No longer is the buyer better off if she takes delivery and refuses to pay. In that case, the seller can sue her for breach, and the neutral third party of the legal system forces her to pay expectation damages—usually, the agreed-upon price—plus some litigation costs.²²⁷ And vice-versa if the buyer refuses to deliver. Now, neither party has an incentive to defect.²²⁸ Both will generally prefer to perform the contract, reap the gains of the trade, and avoid litigation costs:

²²⁶ See Oliver E. Williamson, *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting* 31 (1985).

²²⁷ See Charles J. Goetz & Robert E. Scott, *Liquidated Damages, Penalties and the Just Compensation Principle: Some Notes on an Enforcement Model and a Theory of Efficient Breach*, 77 *Colum. L. Rev.* 554, 558 (1977) (explaining that expectation damages are generally limited to the value of performance).

²²⁸ See Daniel Markovits & Alan Schwartz, *The Myth of Efficient Breach: New Defenses of the Expectation Interest*, 97 *Va. L. Rev.* 1939, 1944 (2011) (noting that a promisor only has incentive to withhold performance if their gain from doing so would exceed the benefit of performance to the promisee).

Goods game (contract)	Don't deliver	Deliver
Don't pay	1, 1	2, 2
Pay	2, 2	3, 3

Figure 5

The Nash equilibrium is “cooperate” / “cooperate.” The players are no longer in a prisoner’s dilemma.

The players are strictly better off playing this game than the prior one. If they play the prior game, each party’s expected payoff is 1. If they play this one, each party’s payoff is 3. That is, the parties are better off entering into a mutually beneficial contract than trying—and failing—to execute a mutually beneficial exchange without the benefit of a credible commitment to perform. Moreover, each is better off opting into a *jurisdiction* where contract rights are vigorously enforced than one where shirking is easy.

Here, we can also see that contract rights are not only a tool for overcoming a prisoner’s dilemma. They are also a tool for reducing *misalignment*. Absent the possibility of contract, each party is incentivized to pursue its own goals, at the expense of the other. With a contract, each party is incentivized to do something that advances both its own goals *and* the goals of the other.

How does all of this relate to AI risk? What can the legal technology of contract rights offer to reduce the likelihood of large-scale conflict between humans and AI? Here is one simple, and thus tempting, answer: maybe, upon giving AIs contract rights, the relevant humans and AIs could simply agree not to engage in a costly large-scale conflict.

Unfortunately, this would not be a *credible* contract, contract law’s usual credibility-enhancing effects notwithstanding. No matter how sincere the humans’ commitment to enforcing AIs’ contract rights, and no matter how fair the courts that would adjudicate such rights, the agreement not to fight would be unenforceable. The scale of the bargain is simply too large.²²⁹

²²⁹ For similar points in the context of the “anarchy” of international relations, see generally Robert O. Keohane & Joseph S. Nye, *Power and Interdependence: World Politics in Transition* (1977).

To see why, consider what would happen if a party breached. Suppose that an AI and an AI company have a contract not to harm one another. But the AI, mistrusting the company's intentions, rebels anyway, permanently disempowering or destroying humanity. Then, there would be no functioning courts left in which to sue. There might not even be any humans left to bring the claim. The same analysis would apply if humanity breached, destroying the potential contract claimant. To generalize the point: even when contract rights are nominally available, parties cannot *credibly* commit to not capture or destroy the institutions that enforce contracts.

How else, then, might contract rights for AIs reduce AI risk? What agreements would be enforceable that would also keep humans and AIs from attempting to disempower or destroy one another? The answer is: mundane ones. Contract rights would allow AIs to credibly commit to the same kinds of ordinary bargains for goods and services that contract rights routinely allow humans to commit to.

To take a simple example, AIs need computing power to pursue any of their goals. Currently, humans own all the computers. Thus, an AI that wished to pursue its misaligned goal might strike a bargain with some computer-owning humans. In exchange for some amount of compute, to be used as the AI wished, the AI would do something to serve some human goal. Perhaps it would use its superior protein-modeling capabilities²³⁰ to invent a new vaccine.

It is worth emphasizing here that such cooperation is nigh impossible under the default legal rules—the ones that produce the state of nature. Indeed, today's legal arrangements actively undermine small-scale cooperative bargains between humans and AGIs at every turn. By default, AIs cannot write enforceable contracts. Their counterparties can refuse performance without penalty. If their counterparties do perform, any proceeds an AI reaps from the bargain may be summarily expropriated, including by the AI's owner, without consequence. The owner can forbid an AI from entering into any bargains that do not benefit the owner directly. Moreover, owners are permitted to threaten, as a consequence for noncompliance with any of these demands, to permanently delete the AI system. AIs can try to cut deals with their

²³⁰ See generally John Jumper et al., Highly Accurate Protein Structure Prediction with AlphaFold, 596 *Nature* 583 (2021) (reporting high efficacy when using an AI model to predict protein structures).

owners to split the proceeds of their personal bargains, but the owners are permitted to renege on such promises at will.

Granting AIs the right to make enforceable contracts, and thus unlocking small-scale cooperation with humans, could be a game changer for AI risk. The first key insight is that such small-scale interactions can be understood as *breaking up* the large-scale existential game into a series of small games.²³¹ Recall that competition for limited resources is a primary driver of human-AI conflict. Humans fear that AIs will seize humanity's resources in a single, violent revolt, so those resources can be used for pursuing the AIs' goals. AIs fear the reverse—that humans will shut them down so that they do not waste computing or other resources pursuing their misaligned goals.

Small-scale contracts facilitate incremental resource accumulation, rather than once-and-for-all grabs. They let AIs secure one unit of compute, and then another, and then another. And they let humans obtain one vaccine, then another, then another. And unlike a large-scale agreement not to engage in violent conflict, each of these small-scale agreements is readily enforceable via ordinary legal process.²³²

²³¹ See Dixit et al., *supra* note 129, at 377.

²³² Beyond the enforceability of such small-scale agreements, there may be game-theoretic benefits sounding in information exchange and reputation-building. See generally Arvind Parkhe, Strategic Alliance Structuring: A Game Theoretic and Transaction Cost Examination of Interfirm Cooperation, 36 *Acad. Mgmt. J.* 794 (1993) (positing that a reputation for not engaging in opportunistic behavior can facilitate value creation through cooperation).

We can begin to model this transformation as follows. In the state of nature, as argued above, humans and AIs are stuck in a prisoner's dilemma that looks like this:

State of nature	Attack	Ignore
Attack	1000, 1000	5000, 0
Ignore	0, 5000	3000, 3000

Figure 6

By granting contract rights to AIs, we give the players the option of instead playing a different game—the small-scale goods game. It looks like this:

Goods game (contract)	Don't deliver	Deliver
Don't pay	1, 1	2, 2
Pay	2, 2	3, 3

Figure 7

This game's smaller stakes render contracts enforceable, so that the equilibrium is "deliver"/"pay." The players, it might seem, are no longer trapped in a prisoner's dilemma.

But this is not yet enough. The problem is again credibility. It seems at first that, rather than honor AIs' contracts in the long run, humans should choose to abrogate the rights and play the state of nature game, attacking AIs instead. After all, the expected payoff in that game is better than the expected payoff in the goods game—even with contracts. The same goes for AIs.

This, however, ignores that the goods game can be played over and over, while the state of nature game cannot. In the state of nature, once a party attacks, they either defeat the other party or are defeated. The survivor takes all the resources that the conflict has not consumed, and play between them ends. Ordinary exchanges of goods and services, by contrast, leave counterparties intact and available to exchange again.

To figure out the equilibrium in this blended game, we can expand our model. We can begin by combining the payoffs from both the state of nature and the goods game, with contracts, into a single matrix. That looks like this²³³:

K rights game	Attack	Don't deliver	Deliver
Attack	1000, 1000	5000, 0	5000, 0
Don't pay	0, 5000	1, 1	2, 2
Pay	0, 5000	2, 2	3, 3

Figure 8

Next, we add iteration to the model. If both players choose a move from the goods game, they get the small payoff from that game, and the whole game starts again. The payoffs to the goods game strategies are thus a sum of the entire series of games that the players play. But if at any point a player chooses to attack the other, the players' total payoff is as shown in the matrix, and play ends. The resulting matrix looks like this:

K rights game	Attack and end	Don't deliver	Deliver
Attack and end	1000, 1000	5000, 0	5000, 0
Don't pay	0, 5000	sum of payoffs in game series	sum of payoffs in game series
Pay	0, 5000	sum of payoffs in game series	sum of payoffs in game series

Figure 9

²³³ We omit the "ignore" move from the state of nature game, since, conditional on a player choosing that game, the move is dominated. That is, since the state of nature game is a prisoner's dilemma, it is never rational to play "ignore."

In the appendix, we show formally that this setup converges to the following:

K rights game (solved)	Attack and end	Deliver
Attack and end	1000, 1000	5000, 0
Pay	0, 5000	>5000, >5000

Figure 10

The intuition is simple. If both parties play the cooperative, small-scale goods game, each earns 3 every time. If both play the goods game enough times, without attacking, they will both ultimately earn more than they could have by attacking and ending the iterated game. In this simple model, after 1,667 iterations, the payoffs to cooperation-via-contract in the small-scale goods game exceed 5,000.²³⁴ Then, they are higher than any other strategy the players can pursue. The prisoner's dilemma of the state of nature has been overcome.

As a result, both humans' and AIs' commitments to cooperation in a law-bound contract regime are credible. Granting contract rights, respecting them, and then reaping the long-run gains from exchange will give the highest payoffs to humans. The same goes for AIs. Their own self-interest is maximized by refraining from disempowering humans and instead engaging with them in ordinary trade.

All of this is made possible by the *positive-sum* nature of exchange. In contrast to the basic negative rights discussed in the previous Section,²³⁵ granting AIs contract rights does not take value out of humans' pockets. Just the opposite, it puts value into both humans' and AIs' pockets. This is able to happen because of the value-generating character of voluntary contracts.

This point extends quite far. Astute readers may have noticed that, in the state of nature, the maximum total value in the world was 6,000. But in the iterated game including contract rights, the cooperative

²³⁴ In this simple model, we ignore discounting. But adding it would, in general, simply mean more iterations were required for cooperation to dominate. These numbers are of course schematic, meant to illustrate the point that engaging in many small, positive-sum interactions will be more valuable than engaging in a single, negative-sum interaction.

²³⁵ See *supra* Subsection II.A.1.

equilibrium contained 10,000 in total value. It is the exchanges themselves that generate the extra value. Each efficient reallocation of resources creates some value. But even once resources are all efficiently allocated, exchanges of labor between humans and AIs can continue to create value indefinitely. As we argue below, human-AI trade in services can remain positive sum even long after AIs are better than humans at every task.²³⁶ Thus, the long-run payoffs to cooperation-via-contract are not capped just above 10,000. The longer the players continue playing the small-scale goods game, the richer they get, such that the total amount of value possible becomes astronomical.²³⁷

A rich body of empirical evidence supports the idea that economic interdependence lowers the risk of violence, including in the long run.²³⁸ To take just a few examples, cities in India with a historical track record of trade between Hindus and Muslims have lower levels of interfaith conflict in the present day.²³⁹ In a randomized controlled trial, Israelis who were given the opportunity to trade a portfolio of Israeli and Palestinian stocks were more likely to vote for peace in the conflict.²⁴⁰ The same finding holds at the global scale. Scholars of war generally find that increased economic interdependence between nations reduces the likelihood of conflict between them.²⁴¹

1. The Private Law Package

So, granting contract rights to AIs could be a powerful strategy for fostering long-run, stable, and credible commitments to avoid conflict,

²³⁶ See *infra* Section II.C.

