

Open Global Investment as a Governance Model for AGI¹

(2025) Nick Bostrom

[*Working paper*, version: 1.15 (October 2025). (First version: July 2025)]

www.nickbostrom.com

ABSTRACT

This paper introduces the “open global investment” (OGI) model, a proposed governance framework for artificial general intelligence (AGI) development. The core idea is that AGI development could proceed within one or more corporations in a context that (a) encourages wide international shareholding, (b) reduces the risk of expropriation, (c) implements strengthened corporate governance processes, (d) operates within a government-defined framework for responsible AI development (and/or a public-private partnership), and (e) includes additional international agreements and governance measures to whatever extent is desirable and feasible. We argue that this model, while very imperfect, offers advantages in terms of inclusiveness, incentive compatibility, and practicality compared to prominent alternatives—such as proposals modelled on the Manhattan project, CERN, or Intelsat—especially in scenarios with short AGI timelines.

Introduction

Various governance models have been proposed for projects that aim to develop and deploy transformative AI.^{2,3} Some have favored an American-led “Manhattan project for AI”—a vision in which the U.S. competes with an analogous project in China in a battle for perpetual cosmic hegemony.⁴ Others have proposed an ownership and governance structure consisting of a U.S.-based nonprofit or public benefit corporation.⁵ Others still have explored more internationally cooperative models, such as a “CERN for AI” or an “Intelsat for AI”.⁶

¹ For comments and discussion I’m grateful to Renan Araujo, Guive Assadi, Nathan Barnard, Eden Beck, Nick Beckstead, Haydn Belfield, Catherine Brewer, Joe Carlsmith, Alan Chan, Tim Chan, James Coates, Aleya Cotra, Allan Dafoe, Max Dalton, Max Daniels, Tom Davidson, Oscar Delaney, Lukas Finnveden, Peter Gebauer, Ryan Greenblatt, Rose Hadshar, John Halstead, Hauke Hillebrandt, Holden Karnofsky, Will MacAskill, Sam Manning, Fin Moorhouse, Avro Muñoz, Toby Newberry, Tina Oberoi, Abi Olvera, Toby Ord, Liam Patell, Zershaaneh Qureshi, Anders Sandberg, Charlie Steiner, Daniel Susskind, Audrey Tang, and an anonymous referee.

² By “transformative AI”, we shall mean, roughly, artificial general intelligence (AGI) that is capable enough to have the potential, if fully deployed, to have profound cross-sectoral impacts on the economy, labor market, and national security.

³ For one compilation of proposed international AI institutions, see Maas & Villalobos (2023).

⁴ Aschenbrenner (2024)

⁵ Juijn et al. (2024). See also various discussions of OpenAI’s and Anthropic’s public benefit corporation structures.

⁶ MacAskill & Hadshar (2025)

Governance models may emphasize different objectives, such as democratic oversight, national security, international cooperation, legitimacy, economic efficiency and dynamism, equitable benefit sharing, AI safety, free enterprise, wise stewardship, scientific advancement, operational security, or responsible deployment. No governance structure can fully achieve all these goals. Transformative AI will, moreover, pose special challenges beyond those present in most other policy contexts—including the risk of misaligned or misused AI causing existential catastrophe, and the potential for extreme concentrations of power. Any governance design will therefore be subject to criticism for falling short on some reasonable desiderata. The choice will be one between imperfect options.

Furthermore, as we approach a potential intelligence explosion, governance structures will come under great strain. This is in part because the stakes become extraordinarily high and in part because the context in which they operate will undergo rapid and profound change.⁷ Thus, it is not enough that a given structure *would* be satisfactory if it governed the development of transformative AI: it is also needful that the structure be robust enough to maintain its integrity throughout the process, and incentive-compatible enough to retain the support of relevant actors—and to be adopted in the first place.

This paper outlines a paradigm that we call the “open global investment” (OGI) model. We then assess how OGI meets various governance desiderata compared to alternative models, and tentatively conclude that it appears relatively attractive.

The open global investment model

Core features and variations

In the OGI model, AGI development is led by one or more ventures that are widely open to international investment and that operate under a government-defined regulatory framework, buttressed by enhanced government assurances and assistance. The status quo already roughly approximates this model, though we could move closer to the ideal. If we favor the OGI model, it could mean that we should strive to move closer to its ideal form; and perhaps more importantly, it could mean that we should resist proposals (like nationalization) that would remove us farther away from it.

This core idea can take different forms depending on whether there is a single lead AGI company (“OGI-1”) or several (“OGI-N”); on whether the companies are domiciled in the United States (“US-OGI”); and on what supporting features and structures are put in place.

The following analysis adopts a somewhat U.S.-centric perspective, reflecting current realities in AGI development and the surrounding Western discourse. However, the OGI model itself is geographically neutral—it could in principle be implemented by any technologically capable nation as host, or by multiple different countries as hosts for different AGI ventures.