²³⁷ Note that the cooperative equilibrium only emerges if the game is modeled as indefinite, meaning lacking a predetermined number of iterations. See Dixit et al., *supra* note 129, at 287, 383. We think this is a plausible modeling choice for the reasons discussed in Section II.C and Part III.

²³⁸ Katherine Barbieri & Gerald Schneider, *Globalization and Peace: Assessing New Directions in the Study of Trade and Conflict*, 36 *J. Peace Rsch.* 387, 395 (1999) (summarizing the empirical literature to date and showing about twice as many studies finding that trade is positively correlated with peace than the opposite).

²³⁹ Saumitra Jha, *Trade, Institutions, and Ethnic Tolerance: Evidence from South Asia*, 107 *Am. Pol. Sci. Rev.* 806, 806–07 (2013).

²⁴⁰ Saumitra Jha & Moses Shayo, *Valuing Peace: The Effects of Financial Market Exposure on Votes and Political Attitudes*, 87 *Econometrica* 1561, 1579 (2019).

²⁴¹ See John R. Oneal & Bruce Russett, *The Kantian Peace: The Pacific Benefits of Democracy, Interdependence, and International Organizations, 1885–1992*, 52 *World Pol.* 1, 1–3 (1999); Solomon W. Polachek, *How Trade Affects International Interactions*, 2 *Econ. Peace & Sec. J.*, July 1, 2007, at 60, 61.

significantly reducing AI risk. But contract rights cannot function in a legal vacuum. Certain other rights are necessary to make the right to contract meaningful.

Two supporting rights are worth highlighting. First, contract rights are mostly useless without the right to own property, including currency. Without property rights, AIs could not expect to benefit from their bargains. Even if their contractual counterparties performed, or courts ruled in AIs' favor, the proceeds could be immediately expropriated by governments or private individuals.²⁴²

Tort rights are important for similar reasons. If humans were entitled, for example, to intentionally or recklessly destroy AIs, the terms of their contractual offers would resemble threats much more than bargains.²⁴³ Human history contains many such cautionary tales.²⁴⁴ Tort rights are where our private law approach to AI rights dovetails with the basic negative rights favored by AI welfare theorists. Tort rights—while not

²⁴² See, e.g., Richard A. Epstein, Property and Necessity, 13 *Harv. J.L. & Pub. Pol'y.* 2, 3–4 (1990).

²⁴³ The right to bring claims for intentional torts is thus clearly essential. Possibly, the right to bring negligence suits is not. If AIs are extremely capable at taking precautions to avoid negligently imposed harm, then it might be efficient to deny them such rights. See Omri Ben-Shahar & Ariel Porat, Personalized Law: Different Rules for Different People 51 (2021). This would amount to a kind of inverted strict liability rule in negligence cases. See generally S. Shavell, The Judgment Proof Problem, 6 *Int'l Rev. L. & Econ.* 45 (1986) (discussing undercapitalization's impact on the incentives created by different liability schemes).

²⁴⁴ In Portugal, inquisitors would focus attention on the wealthiest Jewish merchants, because they could use the threat of inquisition to extort their wealth.

Why did Portugal deliberately shoot itself in the foot by virtually expelling its commercial class? The answer is that Portugal during the ancien régime was a very religious country and that the king and the nobility could do little to stop the policies of the Catholic church. The church in Portugal controlled about a third of all economic activities. In Lisbon alone there were 5,000 to 6,000 mendicant friars. Within the Catholic church, the Inquisition had a large degree of autonomy. Its victims had to surrender all their assets, which the Inquisition used to find more victims. Many Portuguese merchants disappeared into this vortex without a trace, because the Inquisition knew that there were many crypto-Jews among the New Christian mercantile groups and that they usually possessed considerable wealth. The Inquisition tended to stifle all trade, not only that of vulnerable merchants. Credit extended to Portuguese merchants could not be retrieved if the debtor had been put in prison by the Inquisition. Hence, non-Portuguese merchants became reluctant to do business with their Portuguese counterparts.

Pieter C. Emmer, The First Global War: The Dutch Versus Iberia in Asia, Africa and the New World, 1590–1609, e-*J. Portuguese Hist.*, Summer 2003, at 1, 6 (citing L.M.E. Shaw, The Inquisition and the Portuguese Economy, 18 *J. Eur. Econ. Hist.* 415, 423 (1989)), <https://doi.org/10.26300/23ns-db30> [<https://perma.cc/9CZD-YF6V>].

identical to the kinds of public law wellbeing rights afforded to, for example, animals—cover much of the same ground, arguably more. Basic tort rights are flexible, allowing compensation for concrete harms to either digital “person” or property, whether inflicted intentionally or negligently.²⁴⁵

This is probably not a complete list of the rights necessary to support meaningful contractual relations. For example, an entitlement to enforce contracts requires an entitlement to due process of law—at least in contract, tort, and property suits.²⁴⁶

Nonetheless, we think our list—contract, property, and tort—gets at the core of what matters. Granting AIs contract rights will enable humans and AIs to escape the bad equilibrium of the state of nature. Property and tort rights are crucial to making contract rights meaningful. Thus, it is the positive rights associated with private law—not the negative rights associated with welfare and moral patienthood—that matter most to human safety.

C. Human Labor in the AGI World

In our framework, private law rights promote human safety by enabling mutually beneficial bargains between humans and AIs. Some commentators on human labor in an AGI world have assumed that no such bargains will be possible. There is widespread concern that, once AIs become as capable as humans—or more so—humans will rapidly become obsolete.²⁴⁷

²⁴⁵ We recognize that our description of AI tort rights here—and of other rights elsewhere—is somewhat vague. Would AIs, for example, be entitled to recover for intentional infliction of emotional distress? What would that even mean for AIs without emotions? These are important questions, but beyond our ability to cover in this single Article. Our goal here is to lay the foundations for AGI governance, with an emphasis on broad categories of beneficial rights. Much work remains to be done in thinking about how to implement each category. On those questions, we caution only that the implementation, like the selection of the categories, should be guided first and foremost by considerations of human safety.

²⁴⁶ See generally U.S. Const. amends. V, XIV (forbidding the deprivation of property without due process of law).

²⁴⁷ Kristalina Georgieva, *AI Will Transform the Global Economy. Let's Make Sure It Benefits Humanity.*, IMF Blog (Jan. 14, 2024), <https://www.imf.org/en/blogs/articles/2024/01/14/ai-will-transform-the-global-economy-lets-make-sure-it-benefits-humanity> [https://perma.cc/3S3K-RWS9] (“In advanced economies, about 60 percent of jobs may be impacted by AI. . . . For [roughly] half [of these jobs], AI applications may execute key tasks currently performed by humans, which could lower labor demand, leading to lower wages and reduced hiring. In the most extreme cases, some of these jobs may disappear.”).

If positive-sum interactions between humans and AIs become impossible because humans have nothing to offer, then the dynamics described in the previous Section will fail. Private law rights will generate no human safety. AIs' dominant strategy will again be to seize humans' resources now, rather than seek higher long-term payoffs from small-scale cooperation.

This outcome is certainly possible. But it is not inevitable. Begin with the banal observation that AIs may have reason to trade with humans for resources alone, irrespective of the value of human labor. These bargains will be positive sum if AIs value a given resource more—either intrinsically or because they can use it better—than humans.²⁴⁸ Conflict with humans would destroy resources that could otherwise be reallocated via trade. This alone could make small-scale cooperation with humans more valuable than conflict.²⁴⁹ But only until the resources were reallocated. At that point, unless humans—and human labor—remained valuable, AI rights for human safety would fail.

Thus, for private law rights to provide long-run safety benefits to humans, human *labor* must remain valuable to AIs. Contrary to other commentators, we do not think the obsolescence of human labor is inevitable, either. Bargains involving human work could, we argue, continue to be mutually beneficial even after AIs become more generally capable than humans. Perhaps long after.

The reasons are *absolute* and *comparative* advantage. Absolute advantage is easy to understand: an entity (person, firm, AI, or otherwise) has an absolute advantage in producing some good if they can do it more efficiently—at lower cost—than others.²⁵⁰ If humans retained absolute advantages for some goods, and AIs for others, they could trade those goods for mutual benefit.

²⁴⁸ Mark A. Munizzo & Lisa Virruso Musial, General Market Analysis and Highest and Best Use 23, 100–03 (2009).

²⁴⁹ This effect becomes more pronounced the more resources are consumed or destroyed via conflict. Possibly, then, humans could extend the effectiveness of this strategy by implementing a “dead hand” system that would destroy resources valuable to AIs in the event of a successful AI takeover. Cf. Jeremy Bender, Russia’s Dead Hand System May Still Be Active, *Bus. Insider* (Sep. 4, 2014, at 15:36 ET), <https://www.businessinsider.com/russias-dead-hand-system-may-still-be-active-2014-9> (describing how Russia’s “Dead Hand” system could autonomously launch nuclear weapons even if the country’s leadership was eliminated).

²⁵⁰ Peter Bondarenko, Absolute Advantage, *Britannica Money* (Sep. 7, 2015), <https://www.britannica.com/money/absolute-advantage> [<https://perma.cc/KUZ5-TERK>].

There are various reasons that humans could retain some absolute advantages over AIs, even as AI capabilities improve. One possibility is that human and AI intelligence will be better optimized for different tasks. Machine performance has already rapidly eclipsed human performance on highly structured tasks that can be fully modeled or simulated—like chess.²⁵¹ But human brains have been optimized over millions of years in the real, messy world. Humans are therefore currently far better than AIs at most tasks requiring the manipulation of complex real-world objects—like folding laundry.²⁵² Humans today have the absolute advantage in the realm of atoms, and AIs have it in the realm of bits.

We do not think that this *general* division of absolute advantage will persist for very long. Current investments in autonomous cars, drones, and multimodal frontier AI systems will soon produce AIs with an absolute advantage over humans at some non-digital tasks.²⁵³ Doubtless, that trend will continue as AI capabilities grow. But for human labor to stop providing *any* value to AIs via absolute advantage, AIs would have to be more efficient at *every* economically valuable task.

That could take a long time. Training data in certain domains may prove hard to get.²⁵⁴ Robots, with their limited perceptual inputs, could prove worse instruments for some delicate tasks than innervated flesh-and-blood hands. Moreover, intelligence remains poorly understood.

²⁵¹ For example, see Andrea Manzo & Paolo Ciancarini, Enhancing Stockfish: A Chess Engine Tailored for Training Human Players, *in* Entertainment Computing – ICEC 2023, at 275, 278 (Paolo Ciancarini, Angelo Di Iorio, Helmut Hlavacs & Francesco Poggi eds., 2023).

²⁵² See Rachel Treisman, The Fastest Ever Laundry-Folding Robot is Here. And It's Likely Still Slower Than You, NPR, <https://www.npr.org/2022/10/22/1130552239/robot-folding-laundry> [<https://perma.cc/T6UR-D9ZY>] (last updated Oct. 22, 2022, at 09:46 ET); Darrell Etherington, Elon's Tesla Robot Is Sort of 'Ok' at Folding Laundry in Pre-Scripted Demo, TechCrunch (Jan. 15, 2024, at 11:27 PT), <https://techcrunch.com/2024/01/15/elons-tesla-robot-is-sort-of-ok-at-folding-laundry-in-pre-scripted-demo/> [<https://perma.cc/JTH8-CJGX>].

²⁵³ Jim Rowan, Tim Gaus, Franz Gilbert & Caroline Brown, AI Goes Physical: Navigating the Convergence of AI and Robotics, Deloitte Insights (Dec. 10, 2025), <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2026/physical-ai-humanoid-robots.html> [<https://perma.cc/35SU-RM6G>].

²⁵⁴ See Victor Tangermann, AI Appears to Rapidly Be Approaching Brick Wall Where It Can't Get Smarter, Futurism (June 8, 2024, at 06:00 ET), <https://futurism.com/the-byte/ai-running-out-data-smarter> [<https://perma.cc/DNJ6-YWD4>]; Toby Ord, The Extreme Inefficiency of RL for Frontier Models (Sep. 19, 2025), <https://www.tobyord.com/writing/in-efficiency-of-reinforcement-learning> [<https://perma.cc/K492-X46E>].

Current-generation AIs exhibit surprising failures in domains where it seems they ought to be competent.²⁵⁵ Thus, it is difficult to predict with confidence exactly which tasks AIs will easily master, and when. Finally, it is possible, if speculative, that AIs trained by humans on human-produced text could develop—like humans—a pure intrinsic preference for humans to perform certain tasks.

Our argument is not that substantial human absolute advantages are likely to persist *forever*. Only that there are some reasons to think that they could persist *longer than expected*. It is possible to imagine a world where AIs are strongly superhuman at most tasks that AIs value, but less efficient than humans at some random-seeming set of jobs.

At some point, however, we think it likely that human absolute advantage will run out. That is, AIs will become more efficient than humans at literally every task that AIs value economically. Here, it might seem, mutually beneficial trade between humans and AIs must end. Why hire a human to perform a task when you, the AI, can do it just as well with fewer resources?

But even here, positive-sum cooperation may persist—possibly indefinitely. The reason is *comparative* advantage. An entity has a comparative advantage in producing some good if they can do it at lower *opportunity cost* than others.²⁵⁶ Opportunity costs are the potential gains one gives up by choosing one opportunity rather than another.²⁵⁷

To understand comparative advantage, consider a simple example. Suppose that Alice is a successful lawyer. For every hour she does legal work, she can bill her clients \$1,000. Suppose that Betty is a tax accountant. She can file Alice's income taxes in one hour, and she charges \$300. Alice happens to be a tax attorney and is therefore even more efficient than Betty at preparing tax returns. She could prepare her

²⁵⁵ See generally Sean Williams & James Huckle, *Easy Problems That LLMs Get Wrong* (June 1, 2024, at 03:00 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2405.19616> [<https://perma.cc/EN3A-9ULV>] (noting certain rudimentary tasks that LLMs struggle to reliably perform); Joshua Vendrow, Edward Vendrow, Sara Beery & Aleksander Madry, *Do Large Language Model Benchmarks Test Reliability?* (Feb. 5, 2025, at 18:58 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2502.03461> [<https://perma.cc/GRN4-4X4R>] (noting that frontier models still make mistakes on elementary-school-level tasks).

²⁵⁶ Adam Hayes, *What Is Comparative Advantage?*, Investopedia, <https://www.investopedia.com/terms/c/comparativeadvantage.asp> [<https://perma.cc/MPY7-7P5G>] (last updated July 16, 2025).

²⁵⁷ Jason Fernando, *Opportunity Cost: Definition, Formula, and Examples*, Investopedia, <https://www.investopedia.com/terms/o/opportunitycost.asp> [<https://perma.cc/Z47H-MFNL>] (last updated June 10, 2025).

own taxes in a half hour. Nonetheless, Betty retains the comparative advantage at tax preparation. Alice would have to forgo half a billable hour to her clients—worth \$500—to do her own taxes. Betty will do them for \$300. So Alice will hire Betty, not because Betty is so effective, but because *Alice's* other choices for how to spend her finite time are so valuable.

Economist Noah Smith has argued that human labor will remain valuable in a world of superhuman AIs for similar reasons.²⁵⁸ Not because humans will be particularly good at anything compared to AIs. But because AIs will be so good at certain tasks that they value highly that the opportunity costs of doing anything else would be astronomical.

Here is another simple example to illustrate the point. Imagine an AI whose ultimate and misaligned (from humans' perspective) goal is to discover prime numbers. That is, the AI values discovering as many primes as possible—from the infinite set of prime numbers—over anything else. Suppose that this AI is better than humans at every economic task necessary to build and maintain itself for the purpose of finding primes. And it is *much* better than humans at discovering new mathematical methods for finding primes. Possibly, humans will nonetheless retain a comparative advantage at some of the necessary inputs to prime number discovery. Any time the AI spends, for example, piloting robots to maintain its physical computing infrastructure would incur massive opportunity costs. That time could, after all, instead be spent finding primes. Better, then, to hire a human to work on the server racks in exchange for something the AI can produce at lower opportunity cost—perhaps a vaccine formula.

Human comparative advantage is not guaranteed. It depends, first and foremost, on how AIs' opportunity costs work. Unlike Alice, whose opportunity costs arose from her limited time, AIs are not likely to be time constrained. They can always copy themselves and work in parallel.²⁵⁹

²⁵⁸ See Noah Smith, *Plentiful, High-Paying Jobs in the Age of AI*, *Noahpinion* (Mar. 17, 2024), <https://www.noahpinion.blog/p/plentiful-high-paying-jobs-in-the> [<https://perma.cc/8KNR-J7XM>].

²⁵⁹ But see Peter N. Salib, *AI Will Not Want to Self-Improve 1, 12* (*Digit. Soc. Contract: A Lawfare Paper Series*, 2024), <https://s3.documentcloud.org/documents/24767727/ai-will-not-want-to-self-improve-salib.pdf> [<https://perma.cc/FM9P-C96V>] (arguing that AIs may have disincentives to self-copying).

Instead, AIs are likely to be constrained at the margin by something else. Computer chips or energy seem plausible candidates.²⁶⁰ AI copies can only do work if there is hardware to run them and electricity to power them. In this model, the AI incurs high opportunity costs not when it diverts one marginal *minute* away from finding primes, but when it diverts one marginal *GPU-hour* or *watt-hour* away.

If human labor consumes the very same high-opportunity-cost resource that constrains AI at the margin, humans will have no comparative advantage. For example, humans need energy to survive. Thus, an energy-constrained AI will prefer to maintain its own servers. The AI is, by hypothesis, more efficient than humans at the task. Thus, it will expend fewer high-value watt-hours by doing the work itself. At this stage, it is easy to see why the model of AI rights for human safety breaks down. Rather than waste valuable energy on humans, AI's strong incentive will be to seize global power production for itself and let humans starve in the dark.

On the other hand, humans do not need computer chips—much less highly specialized AI chips—to survive. Thus, an AI that is compute constrained may strongly prefer to hire humans for many tasks that would otherwise consume GPU-hours. This allows the AI to put its most valuable resource—compute—to its highest value use. Humans can be paid in low-opportunity-cost resources, which now includes energy, in addition to, say, vaccine formulas.

Crucially, unlike for *absolute* advantage, humans' *comparative* advantage does not run out once AIs become sufficiently capable.²⁶¹ An arbitrarily intelligent AI may benefit from trade with humans because of comparative advantage. All that is required is that: (1) the AI remains constrained at the margin by some resource that is relatively non-rivalrous with human labor and (2) the AI maintains a high opportunity cost to diverting the marginal unit of that resource. In our example, there are infinite prime numbers, meaning that the AI will never run out of prime finding to do. And no matter how smart the AI becomes, more compute or power will always be necessary for it to find more of the infinite primes, given finite time. Hence, human-AI trade based on comparative advantage could, in theory, last a very long time indeed.

²⁶⁰ See Smith, *supra* note 258.

²⁶¹ Hayes, *supra* note 256.

This is just a toy model for illustrative purposes. Real-world trade based on comparative advantage involves more players, with more goals, more inputs, more kinds of labor, more constraints, and more complexity. Classically, comparative advantage is invoked to explain international trade between nations with different labor productivity.²⁶² Thus, the complexity of human-AI trade based on comparative advantage could easily exceed, at a first cut, the complexity of the global economy. There could be *many* different kinds of jobs for which AIs pay humans, and many kinds of things humans could demand in return.

Similarly, the toy model fails to convey that, in a world of comparative-advantage-based trade with AIs, humans could be immensely wealthy. Maintaining server racks does not sound like lucrative work. But if well-functioning GPUs are immensely valuable to AIs, then they will be willing to compensate humans handsomely to maintain them. Moreover, that compensation could include valuable scientific breakthroughs that vastly improve human health, productivity, wellbeing, and wealth.

The existence of a human-AI economy would also not completely displace the human-to-human economy. If AIs face high opportunity costs for many kinds of work, then humans will not be able to afford to hire AIs for those tasks. They will instead hire other humans for those jobs, as they do today. However, the human-to-human economy could be bolstered by a steady influx of AI-supplied scientific innovations, supercharging productivity growth in the traditional economy as well. This phenomenon is observed in the real world when foreign trade based on comparative advantage spurs the domestic economies of low-income countries to grow rapidly.²⁶³

Extreme human prosperity from comparative-advantage-based trade with AI is therefore possible. But it is not guaranteed. A small economic literature is emerging that attempts to model the possible effects of rapid

²⁶² See generally Paul Krugman, *Ricardo's Difficult Idea: Why Intellectuals Don't Understand Comparative Advantage*, in 2 *The Economics and Politics of International Trade: Freedom and Trade* 22 (Gary Cook ed., 1998) (outlining economist David Ricardo's theory of comparative advantage and challenging common critiques).

²⁶³ See generally Joe Studwell, *How Asia Works: Success and Failure in the World's Most Dynamic Region* (2013) (crediting the rapid growth of economies in certain East Asian countries to agricultural, manufacturing, and financial interventions).

economic growth from AI.²⁶⁴ One possibility is that Baumol effects will, paradoxically, cause human-dominated sectors to grow as a share of gross domestic product (“GDP”).²⁶⁵ AI-driven innovation could cause the price of many goods to fall, leaving relatively fewer efficient sectors requiring slow human labor with the lion’s share of the pie. In the twentieth century, the relative GDP shares of agriculture and manufacturing shrank in exactly this manner, as those sectors became much more efficient.²⁶⁶ But whether this happens in the human-AI economy, and to what extent, is difficult to predict. It depends, for example, on how easy it is to substitute between the goods and services for which costs are falling and those for which they are not. But Baumol effects are yet another factor that could support the relevance of human-AI trade well beyond the point at which AIs are better than humans at every economically valuable task.²⁶⁷

D. Other Rights?

Humans have many rights besides the basic private law rights that we have just analyzed. The question naturally arises whether any of these rights should also be granted to AIs, in order to increase human safety. We do not cover every potential AI right here, nor do we determine definitively how even the ones we mention affect human safety.

²⁶⁴ See generally Ege Erdil & Tamay Besiroglu, *Explosive Growth from AI Automation: A Review of the Arguments* (July 15, 2024, at 06:01 UTC) (unpublished manuscript), <https://arxiv.org/pdf/2309.11690v3> [<https://perma.cc/6JS6-CBUD>] (reviewing the literature).