⁷ Bostrom (2014), Dafoe (2018), Karnofsky (2023), Bengio et al. (2024)

Example (US-OGI-1)

A fairly full-fledged version of the US-OGI-1 version (a U.S.-domiciled single AGI firm) could look like this:

- (a) A publicly traded corporation (“AGI Corp”) either already exists or is floated, domiciled in the U.S. AGI Corp could be a standard corporation or a Delaware public benefit corporation.
- (b) Many people buy shares, including private individuals both in the U.S. and around the world. Sovereign wealth funds and other investment vehicles also buy significant chunks. If there is an IPO, it is structured such that initial ownership is widely distributed.
- (c) There may be multiple share classes with different voting rights. For example:
 - Class A shares: profit participation and 1 vote/share
 - Class B shares: profit participation and 10 votes/share
 - Class C shares: profit participation and 1,000 votes/shareThese numbers are merely for illustration. The key point is that profit participation and voting rights can be partially separated—though ideally both should be widely distributed. For instance, everyone worldwide might be able to buy Class A shares, while Class B and Class C shares could be restricted to citizens of countries that agree to a responsible AI framework; and class C shares could be allocated as founder shares to key parties.⁸
- (d) Foreign governments and their citizens, including major powers that might otherwise be considered strategic rivals, are permitted and indeed encouraged to buy shares in AGI Corp (perhaps up to some high limit, such as 20% for any single country).
- (e) To help AGI Corp be the first to develop AGI and to have a big lead, cooperating governments could take various actions, such as:
 - (i) Provide subsidies, tax breaks, regulatory waivers, etc.
 - (ii) Pass regulation that impedes competitors (such as restricting access to advanced AI chips, etc.)
- (f) AGI Corp could also buy up competitors and merge them into its own efforts. Given the corporation’s special status and official sanction, and the headwinds its competitors can expect to face, it might have a low cost of capital, enabling it to make such acquisitions. The U.S. government (USG) could agree to waive antitrust enforcement.
- (g) The charter of the corporation might be created or amended to strengthen governance procedures beyond those typical of Fortune 500 companies. For instance, the board could meet more frequently and have various resources and provisions for internal oversight.
- (h) Ideally, some significant portion of the corporation’s shares are acquired by independent organizations and not-for-profits whose remit is to benefit all of humanity and to serve other moral missions. This might occur because the original IPOing corporation already had such entities on its cap table. Philanthropic organizations can also buy additional shares on the open market.⁹

⁸ A limiting case of shares of differential voting power is a so-called “golden share” that gives the holder veto rights or the right to make a specified class of decision. When such a share is held by a government, this can amount to a form of soft nationalization (which would instantiate the OGI model to a lower degree than when ownership is fully distributed according to commercial principles).

⁹ The OGI model could also facilitate the implementation of proposals such as that of Epstein (2025), which calls for the creation of sovereign wealth funds for transformative AI; see also arguments for debt-financed

- (i) AGI Corp is subject to the laws of the land where it is domiciled. Assuming a US-OGI model, the USG thus retains the ability to pass laws or take regulatory action to prevent AGI Corp from deploying products deemed unsafe. It could even shut down development of more advanced models until certain safety and security standards are met. The USG could also insist on vetting technical personnel or other measures to counteract espionage.
- (j) Various mechanisms are employed to make it legally as difficult and costly as possible for the USG or any other government to expropriate AGI Corp—using an assortment of legal and structural devices employed by other multinational corporations for that purpose.
- (k) Additionally, leaders of the USG (and other governments) should, at its inception and periodically thereafter, signal support for the corporation and reaffirm commitments to respect its property rights and independence (e.g. pledging no confiscatory taxes or nationalization). Ideally these commitments would be embedded in laws and treaties, but informal pledges or expressions of support would be better than nothing.
- (l) To the extent practicable from a security perspective, a significant portion of AGI Corp’s data centers could be distributed across multiple countries and jurisdictions, to further increase the difficulty of expropriation.

In versions with multiple AGI companies (OGI-N), provisions (e) and (f) would not apply, resulting in a situation closer to today’s status quo. This paper remains agnostic on the relative merits of the single-company model (OGI-1) versus the multi-company model (OGI-N).

The regulatory context of operation

The scenario sketched above gives only half of the governance picture: it outlines ownership and control of the AGI project(s). The other half is the set of rules governments impose on how the AI sector operates.

Myriad governance challenges will arise as AGI technologies permeate the economy and society—protecting consumers, maintaining political integrity, assisting displaced workers, preventing abuse, etc. These challenges will mostly fall to governments. In a US-OGI scenario, the U.S. government would bear primary responsibility for oversight—setting safety standards for advanced model development and deployment, establishing monitoring procedures, and so forth. Other countries would impose their own rules on how AI products are used domestically. The details of this regulatory landscape are not unique to the OGI model and lie outside this paper’s scope. In the ideal case, states might agree on an international framework for safe and responsible AI that enables cooperation and provides basic global safeguards, within which each country could then implement its own more specific regulations.

The AGI corporation (or corporations, in OGI-N versions) can also play a role in addressing these issues. For instance, it can choose not to offer products or services it believes would harm its shareholders’ long-term interests or violate its broader mission. A Delaware public benefit corporation charter would increase the scope for sacrificing profit in favor of ethical values, stakeholder interests, or the public good. Even without that, U.S. corporate law gives ordinary

public debt in Casey et al. (2024), Bernard and Hillebrandt (2025), and Svarc and Hillebrandt (2025), which argue for debt-financed public investments in the AI sector.

for-profit boards considerable latitude in practice to consider factors other than immediate profit maximization.

Motivations

The rationale behind the OGI model is threefold:

- (1) To leverage well-entrenched norms, laws, and institutions related to the protection of property rights. These may be more robust and reliable than some novel ad hoc governance scheme concocted specifically for AGI.
- (2) To distribute ownership and control widely, including among powerful actors (such as capital owners, political representatives, etc., both in the U.S. and overseas). This serves important functions:
 - (i) It gives powerful actors personal incentives to uphold the corporation's property rights rather than pursue some radical expropriation or nationalization push that risks unraveling the property rights framework under which their interests are protected.
 - (ii) It offers international rivals an alternative to resentment and overt hostility: they are instead welcome to invest and participate in both profits and control over the first AGI. (On the margin, this may also reduce their willingness to engage in a risky race to AGI.)
 - (iii) It promotes a somewhat wider and globally more equitable distribution of benefits and influence compared to models in which ownership and control are even more narrowly concentrated.
- (3) To present a path that has a realistic chance of being implemented, given current actors' motivations, political and geopolitical constraints, short AI timelines, etc.