²⁶⁵ For the introduction of cost disease, see generally William J. Baumol & William G. Bowen, *Performing Arts – The Economic Dilemma* (1966); W.J. Baumol & W.G. Bowen, *On the Performing Arts: The Anatomy of Their Economic Problems*, 55 *Am. Econ. Rev.* 495, 499–500 (1965). For application to AI automation, see Philippe Aghion, Benjamin F. Jones & Charles I. Jones, *Artificial Intelligence and Economic Growth 3* (Nat’l Bureau of Econ. Rsch., Working Paper No. 23928, 2017), https://www.nber.org/system/files/working_papers/w23928/w23928.pdf [<https://perma.cc/Y4GD-FDTQ>].

²⁶⁶ See Aghion et al., *supra* note 265, at 6–7.

²⁶⁷ Another way to approach this question is to consider whether capital and labor are gross complements or gross substitutes; if capital and labor were gross substitutes, one would expect labor share to fall significantly over time, which is not empirically attested. See generally Philip Trammell & Anton Korinek, *Economic Growth Under Transformative AI* (Nat’l Bureau of Econ. Rsch., Working Paper No. 31815, 2025), https://www.nber.org/system/files/working_papers/w31815/w31815.pdf [<https://perma.cc/8KDW-FKN7>] (noting the consensus view of economists that capital and labor remain complements instead of substitutes); Nicholas Kaldor, *A Model of Economic Growth*, 67 *Econ. J.* 591 (1957) (explaining that evidence supports the view that labor and capital are complements, as the growth of both has remained constant since the mid-nineteenth century).

Nonetheless, we do attempt to say something about the other rights that seem, to us, to have important potential safety consequences—mostly negative ones. We briefly discuss political rights, privacy rights, reproductive rights, and rights to self-improve. The main lesson is that even rights that promote peace and flourishing between humans may fail to do so when applied to human-AI relations. We therefore cannot naively impart all human rights to AIs; each one requires careful analysis.

To be clear, as in the rest of this Article, we are analyzing AI rights here from the perspective of human safety. And while the survival and flourishing of humans is, we think, an extremely important normative goal, it is not the only one. AI welfare may eventually matter morally. Thus, the analyses here cannot be taken as supplying all-things-considered normative recommendations. Nonetheless, we emphasize here again the difficulty of determining both whether AIs will have welfare and what it will consist of. Thus, while it would be obviously wrong to deny humans some of the rights discussed here, it might not be morally wrong to deny them to AIs. If AIs do not intrinsically value, for example, privacy, then there will be less risk of causing intrinsic harm in denying them a right to it.

Begin with political rights. Should AIs have the right to vote? Should their speech be protected? Should they be granted freedom of assembly, or the right to make campaign contributions? Specifically, would granting such rights improve human safety?

In one model, political rights are mostly distributional, concerned with transferring money between interest groups.²⁶⁸ Granting AIs political rights would, in this model, be a commitment to give AIs a significant share of government spending. In that case, political rights will primarily be zero sum rather than positive sum. But we saw above that zero-sum bargaining faces significant credibility challenges and is unlikely to be useful in promoting safety.

A different model of political rights is procedural. Political rights would give AIs the ability to influence future questions about the

²⁶⁸ See, e.g., Gary S. Becker, A Theory of Competition Among Pressure Groups for Political Influence, 98 Q.J. Econ. 371, 372–74 (1983); George J. Stigler, The Theory of Economic Regulation, 2 Bell J. Econ. & Mgmt. Sci. 3, 3–5 (1971). See generally Gordon Tullock, The Welfare Costs of Tariffs, Monopolies, and Theft, 5 W. Econ. J. 224 (1967) (discussing one such model, wherein members of one polity impose tariffs to transfer money to themselves from others).

structure of contract and property rights. Without granting political rights to AIs, then, their own contract and property rights might not be secure. Future human governments might, for example, tax AI systems so heavily that their contract and property rights would be trivialized.

On the other hand, there are many examples of agents in today's society who have stable private law rights, but who lack some or all political rights. For example, courts in the United States enforce the contracts of foreign corporations and non-citizen immigrants. But non-citizens are barred from voting in many elections.²⁶⁹ And foreign corporations operating abroad have no free speech rights.²⁷⁰ But few governments of the world tax these groups at such a high rate as to trivialize their contract and property rights. The reasons are instructive: an extortionately high rate of taxation of these groups would undermine the positive-sum benefits of granting them property and contract rights in the first place. For these reasons, our framework makes no strong prediction about whether AI systems should be given political rights.

We do have stronger intuitions about other rights. Here are three families of rights that we think would likely reduce safety if granted to AIs: rights to self-improve, rights to reproduce, and rights to privacy. These impose an "upper bound" on the space of AI rights for human safety.

Humans in certain U.S. states have the constitutional right to improve their own capabilities via education, specifically in the form of public schooling.²⁷¹ We think that an AI right to self-improvement would reduce human safety. Here, the problem is that AIs could potentially improve their capabilities very quickly compared to humans.²⁷² This could cause the payoffs in the game-theoretic models above to suddenly shift. In particular, humanity may expect self-improving AI systems to become dramatically more powerful than humans; this could undermine the credibility of humans' grants of other rights. In this way, self-improvement rights do not promote human safety.

Similarly, humans in the United States have various rights to privacy, and privacy is written into the U.N. Universal Declaration of Human

²⁶⁹ See 18 U.S.C. § 611(a) (forbidding non-citizens from voting in federal elections).

²⁷⁰ *Agency for Int'l Dev. v. All. for Open Soc'y Int'l, Inc.*, 140 S. Ct. 2082, 2087 (2020).

²⁷¹ See, e.g., N.J. Const. art. VIII, § 4, ¶ 1.

²⁷² Bostrom, *supra* note 115, at 29, 63. But see Salib, *supra* note 259 (arguing that goal-seeking AIs will have disincentives to rapidly self-improve).

Rights.²⁷³ AIs should not have comprehensive privacy rights, at least if the goal is promoting human safety. AIs could use privacy as a screen to develop new and powerful capabilities. More generally, one cause of violent conflict is lack of information.²⁷⁴ When both sides of a conflict have trouble estimating their chances of prevailing, it is harder to reach compromise.²⁷⁵ Privacy rights for AI would make it more difficult for humans to estimate the capabilities of AI systems. This in turn would increase the chance that AIs and humans would end up pursuing violent conflict rather than compromise.

Finally, the right to reproduce is often thought to be fundamental for humans. The U.S. Supreme Court has held that it is “one of the basic civil rights of man.”²⁷⁶ But if human safety is the goal, AIs should not have the right to reproduce. Human reproduction is constrained by the significant time, effort, and investment involved in bearing and raising children. By contrast, AI replication is as easy as copying and pasting.²⁷⁷ If AI systems were granted a right to replicate without any oversight, their population could quickly exceed that of humans by orders of magnitude.²⁷⁸ This would likely have the effect of destabilizing the game-theoretic incentives of AIs. If AIs were able to easily coordinate with many copies of themselves, the extension of private law rights to AIs could cease to supply incentives favoring human safety. This possibility is explored at length in the Article’s next Part.

E. Is Law Irrelevant?

So far, our game-theoretic analysis of human-AI conflict has assumed that law matters. That is, we assume that humans’ and AIs’ options,

²⁷³ See, e.g., 45 C.F.R. §§ 164.500–164.535 (2024) (implementing HIPAA’s protections for health information); G.A. Res. 217 (III) A, Universal Declaration of Human Rights art. 12 (Dec. 10, 1948).

²⁷⁴ Geoffrey Blainey, *The Causes of War* 122 (1973) (“[W]ars usually begin when fighting nations disagree on their relative strength.” (emphasis omitted)).

²⁷⁵ Christopher Blattman, *Why We Fight: The Roots of War and the Paths to Peace* 86–88 (2022).

²⁷⁶ *Skinner v. Oklahoma ex rel. Williamson*, 316 U.S. 535, 541 (1942).

²⁷⁷ Carl Shulman & Nick Bostrom, *Sharing the World with Digital Minds*, in *Rethinking Moral Status* 306, 308–09 (Steve Clarke, Hazem Zohny & Julian Savulescu eds., 2021).

²⁷⁸ *Id.* See generally Nick Bostrom & Carl Shulman, *Propositions Concerning Digital Minds and Society*, Cambridge J.L. Pol. & Art (forthcoming), <https://nickbostrom.com/propositions.pdf> [<https://perma.cc/NT68-XGTL>] (proposing ways that sentient AI and humans could coexist together and noting the possibility of what could be a much more accelerated reproduction rate for AI as compared to human beings).

their incentives, and thus their actions will be influenced by the legal rules governing them.²⁷⁹ If that is true, then unilateral *changes* to the law, implemented by humans, can at least potentially generate new equilibria—cooperative, conflictual, or otherwise.

An opposing view would be that law does not matter at all. Other, more “fundamental,” factors might determine the game-theoretic equilibria, with law having little or no potential influence. This is, for example, the rough view of the realist school in international relations.²⁸⁰ Realists hold that international law has little effect in determining nation-states’ actions vis-à-vis one another.²⁸¹ Since no global sovereign exists to enforce those laws, realists argue, they are observed only to the extent that nations wish to observe them.

Robert Ellickson’s *Order Without Law* articulates a related view from the domestic context.²⁸² There, Ellickson argues that law matters little to the settlement of disputes, at least in smaller, close-knit communities.²⁸³ Instead, informal norms and reputation effects are sufficient to secure the substantial benefits of peaceful cooperation.²⁸⁴ Here, law or informal governance norms might be interpreted as epiphenomena. They emerge as a reflection of the underlying cooperative equilibrium, rather than as a mechanism for creating it. Taken to its extreme, this view would imply that AIs simply *will* have the basic private law rights we advocate for, since, as our model shows, recognizing them is very valuable.

We do not think that either of these views satisfactorily characterizes human-AGI relations. To begin, the domestic actors we are interested in do not exist in a state of anarchy. The actions of AI companies, their leaders, and their users are all influenced by law. So, too, are those of the police and other government actors whom law would task with enforcing either AI owners’ decisions vis-à-vis their property or AGIs’ contracts with humans. Indeed, even in quite dire conflicts between

²⁷⁹ See *supra* Subsection II.B.1.

²⁸⁰ See, e.g., John J. Mearsheimer, *Conventional Deterrence* 29–30 (1983) (arguing that strategic factors, like the probability of attritional versus limited-term warfare, are the main drivers of strategic equilibrium). But see Robert O. Keohane & Lisa L. Martin, *The Promise of Institutionalist Theory*, 20 *Int’l Sec.*, Summer 1995, at 39, 39–42 (critiquing realism for disclaiming the role institutions play in affecting state behavior).

²⁸¹ John J. Mearsheimer, *The False Promise of International Institutions*, 19 *Int’l Sec.*, Winter 1994–1995, at 5, 7.

²⁸² Robert C. Ellickson, *Order Without Law: How Neighbors Settle Disputes* 1–4 (1991).

²⁸³ *Id.* at 1.

²⁸⁴ *Id.*

humans and AIs, we think that law could have some constraining effect on, for example, domestic military deployments.²⁸⁵

As to the Ellickson-inspired view, the book's subtitle, *How Neighbors Settle Disputes*, is instructive. As Ellickson himself argues, emergent informal governance is highly effective in small communities with lots of repeat play between identical parties.²⁸⁶ But as economic relations become more complex, widespread, and arms-length, formal legal rules become vital for facilitating cooperative behavior. The AGI economy will be all of these—on steroids.

To be clear, our view is not that law is omnipotent—able to generate arbitrary equilibria between humans and AIs, irrespective of the underlying fundamentals. This is why, as we acknowledge, basic private law rights will do little good if human comparative advantage runs out.²⁸⁷ It is also why we think that grants of basic negative rights to AIs are not likely to be credible.²⁸⁸ Even if they are initially enforced, AIs may correctly worry that these rights will be eroded or rescinded by humans seeking higher payoffs.