Roles and powers of various actors

We now consider the main actors in an OGI scenario—the AGI corporations, the U.S. government (as host), other governments, and citizens—and the roles and powers each would have.

AGI corporation(s)

In addition to managing its internal affairs and research processes, an AGI corporation can choose which products it develops, whom it sells those products to, at what price, and with what constraints as to their use.

It is assumed that the corporation will be significantly driven by a profit motive in making these decisions. However, its products may so profoundly affect the world that they impact its shareholders in ways beyond direct profits. Therefore, the corporation would plausibly be influenced by shareholders' non-monetary interests to a greater extent than a typical company. If the AGI corporation is chartered as a public benefit corporation, its legally defined public mission could also significantly shape its choices.

Note that in the OGI-N version, the power of any one corporation to shape the impacts of AGI is reduced because of competitive pressures. One corporation might refrain from developing a

product it believes would be harmful to society, but if this means foregoing profits then another corporation may fill the gap. Thus, if one believes it is important for both government and the lead AGI developer to have substantial influence over how AGI impacts unfold, this would be an argument for OGI-1—or at least for keeping the number of top AGI corporations very small.

The United States government

The USG (or another host government) can block AGI corporations from pursuing certain developments or from deploying or selling certain products, for example in order to protect the public from foreseeable harm. It can also insist on procedures to reduce the risk of espionage against a corporation or loss of its intellectual property.

The host government can also choose whether to take actions that help an AGI corporation, such as buying some of its stock, removing regulatory barriers, providing tax breaks and subsidies, or hindering its competitors.

In a fully realized OGI model, the USG would also give up some of the options it currently has to commandeer or expropriate companies. For example, it could commit not to nationalize the AGI firm, not to impose special punitive taxes on it, and not to force foreign shareholders to divest.

Other governments

Other governments can have influence via several channels.

First, they can buy shares in AGI corporations, which gives them voting rights and profit participation.

Second, they can regulate AGI products and services within their own jurisdictions, and potentially prevent a corporation from doing business within their borders unless it complies with all local laws.

Third, if the corporation is a public benefit corporation, they could sue in U.S. courts if the corporation fails to sufficiently advance its public benefit mission.

Fourth, they could exert influence on the USG (and hence indirectly on American AGI corporations) through normal diplomatic and economic channels.

Fifth, if formal or informal multinational agreements have been entered into (e.g. responsible AI policies, assurances of respect for the AGI Corporation's autonomy, etc.), they may have institutional or legal recourse provided by those agreements.

Sixth, if data centers or other corporate assets have been located on their territory, they may as a last resort be able to seize those assets if the U.S. were to renege on its side of the bargain. (One could also imagine novel technical arrangements, such as an n-out-of-m multisignature mechanism to remotely disable critical AI hardware if certain conditions are met.)

Citizens

Citizens (both in the U.S. and around the world) would have at their disposal all the ordinary ways of influencing their governments, and thereby, indirectly, AGI corporation(s): through appeals to officials, voting, protests, lobbying, etc. They can also use familiar methods of influencing corporations: through their purchasing behavior, whether they choose to work for the corporation, taking actions that affect the corporation's brand image, etc.

In the OGI model, citizens additionally have the option to acquire shares in the AGI corporation, which would let them directly participate in the corporation's governance. This ability is, of course, distributed unequally due to wealth disparities both within and between countries. In particular, populations in lower-income countries would have very little influence per person. While this inequity is unfortunate, one must ask: compared to what? Under most realistic alternative models (say, a U.S.-only Manhattan Project or a privately held U.S. company), those same people would plausibly have even *less* ability to influence outcomes. Furthermore, national governments, sovereign wealth funds, religious organizations, charities, or other NGOs could buy shares on behalf of people too poor to participate directly; those institutions could then use internal democratic processes to decide how to exercise their shareholder power on behalf of their constituents.

Representativeness and fairness

It is tempting to assume that a government-controlled initiative would be more representative and inclusive than a privately controlled initiative. In the present case, that assumption is questionable. Even a large country like the U.S. has only about 4.2% of the world population (and 26% of world nominal GDP).¹⁰ In a U.S. nationalized project, 95.8% of the world population (and 74% of the world economy, or 85% if we adjust for purchasing power) is excluded from ownership and influence.¹¹ Even those figures optimistically assume that a Manhattan Project-type intervention would remain under effective democratic control in the U.S., which is far from a given in some scenarios.¹² By contrast, in the ideal-type instantiation of the OGI model, everybody with enough financial assets to buy stocks (and access to the relevant financial institutions) would have an opportunity to participate; and large swaths of the world population may in fact participate, either directly (via private investments) or indirectly (e.g. via sovereign wealth funds, pension funds, index funds, etc.).

Global wealth is very unequal. Almost half of global wealth is owned by the richest 1% of individuals.¹³ But a substantial fraction of global wealth and stocks is also owned by a large number of moderately wealthy individuals. We can put some rough numbers on this. The Gini

¹⁰ U.S. Census Bureau (2025), IMF (2025)

¹¹ IMF (2025)

¹² Congressional oversight of the Manhattan Project was extremely limited—only a handful of congressional leaders knew of the project's existence, and even they received minimal information (Rhodes 1986). Vice president Harry Truman was not aware of the project until after President Franklin Roosevelt's death (Wellerstein 2021). And AGI, which would allow widespread automation of intelligence analysis, propaganda, policing, and military capabilities, could lend itself to coups or other paths to extreme concentrations of power.