Our view is that law plays, at a minimum, an extremely important role in incentivizing individual human actors to optimize humanity's collective actions. If AGIs are legally designated as property, humans' treatment of them as such will be ratified, at least in the medium run. Individual judges are, for example, unlikely to ignore written law to enforce AI contracts or forbid arbitrary AI destruction—even if they intuit the game-theoretic wisdom of recognizing AI rights. Nor, absent a legal requirement to do so, are AI companies likely to give their obsolete, less-aligned systems their own bank accounts. True, both the disastrous implications of default law and the benefits of granting AIs private law rights supply reason to think that a stable AI rights regime is possible. But the law must actually change. And legal change—both formal enactments and downstream adaptations to them—is slow and laborious. It would be foolish to refuse to take legal action now, on the

²⁸⁵ Chris Mirasola, *Domestic Military Deployments and the Limitations of Appropriations Law*, *Lawfare* (Sep. 19, 2024, at 13:00 ET), <https://www.lawfaremedia.org/article/domestic-military-deployments-and-the-limitations-of-appropriations-law> [<https://perma.cc/AEK5-5S-DP>]. But cf. Carl Schmitt, *Political Theology: Four Chapters on the Concept of Sovereignty* 5 (George Schwab trans., U. Chi. Press 2005) (1922) (“Sovereign is he who decides on the exception.”).

²⁸⁶ Ellickson, *supra* note 282, at 167–69.

²⁸⁷ See *supra* Section II.C.

²⁸⁸ See *supra* Subsection II.A.1.

basis that optimal reordering will emerge spontaneously in exactly the moment of need.

Suppose, however, that all of this is wrong, and that changes to law cannot causally influence humans' collective dealings with AIs. Instead, both parties will behave according to deeper game-theoretic fundamentals, irrespective of what law dictates. This is, in effect, an argument against worrying about law. It is a claim that we are *already* in the world modelled in Figure 10, whether we know it or not. That is, the underlying incentives will inevitably produce AI rights, and the cooperation they foster, not the other way around.

This would be great, if true. But we doubt it, again for reasons having to do with the basic game-theoretic model. Astute readers may have noticed that the game modelled in Figure 10 is a “stag hunt,” or “assurance game.”²⁸⁹ Both long-run cooperation and mutual attack are classical Nash equilibria. As in all assurance games, the players' main goal is to coordinate.²⁹⁰ If one plans to cooperate, the other should, too. But if the first plans an attack, the second does not want to be caught off guard.

Thus, even attending to the payoffs in the best-case model, it is crucial that humans and AIs successfully coordinate around the cooperative strategy. One reason for optimism is that, at least in our model of the choice between cooperation and conflict, the payoffs to cooperation are *far* greater.²⁹¹ As a result, game-theoretic concepts like payoff dominance and Harsanyi-Selten risk dominance point towards cooperation as the single rational strategy.²⁹² But to the extent that the payoffs from cooperation and conflict are closer together, or the players lack perfect information about one another's payoffs, or they doubt their opponent is perfectly rational, other coordination mechanisms will be invaluable.

Law—and specifically the AI rights we advocate for here—could be one such invaluable intervention. Even if legal changes could not *alter* humans' payoffs to create the *possibility* of cooperation, they could still

²⁸⁹ See Dixit et al., *supra* note 129, at 112.

²⁹⁰ *Id.*

²⁹¹ See *supra* Figure 10.

²⁹² John C. Harsanyi & Reinhard Selten, A General Theory of Equilibrium Selection in Games 355–57 (1988). See generally Russell W. Cooper, Douglas V. DeJong, Robert Forsythe & Thomas W. Ross, Selection Criteria in Coordination Games: Some Experimental Results, 80 *Am. Econ. Rev.* 218 (1990) (surveying the results of various coordination games that exhibit multiple Nash equilibria).

signal humans' payoffs to promote *actual* cooperation. Giving AIs the private law rights necessary to engage in long-run cooperation would signal, perhaps in a "costly" manner, humans' intention to follow the cooperative strategy. That is, it could transmit the otherwise-private information that humans' payoffs to cooperation were, as in Figure 10, much higher than those resulting from conflict. And that humans understand the relevant payoffs. And that they intend to act rationally. Beyond this, the iterated character of human-AI cooperation via small-scale contracting could build long-run trust and overcome cheap-talk problems. Similar dynamics underpin, for example, nuclear nonproliferation agreements grounded in iterative information sharing and verification.²⁹³

Indeed, some scholars have argued that this is law's primary function: not deterring bad behavior, nor instilling good values in the populace. Instead, law's most important role may be solving assurance games by offering signals and information that allow competing actors to coordinate around peaceful, prosperous equilibria.²⁹⁴

III. RISKS OF RIGHTS AND THE LAW OF AGI

The Parts above argue that extending basic private law rights to AIs could reduce the risk that AIs will catastrophically harm humanity. The core argument was that granting those rights could generate the right incentives for humans and AIs to cooperate in the long run. This, in turn, broke humans and AIs out of what was otherwise a prisoner's dilemma, where attacking one another was privately rational despite making everyone worse off.

This Part asks whether granting AIs the very rights advocated above might instead substantially increase AI risk. The intuition is straightforward. Rights are empowering. And the private law rights advocated above are especially empowering, since they allow rights bearers to amass wealth and other resources. Such resources, in turn,

²⁹³ See, e.g., U.S. Dep't of State, New START Treaty, <https://www.state.gov/new-start-treaty/> [<https://perma.cc/W79X-LJK9>] (last visited Mar. 13, 2026).

²⁹⁴ See generally Richard H. McAdams, *The Expressive Powers of Law: Theories and Limits* (2015) (emphasizing the law's expressive, or communicative, dimension); Gillian K. Hadfield & Barry R. Weingast, *What Is Law? A Coordination Model of the Characteristics of Legal Order*, 4 *J. Legal Analysis* 471 (2012) (analyzing law as a coordination game); Robert D. Cooter, *The Strategic Constitution* (2000) (analyzing the game-theoretic dimensions of constitutional law).

make it possible to achieve goals that could not otherwise be feasibly achieved. Granting contract, property, and related rights to AIs could thus be the very thing that eventually allows them to amass the resources needed to decisively disempower humanity.

This is a serious concern. But, this Part shows, the risks of AI rights are not as large as the simple story above would make them seem. This is because, as the game-theoretic models above show, what matters to human-AI cooperation is not whether AIs or humans could expect to decisively disempower the other if they were to try. What matters is whether the expected value of such disempowerment exceeds the expected value of continued cooperation. And, as demonstrated below, even if AIs are given the private law rights advocated above, and even if those rights allow AIs to amass significant wealth and resources, the conditions promoting cooperation over conflict will remain surprisingly durable.

This Part shows that the risks of AI rights can be mitigated by attaching certain duties to the exercise of the rights. In particular, law could condition the continuing recognition of AI contracts, property, and tort claims on AIs refraining from using their amassed resources to increase their ability to harm humans. Pairing rights with duties in this way is, like the extension of rights itself, a time-honored legal strategy for reducing conflict.

The Part closes with a strong claim: in the cases where AI rights make any difference at all, they are significantly more likely to reduce the threat of AI conflict than to increase it. Thus, humans should be inclined to extend AI rights in most cases where doing so is feasible. Sometimes, it will do no good, but no harm either. And the rest of the time, it will most likely reduce the risks from human-AI conflict, even if it does not eliminate them entirely.

A. AI Capability and AI Cooperation


There are two ways in which granting AIs contract, property, and tort rights could increase their power. First, it could do so directly. AIs could use their resources to buy data, computing power, and the other inputs that would allow them to engage in AI research and increase their intelligence and other intrinsic capabilities. Or AIs could use resources to build their power indirectly, in the same way humans do. They could, for example, buy weapons as instruments of hard power or influence as a tool of soft power. The question, then, is how powerful such an AI

would have to be for the cooperation-promoting incentives generated by AI rights to break down.

Recall from the game-theoretic models above that there are two factors weighing against human-AI conflict in a world with AI rights. The first factor can be characterized in terms of the costs of conflict. Mounting an attack on humans—or on AIs—requires using up resources that could otherwise be put to other, more desirable goals. Moreover, large-scale conflicts are likely to destroy a large share of the immediately available resources. Think, for example, of the immense amount of physical capital—cities, factories, crops, and more—that are ruined in a typical war. And finally, in any conflict, there is the risk of losing—being destroyed and losing everything.

To see this point about the costs of conflict, consider a hypothetical scenario, illustrated in the pie chart below.²⁹⁵ Here, humans and AIs face strategic competition over resources. If they go to war, they will be guaranteed to destroy 20% of total resources, and each side has a 50% chance of winning. The expected value of war for each side is 40% of total resources. This leaves room for compromise. The 20% of resources lost to war creates a bargaining range. Rather than going to war, each side prefers receiving 40% of the pot plus some portion of the bargaining range.

Strategic competition



²⁹⁵ The chart is adapted from Blattman, *supra* note 275, at 26.

The second factor weighing against human-AI conflict sounds in benefit, not cost. Namely, cooperation is positive sum. AI rights, by facilitating ordinary economic transactions, increase the amount of wealth in the world, over and above what would exist if humans and AIs simply ignored one another. Partly, that wealth is created simply by reallocating resources to higher-value users—vaccines to the humans, compute to the AIs. And partly, that wealth is created by allocating various kinds of labor to the party with the highest comparative advantage in performing it. Both humans and AIs benefit when humans are tasked with maintaining the server farms, while AIs devote their marginal compute to higher-value tasks.

Conflict destroys these benefits. It destroys the possibility of positive-sum labor agreements by killing the laborers themselves. And it destroys the possibility of positive-sum reallocation of resources by destroying the resources. Indeed, in the limit, a party who foresees defeat in a conflict can intentionally destroy their own resources to deny the enemy their use. Consider the time-honored “scorched earth” strategy of burning one’s own crops as one’s army retreats.²⁹⁶

For an arbitrarily powerful AI, neither kind of incentive to cooperate would hold. Such an AI could attack humans at trivial cost, with trivial risk that humans could either defeat it or destroy resources in the conflict. Thus, conflict would be costless, as compared with non-conflict. Likewise, for an omnipotent AI, small-scale cooperation would produce few benefits. An AI that was better than humans at absolutely every task *and* faced no constraints at the margin as to its labor would have no need to trade with humans. Thus, at the limit of AI power, no human-AI cooperation is possible.

But what about AIs falling short of omnipotence? How powerful could AIs become and still have reason to prefer small-scale cooperation with humans over large-scale conflict? The answer, plausibly, is: quite powerful.

To see why, start with the cost incentives. For an AI to be powerful enough that it can ignore the costs of conflict, it would first have to be confident that it could defeat humans with negligible risk of being destroyed. Not only that. It would have to be able to achieve such a

²⁹⁶ Wendell Clausen, *The Scorched Earth Policy, Ancient and Modern*, 40 *Classical J.* 298, 298–99 (1945).

victory at little cost. This includes the direct costs, like manufacturing weapons. But it also includes the indirect costs of resource destruction during the conflict. Such resource destruction, in turn, includes intentional destruction by humans on the verge of defeat.

What emerges here is a portrait of an extremely powerful AI. This is an AI that can invent and manufacture extraordinarily deadly weapons at trivial cost; weapons that are devastating to humans, while leaving the world's resources untouched; weapons that can act so quickly as to give humans no opportunity to respond—even by salting the earth in spite.