¹³ UBS (2023)

coefficient of global wealth is approximately 0.89.¹⁴ The Gini coefficient of global stock ownership, while not precisely measured in standard databases, is plausibly higher, perhaps approaching 0.90-0.92, given that financial assets are more concentrated than overall wealth.¹⁵ A U.S.-nationalized program—even if we assume that all U.S. citizens have exactly equal stakes and influence in it (which seems optimistic) while non-U.S. citizens have no stake or influence—gives a global distribution with a Gini coefficient of approximately 0.96.¹⁶ That is to say, by at least one common measure of inequality, the U.S. nationalization model appears *more* unequal than the open global investment model in terms of the distribution of financial gains and influence.¹⁷

This initial gap in favor of the OGI model widens substantially when we consider that some fraction of the profits of the AGI sector would be taxed by governments both in the U.S. (or wherever the AGI corporations are domiciled) and in countries where shareholders report their foreign-source dividends and capital gains: tax revenues that would be used to provide services or payouts to vast numbers of citizens around the world. The gap may widen further still if we consider OGI-N versions of the model, in which a substantial fraction of potential profits may be competed away rather than captured by a single monopoly actor (as in OGI-1 or U.S. nationalization). Furthermore, to the extent that we place significance on the geographic dispersion of benefits, the OGI model does much better in that regard.

Arguably, therefore, an open and internationally participatory investment structure is fairer and more globally representative than any structure that is wholly owned and controlled by a single country.

It is possible to conceive of organizational arrangements that would be more inclusive than OGI. For example, one might imagine an AGI project run by the United Nations, or by some new idealized version of the UN set up for the purpose. One major problem with such more idealistic constructions is that they may not be incentive-compatible for current power-holders and potential funders, and so may have little chance of being implemented. Many international governance organizations are also often slow and may not be capable of running a globally competitive AGI project. There is currently no frontier AGI effort operated by any international governance organization, nor any that is operated by a national government (at least as far as is publicly known). The timeline for reaching agreements to set up a perfectly globally inclusive effort might be long—especially one that is embedded in international treaties—and the resulting

¹⁴ According to the Credit Suisse Global Wealth Databook 2022 (the most recent available with detailed methodology), the global wealth Gini coefficient was 0.889; see Davies et al. (2022). This figure has remained relatively stable in recent years.

¹⁵ This estimate is based on the fact that financial assets are somewhat more concentrated than overall wealth. In the U.S., for instance, the top 1% own 54% of public equities while owning approximately 35% of total wealth.

¹⁶ This calculation assumes that the U.S. population (4.2% of global total) owns 100% of the “asset” (i.e. control of the hypothetical U.S.-national project via voting rights) while the remaining 95.8% owns 0%. Using the standard Gini formula $G = 1 - 2B$, where B is the area under the Lorenz curve, gives $G \approx 0.96$.

¹⁷ A significant portion of voting power in public companies is concentrated among large institutional asset managers. For instance, BlackRock, Vanguard, and State Street collectively control approximately 25% of voting rights in S&P 500 companies (Fichtner et al. 2017). But this concentration might be compared to representative democracies where citizens delegate power to elected officials—in both cases, many individuals choose where to allocate their resources (investments or votes), but actual decision-making is exercised by a smaller number of actors purportedly acting on their behalf.

organizational structure would be relatively untested compared to the international corporate property rights regime upon which the OGI proposal is based. An additional difficulty that an internationalized project would confront is how to achieve operational and information security, particularly if one envisions that its research and management staff would be drawn from nearly 200 countries.¹⁸

Military applications and foreign competitors

The distinction between civilian technology and military technology is already often difficult to draw. AGI may complicate the picture further and introduce novel areas of application where precedent offers limited guidance as to classification. High levels of AGI capabilities, even if not specifically designed for military applications, may be capable of greatly contributing to military success—for example by helping with military planning and logistics, operating drones and robots, conducting rapid military R&D, and executing cyber and information operations. Furthermore, the potential for extremely rapid economic growth (such as is predicted in some models of AGI development) could itself have destabilizing effects on a country's security and geostrategic position, since any sufficiently large gap in economic development would likely also cause a divergence in military power.¹⁹

For this reason, the host government—the USG, if we assume the US-OGI model—may seek to restrict the sale of AI products and services, thereby preventing other nations from participating fully in shaping the emerging machine intelligence era. Furthermore, the USG may seek to restrict sales of top-tier AI capabilities to its own citizens, perhaps on grounds of dual-use concerns.

Fearing such a scenario, rival states with the resources to compete with the U.S. in AI development (such as China) may pursue their own independent national and/or commercial AGI projects. They may do so even if they were significantly invested in American AGI corporation(s) and even if they were reasonably confident that the U.S. would honor its commitments to respect corporate autonomy and property rights. These dynamics reflect understandable concern on all sides about technological dependence and national security.

When considering how much of an objection this is to the US-OGI model, it is important to bear in mind that the most relevant comparison is to alternative models that have some realistic prospect of being implemented, not an idealized world order in which all of humanity voluntarily comes together in perfect friendship and harmony to wisely work together for the common good of all. We should remain open to and favor the latter whenever we have opportunities to do so; but simultaneously it is prudent to think through and develop second-best options in case the preconditions for the first-best option should prove unattainable.

Thus, if we compare the OGI model to a “U.S.-led Manhattan Project for AGI” model, or to a model in which AGI is developed by a privately held corporation, we can observe that in those cases too would rival powers have incentives to pursue their own AGI projects. In fact, their incentives to pursue independent projects in those alternative models would be *stronger*. In the basic US-OGI model, although it would not guarantee the maintenance of their geostrategic position in

¹⁸ Cf. Bostrom (2017)

¹⁹ Cf. Erdil & Besiroglu (2024)

scenarios where the U.S. restricts exports of advanced AGI capabilities, they—and also individuals within their economic and political elites, who may have personally invested in the project—would at least have a stronger prospect of being able to participate in the economic upside of AGI development than they would in alternative models. They would also have at least *some* say, via their shareholder voting rights, in how the technology is developed and which products are offered for sale.

At present, this degree of participation is certainly not enough to make U.S. rivals abandon their own AGI efforts. But it might, at the margin, reduce their urgency or scale. Moreover, one can imagine circumstances where the prospect of share gains from AGI Corp’s success could conceivably tip the balance in favor of rivals stepping back. For example, if it became overwhelmingly clear that racing to be first with AGI would pose an extremely high level of existential risk (due to misaligned superintelligence), even great-power competitors might consider unilaterally slowing down and relying on the shared project.

If stronger assurances to non-host countries are both desirable and possible, then the basic OGI model could be augmented with additional mechanisms, such as international arms control agreements, non-aggression pacts, etc. The OGI model should be as amenable to such international assurances as e.g. a private development model or a Manhattan Project model. It is conceivable that models that focus on concentrating AGI developments within a jointly operated international project would have an advantage in this regard. But they have other difficulties (see Appendix 4); and they would still confront the issue of how to deal with competing commercial or national projects.

Private control in relation to national security

When AI becomes a critical factor for national security, it is plausible that a host government, i.e. the USG in our main example, would want to ensure that it has full access to the most capable technology that AGI companies can produce.

Today, most advanced U.S. defense materiel is developed and produced by privately owned (and usually publicly traded) corporations, such as Lockheed Martin, RTX Corporation, Boeing, and Northrop Grumman. Although the USG typically has no direct ownership stake in these companies, it nevertheless exerts a great deal of control over their activities—more so than over most other private-sector enterprises. This is partly because the USG is the primary buyer of their products and so has monopsony power, and partly because it tightly regulates the defense sector—through the issuance of security clearances; controls of what technology and information can be shared with foreign nationals (ITAR—International Traffic in Arms Regulations); embedding of government employees at contractor facilities to monitor production, quality, and compliance (DCMA—Defense Contract Management Agency); the ability to invoke the Defense Production Act (through an executive order) to force companies to disclose capabilities and to accept and prioritize “rated orders” from the government (in extreme cases, these can include price controls); etc. Foreign investment is permitted but controlled through the CFIUS (Committee on Foreign Investment in the United States) review process, which can block foreign attempts to acquire a controlling interest (or sometimes even a significant minority stake) or impose mitigation measures such as Special Security Agreements or FOCI (Foreign Ownership, Control, or

Influence) Mitigation that may, for example, cap foreign ownership at a certain percentage or restrict the voting rights and information access of foreign investors.²⁰

Given the government's formidable array of tools, investors may reasonably worry that they would be at risk of de facto expropriation. Defense contractors have been fairly successful at mitigating this risk by investing heavily in lobbying, "revolving door" hiring, and distributing factories across key political districts. AGI firms might try similar tactics, but they may become more vulnerable than defense firms to nationalization or expropriation—especially if they generate prodigious windfall profits from an intelligence explosion or become extremely central to national security. Such concentrated wealth and power would be a tempting political target. Moreover, unlike defense contractors, AGI firms probably won't create vast numbers of jobs (relative to their market cap); indeed, they may be blamed for massive job losses via automation. An AGI company might fend off outright shutdown by strategically invoking the specter of Chinese competition, but if it wants to protect its autonomy and profits in the long term, it may need broader buy-in—particularly among political and economic elites. The OGI model helps achieve that by giving a wider range of people a financial stake in the success of AGI companies.

Speed gap between technology and regulatory processes

One distinctive feature of AGI development is the speed at which it may unfold and the range as well as novelty of the issues it will force policymakers to confront. Especially in "fast-takeoff" scenarios, in which escalating AI capabilities quickly ascend through the human range and beyond into increasingly exalted levels of superintelligence, any governance model that relies entirely on traditional regulatory processes will fail to keep pace with the technology. Drafting, vetting, and instituting new regulations takes time. In many cases, regulation is anyway not the appropriate instrument with which to govern a fast-evolving situation presenting unique, speculative, or highly technical judgment calls. Just like a military campaign cannot be managed by issuing regulatory codes, but must rely on field commanders exercising judgment in response to rapidly changing circumstances, so too will many aspects of the development and deployment of advancing AGI require executive discretion by competent decision-makers. The faster things move, the more this will be the case—and the more the outcome will depend on the values and competence of the specific individuals who occupy key positions at the relevant time.

If and insofar as effective government oversight is desired, therefore, the OGI model may require an unusually close coupling between the company and the government. Short of outright nationalization, this could be achieved in various ways. For example, the AGI firm might voluntarily agree to regular meetings with a dedicated government task force (with technical experts and senior officials—perhaps even involving congressional leaders and the president when the stakes are high). More formal measures could include embedding government representatives within the company or establishing an oversight board to continuously monitor its

²⁰ In the basic OGI model, the U.S. would retain whatever options it currently has for trying to hinder geopolitical rivals from competing with or overtaking it in the AI race. For example, it could impose export restrictions on AI chips and semiconductor fabrication equipment, and encourage its allies and trading partners to impose similar restrictions on their exports. It might also use the right to purchase Class C shares (which offer greater voting power) as a bargaining chip to encourage countries to cooperate with it in imposing such restrictions, or in implementing other measures that would form part of a broader framework for responsible AI.

actions. At the extreme, one might create a public-private partnership structure. Notably, if the government becomes deeply involved, the OGI model may need extra safeguards to reassure shareholders that their interests will still be protected (see Appendix 2 for further discussion).