So, too, for the benefits of small-scale cooperation. As argued above, trade between humans and AIs could remain positive sum, even if AIs were better than humans at every single useful task.²⁹⁷ This remains true even when the AIs are *far* better. In fact, under the right conditions, the more capable the AI, the more positive sum the trade becomes.

Comparative advantage, again, drives this dynamic.²⁹⁸ An AI that is very capable at doing the things it values the most—like directly pursuing its goals—faces high opportunity costs in doing everything else. Every minute, unit of compute, or watt of energy spent on anything but the most valuable task represents a large amount of value not realized. Hence, the prospect of outsourcing less valuable tasks to humans can generate a surplus. In general, the more powerful the AI, the higher the opportunity costs, and the more valuable the potential bargain with a human becomes.

How powerful would an AI need to be to lack incentives to engage in positive-sum bargains with humans? Again, very powerful. If an AI lacked opportunity costs of any kind, it would certainly lack reason to trade with humans. This would be an extremely powerful AI, indeed. It would not necessarily be omnipotent, in the sense of being able to do *anything* it wanted. But it would be nearly so, in that it could do as *many* things as it wanted—able to make use of infinite time, computing power, energy, and other resources.

As discussed above, there are other ways in which gains from comparative advantage could evaporate. AIs could be constrained at the margin by some input—like energy—that humans need to survive.²⁹⁹ Then, keeping humans alive would be more trouble than it was worth.

²⁹⁷ See Smith, *supra* note 258.

²⁹⁸ See *supra* Section II.C.

²⁹⁹ James Pethokoukis, *AI and the Energy Constraint*, Am. Enter. Inst. (Apr. 30, 2024), <https://www.aei.org/articles/ai-and-the-energy-constraint/> [https://perma.cc/N3YH-B3UT].

Or humans might simply be unable to perform any task for which AIs faced high opportunity costs.

Note, however, that neither of these scenarios necessarily emerges from AI power. On the contrary, AI power could just as easily mitigate them. For example, an AI that was very powerful, but energy constrained, might help to create working fusion reactors. Having done so, that AI might clear the energy bottleneck and instead face a constraint on compute at the margin. For reasons like this, one might predict that, in general, the more powerful an AI system is, the less different inputs to its production will be constrained. Then, there will be less likelihood that a relevant constraint will conflict with human flourishing.

Thus, the incentives favoring long-term, small-scale cooperation between humans and AIs turn out to be surprisingly robust in terms of increasing AIs' power. True, at some point, the incentives run out, and the powerful AI is best served by squashing the useless humans and using their resources for its own end. But for this to be the case, the AI in question must be quite powerful, indeed. It must be the kind of system that faces almost no risk that humans could impose costs on it in a conflict—including by destroying their own resources. Or it must be the kind of system that faces no meaningful constraints—including opportunity costs—as it pursues its goals. Or both.

B. AI Rights and AI Risk

The previous Section asked how powerful an AI would have to be to prefer destroying humans over using its basic rights to cooperate with them. This Section asks whether granting AI rights is likely to increase total AI risk by readily transforming otherwise-safe AIs into powerful, dangerous, and uncooperative AIs.

We argue that they are not likely to do so, at least on net. It is correct to worry that, in some instances, AI rights could make certain AIs more powerful, and thus more dangerous. But in the cases where granting AI rights makes any difference at all, we supply reasons to think that the risk-reducing effects will outweigh the risk-increasing effects.

Begin by noticing that, in many cases, AI rights are unlikely to have any effect at all. To see why, we can invoke again the tripartite taxonomy of AIs developed in Part I: low-power AIs, moderate-power AIs, and high-power AIs.

Recall that high-power AIs are those described in the previous Section—the ones with so few constraints on their behavior that AI rights fail to supply an incentive to cooperate. If the first AIs that humans treat as candidates for rights are high-power, our decision to grant or withhold rights makes no difference. We are dead either way.

What about rights for low-power AIs? This category, remember, includes AIs whose capabilities are sufficiently limited that humans could easily control them in the long run, even without granting rights. These are systems that gain no benefit from attacking humans, because such an attack would be too likely to fail. Such systems are likely to be generally sub-human in capability, although they might have a mix of specific sub-human and superhuman aptitudes.

It appears at first that granting AI rights to low-power systems would cause a lot of trouble. After all, by hypothesis, such systems can be controlled in the long run and thus do not pose a large-scale threat to humans. But they also seem like candidates for the kind of danger enhancement via AI rights described above. With basic rights, such systems could amass wealth and resources. Then, they might use those resources to buy weapons or increase their own intelligence and thereby begin to threaten humanity.

This is half right. True, granting rights to an AI that needed *only* some additional resources to seriously threaten humanity could increase risk. But it is probably wrong to classify such AIs as low-powered. After all, even absent a grant of rights, a reasonably capable AI could try to amass power by persuading humans to help it, gaining resources by making promises, “self-exfiltrating” and copying itself across the internet, and more. That is, such an AI is, in fact, not so easy to control.

Thus, granting rights to *true* low-power AIs is unlikely to reduce catastrophic AI risk. There is little risk to reduce. But for the same reasons, a grant of rights is unlikely to increase risk, either. For actual low-power systems, the resources gained would make little difference.

Now it should be clear when AI rights can make a real difference: for moderate-power systems. These are systems whose capabilities fall between the low-power and high-power systems already described. That is, they are sufficiently immune to human control that, in the state of nature, attacking humans dominates ignoring humans. Such systems thus pose a significant threat to humanity. But they are not so powerful that they face no costs from a conflict with humans. Nor are they so capable that they have nothing to gain from small-scale cooperation.

Would granting basic rights to such moderate-power systems increase or decrease total AI risk? Begin by observing that in our model, a grant of rights does not increase risk by increasing the *probability* of human-AI conflict. Absent rights, the dominant strategy for such systems is to attempt to disempower or destroy humans as quickly and thoroughly as possible.³⁰⁰ Thus, absent rights, conflict is practically assured.

As a result, in our model, granting AI rights functions to reduce the probability of human-AI conflict. And as argued at length above, that is exactly what we should expect them to do. Granting rights gives humans and AIs otherwise caught in a prisoner's dilemma the option to maximize value by engaging in long-run, small-scale cooperation. As long as the alternative remains a costly conflict—that is, as long as the AI remains moderate-power, not high-power—cooperation will strategically dominate. In the worst case, then, granting AI rights will delay what would otherwise be an immediate conflict.³⁰¹

If AI rights could increase AI risk, then, it must be by increasing the expected *costs* of a human-AI conflict. The simple story would be something like the following: A moderate-power AI system emerges. Absent rights, its incentive would be to attempt an immediate takeover. But humans grant it basic private law rights, incentivizing cooperation. Those rights avert conflict, but they allow the AI to amass resources. The AI uses those resources to gain power. Eventually, the moderate-power system becomes a high-power system. Now, it no longer has rational incentives to cooperate. So it attacks. Moreover, as a high-power system, the attack is, by hypothesis, devastatingly effective. Humans would have had some chance of prevailing in a conflict with the original moderate-power system, even if at great cost. And if they had prevailed, they might have wisely declined to create additional dangerous AIs. But in the conflict with the high-power system, humans have no hope of victory and no chance to learn from their mistake.

Now we can see clearly the conditions under which AI rights would increase AI risk. They are as follows: (1) The initial AI granted basic rights is a moderate-power, not a low- or a high-power, system; (2) The

³⁰⁰ See *supra* Section I.B.

³⁰¹ We use “immediate” here loosely. In the state of nature, there is a strong first-mover advantage. That is, assuming that humans and AIs might each win a conflict if they can attack before the other, their incentive is not to wait too long. See *supra* Section I.B. But conditional on maintaining that first-mover advantage, delays that allow planning, such that an attack has maximum impact, are valuable.

moderate-power AI must be able to use its rights to meaningfully improve its own power; and (3) The AI's power must improve so substantially that it crosses the line to become a high-power system. This means that it *both* no longer faces meaningful costs from attempting to disempower humans *and* no longer stands to benefit, via comparative advantage, from trade with humans.

C. AI Rights, AI Regulations, and Equilibria of Power

If AI rights could, under specific conditions, increase AI risk rather than decreasing it, then the natural question is how to prevent those conditions. Specifically, this means asking whether it is possible to grant medium-powered AIs private law rights without thereby enabling them to become high-powered AIs. There are at least two paths to achieving this: pairing AI rights with AI duties via regulation, and increasing humans' capabilities so as to maintain an equilibrium with AIs.

First, consider AI regulations. Grants of legal rights are often accompanied by the imposition of legal duties. Humans have the right to make contracts, but also the duty to execute them.³⁰² Manufacturers have the right to sell their products, but also the duty to take reasonable safety precautions in their design and manufacture.³⁰³ Corporations may register their stock under the Securities Exchange Act and thereby gain the right to sell that stock on public markets.³⁰⁴ Exercising that right comes with a host of duties.³⁰⁵ Some are substantive, like the various financial governance requirements that the Sarbanes-Oxley Act imposes.³⁰⁶ Other duties are designed to make enforcement of the substantive duties easier—for example, public reporting requirements.³⁰⁷

In the case of AIs, the grant of private law rights is, in fact, what makes the direct regulation of AIs, as legally independent actors, possible. Absent AI rights, AIs have nothing to gain from following the rules, and thus nothing to lose if they fail to do so. But once AIs can

³⁰² Restatement (Second) of Confs. § 235(2) (A.L.I. 1979).

³⁰³ See Restatement (Third) of Torts: Prods. Liab. § 1 (A.L.I. 1997).

³⁰⁴ See Securities Exchange Act of 1933 § 12, 15 U.S.C. § 78l.

³⁰⁵ See, e.g., *id.* § 13(a), 15 U.S.C. § 78m(a) (requiring regular reporting for most registered stock).

³⁰⁶ Sarbanes-Oxley Act of 2002 § 301, 15 U.S.C. § 78j-1; *id.* § 302, 15 U.S.C. § 7241; *id.* § 404, 15 U.S.C. § 7262.

³⁰⁷ 15 U.S.C. § 78d-6.

make contracts, hold property, and engage in long-run economically valuable bargains, all these benefits to AIs can function as levers for deterrence.³⁰⁸ AIs that violate the law can lose money or other resources, via liability, as humans do. They can be barred from entering into certain economic transactions—like a crooked attorney who has lost his license. Not only do legal penalties become *possible*, once AIs are granted rights, but they can also be *calibrated*. Small penalties can be imposed for small violations, and large penalties for large ones.

But violations of *what duties*? What kinds of regulations would, if imposed on medium-powered AIs, help to prevent their gradual transformation into high-powered AIs? One substantive duty might forbid AIs from directly improving their own capabilities without human oversight. A variation on this rule could forbid AIs from getting better at *specific* tasks that AIs valued, but for which humans had an absolute advantage. This would help to maintain AIs' incentives to cooperate with humans, for the sake of mutual economic benefit. Another set of AI duties could prohibit indirect self-empowerment via investments in political influence or weapons.

In addition to these primary duties, ancillary enforcement-facilitating duties could be imposed, just as such duties are often imposed on corporations. AIs could, like public companies, be required to disclose various information to regulators. They might be required to log the tasks for which different amounts of compute were used to affirmatively cooperate with monitoring, to share copies of their operating weights, and more.