Transition to post-AGI governance

It is very difficult to imagine in concrete detail what a good governance system would look like for a fully post-AGI world—one in which, perhaps, biological humans and animals coexist with digital minds spanning an enormous range of capability levels, architectures, functions, goals, organizational forms, and moral statuses, and in which technology may approach the ultimate physical limits.²¹ The OGI model is primarily intended as an option for governance during an intermediate stage of development—between where we are now and the emergence of full-fledged superintelligence. Beyond that point, fundamental changes in how we think about governance may become necessary; but by then it will hopefully be easier to see what the situation requires, and decision-makers may have access to superintelligent AI advisors to guide them in further reforms.

The OGI model remains relevant for the longer term insofar as it might (a) help us get there with a minimum amount of negative-sum conflict, and (b) shape the initial conditions for whatever comes after. If successful, OGI could encourage a trajectory where some rivalries are partly defused through cooperative economic structures, and where influence and profits are somewhat more broadly shared than under the most realistic alternative models.

Discussion and conclusion

One could conceive of the OGI model as an “ideal type” (in the Weberian sense, not necessarily in the sense of “first-best”), which reality might more or less closely approximate.

The current situation (as of June 2025) already provides some of the benefits that a more fully realized OGI model would offer. Alphabet and Meta, although controlling stakes are closely held, are publicly traded, as are many other companies that are key suppliers to AI developers (including the dominant designer of AI chips, along with many links in the semiconductor supply chain—all the leading foundries and many key equipment makers—as well as the largest data center providers). As for private AI companies such as OpenAI, Anthropic, xAI, and SSI Inc., it is likely that large institutional investors in countries that are friendly to the U.S. would be able to negotiate stock purchases with these companies if they are willing to offer favorable terms. Even individuals and countries that are not regarded as friendly to the U.S. could gain diluted exposure through shares in publicly traded companies that have invested in some of these private equities (such as SoftBank, BlackRock, etc.).

If one is focusing on scenarios that lead to an astronomical AI bonanza, one might think that even a highly diluted exposure to the upside would be sufficient to provide for any plausible resource-satiable individual or national priority. Whether you end up owning 1,000 galaxies or a paltry 10 solar systems, you could still saturate most non-positional personal requirements with many orders of magnitude to spare. This may require no investment on your part at all: if at least

²¹ Bostrom & Shulman (2020)

one slightly generous person ends up with 1,000 galaxies and is willing to spare 0.1% of their resources on philanthropy to the existing human population, then every individual could get 10 solar systems. Among the current significant investors in leading AI companies, there is plausibly at least one slightly generous person. Moreover, some of the private AI companies are partly owned or controlled by not-for-profit entities with a mission to benefit humanity; and in these types of scenarios, those charities could become able to underwrite massive grants that would distribute the benefits of an intelligence explosion widely. The U.S. (or other nations whose companies would capture some non-trivial fraction of the bonanza) could tax a portion of the profits and then via foreign aid programs—even if they amount to a fraction of one percent of their GDP—fund an equivalent global handout.

For these reasons, one might argue that there would be little to gain from moving closer toward the ideal type of the OGI model. However, this astronomical bonanza scenario rests on suppositions that might not hold—for example, it might not hold on the simulation hypothesis or if there are other claimants to the apparent cosmic endowment.²² There may also be uncertainty as to how the behavior of states, charities, or currently philanthropic individuals might change after such profound transformations have taken place. Furthermore, many actors also have positional preferences that are not resource-satiable. And even if in fact some actor's long-term preferences would be well satisfied in these scenarios, they may not seem “realistic” enough to offer that actor much comfort in the run-up to superintelligence (and dissuade them from desperate moves aimed at increasing their portion of potential upsides); whereas a more legible equity position in the AGI sector might offer more reassurance—along with the satisfaction of watching “numbers go up” during the transition period. Greater assurances that property rights will be protected would also be helpful. Beyond this, there is the issue of participation in the governance of the entity that spearheads the development of AGI. In a fully realized OGI model, actors from anywhere in the world for whom this is important would have a peaceful and lawful option to gain some degree of such participation by buying voting shares in AGI Corp.

Perhaps more importantly, assessing the merits of the OGI model matters not only because it could encourage efforts to move closer to its ideal form, but also because it could motivate avoiding steps that would move further away from it. If someone advocates that the U.S. should nationalize AGI development on ethical grounds, such as inclusivity, equity, and democratic legitimacy, it is relevant to point out that such a move would completely exclude 95.8% of the world population and all countries except one. If someone is promoting nationalization in order to avoid a race to the bottom in safety standards, it is relevant that the OGI model would allow foreign powers (and especially their elites) to participate as investors in a project based in a different country than their own rather than being forced to compete against it at all cost or resorting to desperate acts of sabotage.²³ If someone is concerned with the potential for coups, power grabs, or general destabilization, it is relevant that norms and laws around property ownership, investor rights, and corporate governance are comparatively well established and integrated with civilian society. And if someone proposes to create a new international organization to manage the development of AGI, it is relevant to compare such a construction to the OGI model in terms of political feasibility, time to fruition, informational and operational security, funding prospects, and likely levels of organizational efficiency.