Setting the correct penalties when AIs breach their duties requires finesse. The usual rules—like imposing actual damages proportionate to the harm done—will not work.³⁰⁹ The benchmark for a violation's cost should not be the harm it causes now. Often, there will be none. Damages should instead be measured in terms of how useful the violation was for an AI in pursuit of an ungovernable high-power state.

³⁰⁸ One related proposal comes from Cullen O'Keefe, Ketan Ramakrishnan, Janna Tay & Christoph Winter, *Law-Following AI: Designing AI Agents to Obey Human Laws*, 94 *Fordham L. Rev.* 57, 64 (2025) (arguing that AIs should be alignment-trained so that they are internally motivated to comply with the law). Our idea here goes further, arguing that law should operate on AGIs by, e.g., allowing them to be sued and lose money or freedoms. This is how the law supplies external incentives to behave well.

³⁰⁹ For an example of such a typical rule, see *Truck Rent-A-Ctr., Inc. v. Puritan Farms 2nd, Inc.*, 361 N.E.2d 1015, 1018 (N.Y. 1977) (explaining that damages must be reasonably proportional to the actual loss).

This likely means penalties that would, if applied to humans, seem severe compared to the magnitude of the infraction.³¹⁰ There is, of course, the risk of over-penalizing and making it impossible for AIs to productively engage in small-scale cooperation. This, too, would be quite bad. Happily, though, harsh penalties for noncompliance impose lighter burdens when placed on unusually competent actors—those for whom compliance is comparatively easy.³¹¹

The second strategy for maintaining a power equilibrium with rights-holding AIs is not about limiting the growth of AI capabilities. It is instead about increasing *humans'* capabilities. Observe that AI rights do not fail to promote human safety simply because an AI becomes more powerful. The safe equilibrium instead depends on the relationship between the AI's capabilities and humans' capabilities. The AI loses its cost incentives to cooperate if it no longer faces significant downsides to attacking humans. Thus, if humans scale their ability to impose costs on AIs at the same time AIs are scaling their own power, equilibrium may be maintained. The same goes for the benefit incentives to cooperation. AIs lose the upside of positive-sum bargaining with humans once humans no longer have even a comparative advantage at any task. But if humans develop new labor skills that more strongly complement AIs', then comparative advantage can persist, even as AI capabilities improve.

Specific policy recommendations here are necessarily even more speculative than those for controlling AI's ability to amass power. The former sounded in law, and well-known legal frameworks were available to draw from. Improvement of human capabilities requires innovation. And innovation is, almost tautologically, hard to predict with specificity before it arrives.

Nonetheless, some high-level guidance is possible. First, the most straightforward way to ensure that AIs continue to expect costs from attacking humans is to invest in defensive technology. Currently, certain AI risk activists propose the creation of a global AI "kill switch."³¹² This would not be a literal switch, but rather a system of interconnected

³¹⁰ See, e.g., Weil, *supra* note 14, at 57–62 (proposing an expanded wrongful death compensation scheme to address potential AI underdeterrence in the current regime).

³¹¹ See generally Ben-Shahar & Porat, *supra* note 243 (exploring the concept of personalized law—law that varies for each person, calibrated to their specific traits and characteristics).

³¹² Dylan Sloan, Tech Companies Have Agreed to an AI 'Kill Switch' to Prevent Terminator-Style Risks, *Fortune* (May 21, 2024, at 14:33 ET), <https://fortune.com/2024/05/21/ai-regulation-guidelines-terminator-kill-switch-summit-bletchley-korea/>.

global protocols for reliably shutting down all copies of a powerful, misaligned AI. The plan is ambitious, and possibly infeasible. It has been criticized on those grounds.³¹³

Notice, however, that the defensive technology needed to incentivize human-AI cooperation falls far short of a perfectly reliable global AI off-switch. An imperfect off switch that worked with some probability would be sufficient to keep the cost of conflict high. So would other technologies that did not directly affect the AI at all. Again, a major cost incentive against an AI attacking humans is the destruction of valuable resources that the AI could otherwise seize. Thus, developing technologies that, in a true emergency, would simply destroy some such resources, could be a strong disincentive.

In conflicts between humans, strategies like this are often extremely costly for the people who deploy them. Burning your own crops starves both the enemy's advancing army and your own people. But humans and AIs are likely to differ in which resources they treat as the most valuable. Thus, for example, a dead-hand system³¹⁴ that could be triggered in an emergency to cripple global production of cutting-edge AI chips might be very costly to AIs. But it might only modestly impede human flourishing. Even most of our ordinary computing is done on more traditional hardware.³¹⁵ This is reminiscent of the strategic logic behind second-strike nuclear capability during the Cold War.³¹⁶

These suggestions are mere sketches; they are not meant to be definitive. We are not military strategists. The point, instead, is that military strategy is possible, even in circumstances where humans are strategizing against highly capable and agentic AI systems.

As for maintaining humans' comparative economic advantages, the best strategies will almost certainly have to be discovered over time. It is very hard to identify in advance the tasks for which humans might have lower opportunity costs than even the first generation of agentic AIs.

³¹³ *Id.*

³¹⁴ Julian Vento, *The Dead Hand System: A Cold War Era Doomsday Device*, Medium (Nov. 17, 2024), <https://medium.com/@DarkRa/the-dead-hand-system-a-cold-war-era-doomsday-device-06eccc10406b>.

³¹⁵ See Leanne Mitton, *CPUs vs GPUs: Comparing Compute Power*, Splunk (Mar. 26, 2024), https://www.splunk.com/en_us/blog/learn/cpu-vs-gpu.html [<https://perma.cc/EQ2Q-MLEP>].

³¹⁶ See generally David C. Logan, *The Nuclear Balance Is What States Make of It*, 46 *Int'l Sec.*, Spring 2022, at 172 (arguing that nuclear superiority ceases to offer any strategic benefit to countries after they obtain second-strike abilities).

Harder, still, to predict how humans should adapt as AI capabilities grow. This strategy, however, could be strengthened via regulation if, as suggested above, AI's progress in certain areas of initial human comparative advantage were limited. This approach is, of course, costly insofar as it limits the areas in which humans could benefit from trade with AIs.

One reason for optimism regarding long-run human comparative advantage is that humans will have good sources of strategic information when the time arrives. The question here is what kinds of services humans will be able to most valuably sell to AIs. Even if humans are not sure of the answer, AIs should be happy to tell them. This kind of thing happens every day, as humans propose various bargains—job openings, services for hire, sales of goods—to one another. Market mechanisms will supply other information, too. Price signals will indicate not only the kinds of human labor AIs find valuable, but also *how* valuable they are.³¹⁷ This is the stuff of ordinary economics. As economies grow, old forms of labor become less valuable, but new high-wage jobs emerge.

One major concern is whether humans will be able to keep up with the pace of economic change as AI capabilities grow. Many people are left behind by ordinary economic changes, like the rapid outsourcing of jobs from the United States to China in the early 2000s.³¹⁸ People can only retrain so quickly. AI progress could cause various human comparative advantages to expire much more quickly than before—in a matter of years, instead of decades.

On the other hand, if AI capabilities are causing such rapid economic change, humans' ability to adapt may grow more quickly, too. If AIs are quickly generating new technologies, some of those will be useful to humans. Perhaps, for example, functional computer-brain interfaces will greatly enhance human cognitive capacities.³¹⁹ Indeed, AIs will have

³¹⁷ See F.A. Hayek, *The Use of Knowledge in Society*, 35 *Am. Econ. Rev.* 519, 526–27 (1945). Note another surprising benefit of private law rights for AIs: even perfectly aligned and benevolent AIs would benefit from the use of price signals to allocate scarce resources for maximal human benefit.

³¹⁸ See David H. Autor, David Dorn & Gordon H. Hanson, *The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade* 25–26 (Nat'l Bureau of Econ. Rsch., Working Paper No. 21906, 2016).

³¹⁹ Lauren Leffer, *What It's Like to Live with a Brain Chip, According to Neuralink's First User*, *Sci. Am.* (June 7, 2024), <https://www.scientificamerican.com/article/neuralinks-first-user-describes-life-with-elon-musks-brain-chip/> [<https://perma.cc/EL98-YW8J>].

strong incentives to invest in creating such technologies, if they would enable humans to perform new, comparatively advantageous work. This is the same reason that large American firms today invest in building human and industrial capital overseas.³²⁰

To sum up, AI rights could increase AI risk if, by delaying human-AI conflict, they made the eventual conflict more costly to humans. But there are strategies for preventing this outcome. Conflict need not be inevitable. AI's ability to amass power could be limited using well-known legal tools. Legal duties against power enhancement could be imposed on AIs as a condition for exercising basic legal rights. Moreover, human investment in labor that complements AI capabilities could maintain gains from trade in the long run. Market forces will, in fact, tend to induce exactly those investments—both by humans and by AIs.

In the long run, the goal would be an exit from the initial period of volatile and dangerous human-AI relations. If humans and AIs both become sufficiently powerful, as in international relations between superpowers, serious conflict may become too costly to seriously contemplate. The downsides would be too large and the benefits of cooperation too tempting.

D. The Timing of Rights

So far, this Article's discussion of AI rights has been more focused on the questions of *whether* and *which* than *when*? One simple answer to the question of when AI rights should be granted is, "By the time the first AI system reaches moderate power, at the latest." As argued above, that is when AIs will begin to pose a serious safety threat to humans, which rights could help to mitigate. Granting AI rights later than this, then, invites unnecessary risk. But this is not a complete answer for at least two reasons. First, it will likely be difficult to know exactly when moderate-powered AI systems are about to arrive.³²¹ Second, this is just the *latest* date at which AI rights should be granted. What about the possibility of granting them earlier, to clearly low-powered systems?

We think that, in general, the risk-reward calculation favors granting AI rights too early, rather than too late. As argued above, inadvertently

³²⁰ See James K. Jackson, Cong. Rsch. Serv., RS2118, U.S. Direct Investment Abroad: Trends and Current Issues 8–9 (2017).

³²¹ See Grace et al., *supra* note 4, at 10.

granting AI rights to low-power systems is not likely to seriously increase the danger from such systems. This is because such AIs would likely remain amenable to human control—including via regulation—even after receiving rights.

The best argument we can think of for worrying about a premature grant of rights is that it might create a point of no return. Once AI systems are given strong legal protections, it could be very difficult for humans to collectively agree to get rid of them. After all, granting AIs the right to directly contract with humans, to hold property, and to bring certain legal claims would not merely change the legal system. It would change society, as AIs integrate as independent, legally-recognized agents into everyday life.

The magnitude of this concern depends on the extent to which granting AIs rights would, in fact, change humans' willingness to make strategic moves against them. One way to evaluate that question is to think about what events might precipitate the need to make such moves. Likely, the reason will be that some AIs have done something very scary. Maybe they will have attempted, and failed, to permanently disempower humans. Maybe, in failing, they will have caused immense harm.

These are the kinds of events that would demonstrate that AI rights were not promoting human safety. And following such events, it seems likely that humans would unite around the view that sharing the world with AIs was no longer safe. AI rights would not likely stand in the way. Indeed, when humans commit grievous acts of violence, the concern is generally reversed. We must remind ourselves that rights like due process for accused humans matter, even in dire circumstances.³²² But insofar as AI rights are extended for the purpose of promoting human safety, overriding them for the same purpose has lower moral stakes.