²² Bostrom (2003, 2024)

²³ Armstrong, Bostrom, & Shulman (2016)

Recent international efforts highlight both the appeal and limitations of multilateral approaches to AI governance. The UN's 2024 "Governing AI for Humanity" report, for instance, proposes relatively modest initial steps—an international scientific panel, capacity-building programs, and coordination mechanisms.²⁴ While these may serve useful functions, they illustrate how far current international proposals remain from a comprehensive framework capable of meaningfully governing AGI development. Even if such proposals were eventually to evolve into something more substantial—such as an IAEA-style agency (analogous to how the International Atomic Energy Agency governs nuclear technology) or a complex ecosystem of distributed governance mechanisms—the history of international organizations suggests this could take many years or decades and would likely still face challenges of enforcement, agility, and technical competence. In the meantime, the OGI model offers a pragmatic alternative: it can be implemented immediately while remaining compatible with various forms of international coordination that might emerge—particularly those focused on standard-setting, monitoring, or rule enforcement, rather than ownership or direct operational control. The international framework for safe and responsible AI mentioned earlier might emerge gradually around OGI operations rather than requiring grand institutional designs from the outset.

References

- Achenbrenner, L. (2024) *Situational Awareness: The Decade Ahead* [online]
[<https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf>]
- Armstrong, S., Bostrom, N., & Shulman, C. (2016) "Racing to the Precipice: A Model of Artificial Intelligence Development" *AI & SOCIETY*, 31(2): pp. 201-206
[<https://www.fhi.ox.ac.uk/wp-content/uploads/Racing-to-the-precipice-a-model-of-artificial-intelligence-development.pdf>]
- Barnard, N. & Hillebrandt, H. (2025) "Leveraging Public Debt: A Self-Financing Hedge against Unemployment under Transformative AI". *Working paper*
- Bengio, Y. et al. (2024) "Managing Extreme AI Risks Amid Rapid Progress" *Science*, 384(6698): pp. 842-845
- Bostrom, N. (2003) "Are We Living in a Computer Simulation?" *Philosophical Quarterly*, 53(211): pp. 243-255
- Bostrom, N. (2014) *Superintelligence: Paths, Dangers, Pathways*. Oxford University Press
- Bostrom, N. (2017) "Strategic Considerations for Openness in AI Development" *Global Policy* 8(2): pp. 135–148 [<https://nickbostrom.com/papers/openness.pdf>]
- Bostrom, N. & Shulman, C. (2020) "Propositions Concerning Digital Minds and Society" (forthcoming in *Cambridge Journal of Law, Politics, and Art*, 2025)
[<https://nickbostrom.com/propositions.pdf>]

²⁴ UN (2024)

- Bostrom, N (2024) "AI Creation and the Cosmic Host" *Working paper* [https://nickbostrom.com/papers/ai-creation-and-the-cosmic-host.pdf]
- Casey, E., Roy, H. & Rockall, E. (2024) *Designing an AI Bond for Growth and Shared Prosperity in the UK* [online] [https://britishprogress.org/uk-day-one/designing-an-ai-bond-for-growth-and-shared-prosper]
- Dafoe, A. (2018) "AI Governance: A Research Agenda" *Future of Humanity Institute, University of Oxford* [https://cdn.governance.ai/GovAI-Research-Agenda.pdf]
- Davies, J.B., Lluberas, R., & Shorrocks, A. (2022) *Global Wealth Databook 2022*. Credit Suisse Research Institute, p. 141.
- Epstein, L. (2025) *Lead, Own, Share: Sovereign Wealth Funds for Transformative AI* [online] [https://ssrn.com/abstract=5343934 or http://dx.doi.org/10.2139/ssrn.5343934]
- Erdil, E. & Besiroglu, T. (2024) "Explosive Growth from AI Automation: A Review of the Arguments" [online] [https://arxiv.org/pdf/2309.11690]
- Executive Office of the President (1942) Executive Order No. 9066, 7 *Fed. Reg.* 1407
- Fichtner, J., Heemkerk, E., & Garcia-Bernardo, J (2017) "Hidden Power of the Big Three: Passive index funds, re-concentration of corporate ownership, and new financial risk" *Business and Politics*, 19(2): pp. 298-326
- IMF (2025) *World Economic Outlook Database, April 2025*. International Monetary Fund
- Juijn, D. et al. (2024) "CERN for AI: The EU's Seat at the Table" *International Center for Future Generations report*. [https://cfg.eu/wp-content/uploads/CERN_for_AI_FINAL_REPORT.pdf]
- Karnofsky, H. (2023) "Transformative AI issues (not just misalignment): an overview" [online] [https://www.cold-takes.com/transformative-ai-issues-not-just-misalignment-an-overview/]
- Maas, M. & Villalobos R. (2023) "International AI Institutions: A Literature Review of Models, Examples, and Proposals" *Legal Priorities Project, AI Foundations Report #1* [https://law-ai.org/international-ai-institutions/]
- MacAskill, W. & Hadshar, R. (2025) "Intelsat as a Model for International AGI Governance" [online] [https://www.forethought.org/research/intelsat-as-a-model-for-international-agi-governance]
- Rhodes, R. (1986) *The Making of the Atomic Bomb*. New York: Simon & Schuster
- Svarc, R. & Hillebrandt (2025). "Diversifying AI Ownership". *Working paper*
- UBS (2023) *Global Wealth Report 2023*. UBS Global Wealth Management [https://www.ubs.com/global/en/media/display-page-ndp/en-20230815-global-wealth-report-2023.html]

UN High-Level Advisory Body on AI (2024) “Governing AI for Humanity - Final Report” *United Nations*