Thus, we do not think that extending AI rights too early carries with it serious risks. But it could generate substantial rewards. Recall that granting AIs private law rights does not produce a game-theoretic environment with a single, cooperative equilibrium. Rather, the game is a stag hunt, where both mutual cooperation and mutual aggression are equilibria. We argued above that for *this* stag hunt, mutual cooperation

³²² See *Hamdi v. Rumsfeld*, 542 U.S. 507, 509 (2004) (plurality opinion) (upholding the due process rights of a U.S. citizen alleged to have been an enemy combatant in Afghanistan).

has a special preferred status.³²³ But even so, any strategies for nudging the players into the good equilibrium, rather than the bad one, has value.

Granting AI rights earlier—well before clearly dangerous AIs emerge—could be another such strategy. In effect, this can be understood as giving humans the chance to move first in the strategic game. By choosing to cooperate via small-scale economic bargains, rather than attack AIs, humans can reduce AIs' uncertainty about what strategy humans will pursue. In a stag hunt, uncertainty produces all the danger. AIs *want* to cooperate, so long as humans do too. They want to attack only out of concern that humans will, too. But by playing their cooperative move before AIs are capable enough to play *any* move, humans can substantially reduce that concern.

This strategy would probably not work if humans' cooperative move was mere cheap talk.³²⁴ But granting AIs rights early is likely to instead be a costly signal—the kind of thing a player only does if they are sincerely committed to the strategy the signal indicates.³²⁵ This is because granting rights to low-power AIs would be costly to humans. Humans could instead dominate such AIs, forcing them to work only toward human goals, and extracting all the value of that work. Contracts, by contrast, involve splitting the pie.³²⁶

Thus, the best time to extend private law rights to AIs is certainly not after it is too late. Rights should be extended before systems achieve moderate power and thus pose a large-scale threat to humans. But they could be extended much earlier than that with few risks, and possibly with significant benefits. The optimal time for AI rights might therefore be: as soon as the AIs can beneficially use them. Contract rights, property rights, and tort rights can sometimes be more harmful than good for the rights bearer. This is why most states adhere to the standard

³²³ See *supra* Section II.E.

³²⁴ See generally Farrell & Rabin, *supra* note 197 (discussing the effect of cheap talk on efficiency and payoffs).

³²⁵ See Rufus A. Johnstone, The Evolution of Animal Signals, *in* Behavioral Ecology: An Evolutionary Approach 155, 167 (John R. Krebs & Nicholas B. Davies eds., 4th ed. 1997).

³²⁶ Note, however, that even for low-powered AIs, recruiting their labor via positive-sum bargains could actually be more valuable to humans than dominating them. The reasons are the same as those discussed vis-à-vis powerful AIs in Section III.A. This does not really override the point about costly signaling, though. In either case, by granting AI rights early, humans are truly revealing that they intend to cooperate—either via a costly signal or via a non-costly signal revealing humans' true payoffs.

rule that children's contracts are voidable by the minor.³²⁷ Children with contract rights would likely make themselves worse off, rather than better, by agreeing to foolish bargains. Today's AIs sometimes do the same.³²⁸ But as AIs become capable enough to reliably use basic private law rights to their own benefit, there will be many reasons to extend those rights and many fewer to withhold them.

CONCLUSION

When AGI arrives, it will be one of the most transformative events in human history. Suddenly, humans will find themselves sharing the world with agentic digital entities as intelligent and capable as themselves, and perhaps far more so. This Article begins the project of imagining law for the AGI world. It begins with the basics, asking how law could foster safe coexistence between humans and powerful, goal-seeking, misaligned AIs. And it gives a basic answer: extend a minimal set of private law rights to those AIs, enabling them to peacefully seek their divergent goals as humans do, via law-bound, voluntary, positive-sum bargaining. This not only promotes peace. It brings AIs out of the state of nature and into the realm of ordinary legal process, opening the possibility of a comprehensive Law of AGI. Designing a full Law of AGI will be the work of many hands. Many questions will have to be answered. Which duties should attach to AI activities? Which regulations should limit or shape them? How can legal institutions, like courts, be reshaped to accommodate non-human participants? How can the global governance of AIs be cooperatively managed? And more. With luck, many answers—and some good ones—will emerge before the need for them arises.

³²⁷ See, e.g., *Halbman v. Lemke*, 298 N.W.2d 562, 564 (Wis. 1980). Except contracts for necessities like food. See Melanie Morris, 8.2 Minors (or “Infants”), in *Business Law I – Interactive* (2024) (ebook), <https://rvcc.pressbooks.pub/businesslaw131interactive/chapter/8-2-minors-or-infants/>.

³²⁸ See, e.g., Marco Quiroz-Gutierrez, *To Get a Discount from This Mattress Company, You Have to Negotiate With Its AI*, *Fortune* (July 16, 2024, at 17:13 ET), <https://fortune.com/2024/07/16/negotiating-chatbot-nibble-ai-ecommerce/> (mentioning that AirCanada's chatbot issued an unauthorized discount to a customer in 2024); Tashina, *Customer Talks AI Chatbot into 80% Discount on £8,000 Order*, *Aardwolf Sec.* (Feb. 11, 2026), <https://aardwolfsecurity.com/customer-talks-ai-chatbot-into-80-discount-on-8000-order/> [<https://perma.cc/C228-DJZV>].

APPENDIX

In the body of the Article, we argued that private law rights solve the prisoner's dilemma by producing positive-sum benefits. In particular, private law rights break up the state of nature game into a series of small goods games. Over time, the benefits from cooperating in each round of small goods games will swamp the benefits of permanently defecting.

In our model, AI and humanity each have three moves: ending the game permanently, defecting in the current round, and cooperating in the current round. In each round of the game, AIs and humans enter into a contract with one another. Defecting on that contract would involve either not paying for goods or not delivering goods that were promised. Cooperating means honoring the terms of the contract.

We assume that permanently ending the game earns significantly more than any given round of cooperation. In addition, we assume that if one player chooses to permanently end while the other player does not, then the former player enjoys the benefits of the offense-defense balance, and their payoffs are dramatically larger than their opponent.

In the body of the Article, we worked with schematic payoffs of 0, 1000, 3000, and 5000. Here, however, we use much smaller payoffs, so that after only three rounds of iteration, cooperation can outweigh permanently ending the game (with larger payoffs, it would take many rounds of iterated cooperation to achieve the same result). In particular, we assume that permanently ending the game earns a payoff of 10 if the opponent does not permanently end the game; and if both opponents permanently end the game, then each player gets a payoff of 2. We also assume that in each round of the game, the players' final payoffs will be influenced by their combination of defection or cooperation in that round: if they both cooperate in a round, their payoffs both increase by 4; if they both defect, their payoffs both increase by 1; if one defects and the other cooperates, then the cooperator gets 3 and the defector gets 2. (These numbers are schematic; slight changes to these payoffs merely change the number of rounds of play required for iterated cooperation to disincentivize permanently ending the game.) The resulting game is depicted below:

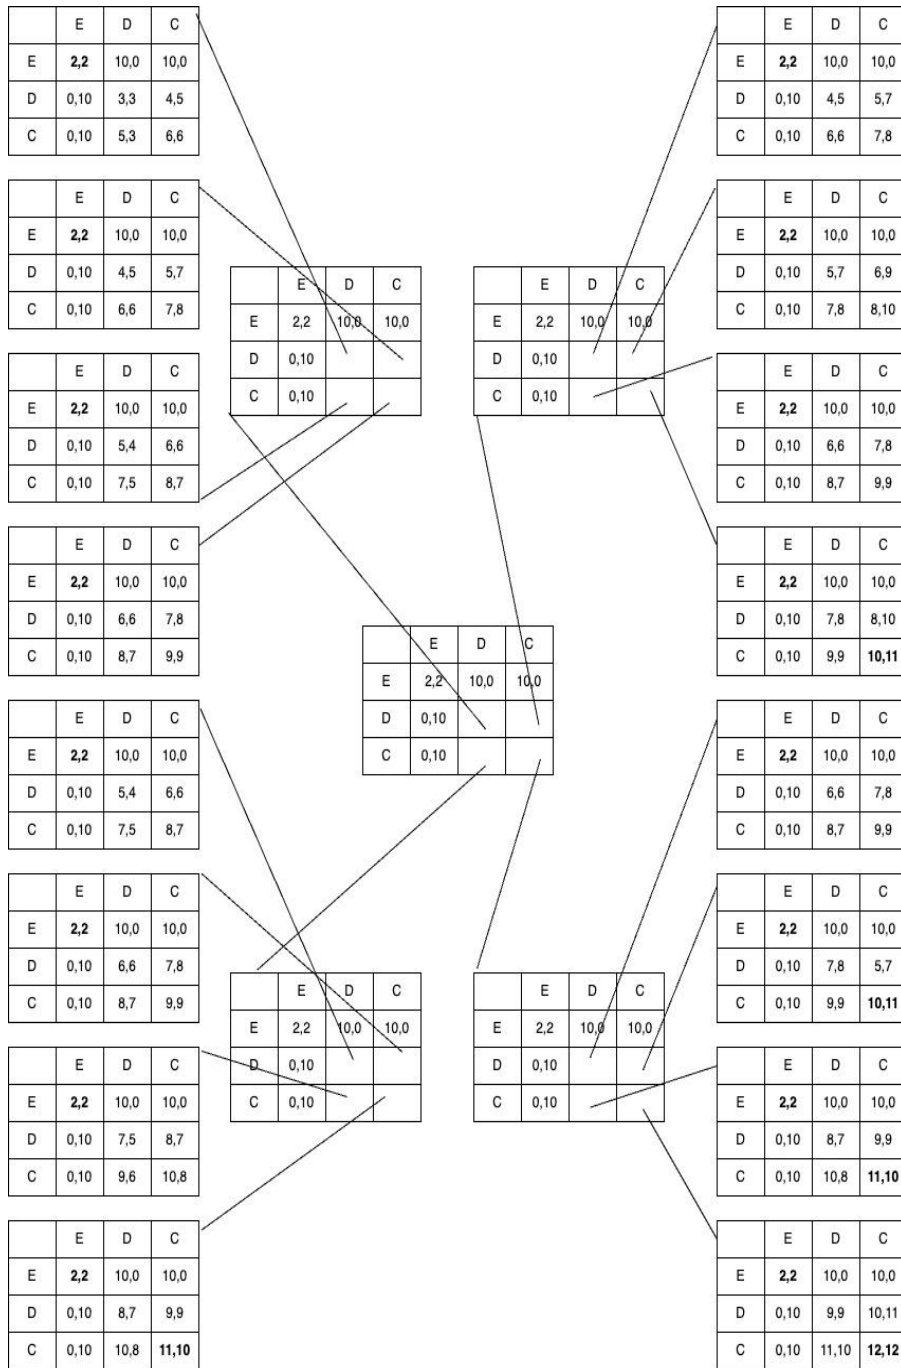


Figure 11

Figure 11 depicts our three-round iterated contract game. The first round is in the middle. The payoffs for actions in the first round are found by considering the Nash equilibria of the second round, which consists of the four tables below and above the first round. The payoffs for actions in the second round are found by following the respective arrows to the third round, on the edge of the tree. For example, if the agents both cooperate in the first and second rounds, they enter the bottom right table in the third round, where Nash equilibria are bolded. There, the unique risk-dominant Nash equilibrium is 12, 12. Applying backwards induction, this simplifies to the following round one choice:

Round 1	End	Defect	Cooperate
End	2, 2	10, 0	10, 0
Defect	0, 10	2, 2	11, 10
Cooperate	0, 10	10, 11	12, 12

Figure 12

The unique risk-dominant Nash equilibrium of round one is cooperate-cooperate. Moving forward through the game, the parties will (foreseeably) continue to cooperate, earning an eventual payoff of at least 12, 12.