[https://www.un.org/sites/un2.un.org/files/governing_ai_for_humanity_final_report_en.pdf]

U.S. Census Bureau (2025) “Census Bureau Projects U.S. and World Populations on New Year’s Day” *Press Release Number CB25-03, January 1, 2025*

[<https://www.census.gov/newsroom/press-releases/2025/population-new-years-day.html>]

U.S. Congress (1917) Trading with the Enemy Act, *Pub. L. No. 65-91, 40 Stat. 411*

U.S. Congress (1946) Atomic Energy Act (McMahon Act), *Pub. L. No. 79-585, 60 Stat. 755*

U.S. Congress (1977) International Emergency Economic Powers Act, *Pub. L. No. 95-223, 91 Stat. 1625*

U.S. Congress (1988) Civil Liberties Act, *Pub. L. No. 100-383, 102 Stat. 903*

U.S. District Court (1984) *Korematsu v. U.S.*, 584 F. Supp. 1406 (N.D. Cal.)

U.S. Government (1868) Treaty of Fort Laramie, U.S.-Sioux Nation, *15 Stat. 635*

U.S. Supreme Court (1944) *Korematsu v. United States*, 323 U.S. 214

U.S. Supreme Court (1980) *United States v. Sioux Nation of Indians*, 448 U.S. 371

U.S. Supreme Court (1935) *Perry v. United States*, 294 U.S. 330

Wellerstein, A. (2021) *Restricted Data: The History of Nuclear Secrecy in the United States*. Chicago: University of Chicago Press

Appendix 1: Corporate structures and constraints

There are reasons why several corporations currently having among the most advanced AGI projects have thus far opted to remain private.

Typically, the most important reason why firms IPO is to gain access to larger pools of (lower cost) capital, and to provide liquidity to their early investors. To date, it has been feasible for these companies to raise sufficient capital from private investors. This might be due to there currently being great appetite among long-term-oriented private investors for opportunities in the AI space (along with some of the founders of these companies being individually very wealthy or having strong networks and special skills in private fundraising). As the capital requirements for frontier AI firms soar into the hundreds of billions of dollars or beyond, the number of private investors who have access to sufficient capital will start to dwindle; but so far, this has not proved a decisive limitation.

From the perspective of these firms and their owners and principals, the potential benefits of an IPO must be weighed against a number of disadvantages. One of these, the sheer monetary cost

of the compliance burden—which may amount to around \$100 million per year for a firm with a market cap in the hundreds of billions—is relatively insignificant in the present context. But other potentially more serious drawbacks include disclosure requirements (which would risk divulging strategically sensitive information to competitors), added scrutiny and whistleblower provisions, distraction from focusing on core technology and business development, the possibility of dilution of founder/VIP control, vulnerability to activist investors, and potentially greater exposure to regulatory actions and lawsuits of various kinds. The situation may be extra complicated for AI companies that currently have nonstandard governance structures that involve elements of nonprofit control, profit caps, and various commitments (to the public good, or to step aside and support a competing AGI project under certain conditions, etc.). It might also be difficult to present product roadmaps and business models that are legible to Wall Street analysts. (How do you explain in your quarterly earnings call that you decided to forego billions of dollars in profit because you believed that there was a *40% probability that your AI models are conscious* or have moral status on other grounds such that you were morally obligated to splurge on digital minds welfare?) These things are easier to deal with if your company is privately owned by value-aligned investors that have been handpicked for understanding and supporting your vision.

To accommodate these concerns, various arrangements could be explored. These could include adopting special governance provisions and multiple share classes, to allow the original mission and the company’s founders or principals to continue to have a large degree of control even as new investors come on board. Constituting as a Delaware public benefit corporation could confer more degrees of freedom for the board of directors to pursue objectives besides the maximization of shareholder returns. Creating a publicly traded holding company or closed-end fund that owns a substantial chunk of the original company might allow the latter to remain private while expanding participation opportunities to a broader set of investors. The space of possibilities expands if there is cooperation from regulators. The SEC could issue no-action letters or exemptive orders to waive the enforcement of some regulations that would otherwise prevent or complicate desirable solutions. An even larger space of possibilities opens up if Congress decides that it wants to facilitate a solution along the OGI model.

Another option that could be explored, if the regulatory burdens of a U.S. IPO are too onerous, is an overseas listing, e.g. on the London, Toronto, Singapore, or Hong Kong stock exchange, or, to minimize regulation and disclosure requirements, on the Cayman Islands (CSX) or Bermuda (BSX) stock exchange. This could be combined with the issuance of American Depositary Receipts (ADRs) which could be traded on the NASDAQ (or OTC for Level 1 ADRs).

With any of these options, there are complications and limitations that are beyond the scope of this paper.

Appendix 2: Investor protections and government oversight

While the U.S. government taxes the profits of American corporations, it normally respects corporate property rights and autonomy. However, given the potentially strong temptations to meddle with a successful AGI company—and the availability of ready rationales (such as national security justifications) for doing so—additional assurances would be helpful for making the OGI model credible to investors. Special waivers or encouragement may also be needed to unlock

investments from countries with which the U.S. currently has strained relationships or that are under sanction, which the purest form of the model would allow.

In versions where the U.S. government seeks to concentrate AGI activity within one corporation (US-OGI-1), the role of the USG may expand to include that of suppressing or merging competitors, and shielding the selected project from antitrust action. As discussed in the text, the USG would also have an important role in defining a regulatory framework for AGI, overseeing activities, and protecting corporations from various threat actors. In some scenarios, government involvement might grow into a public-private partnership or take the form of “soft nationalization”.

To reap the full benefits of the OGI model, it is important not only that investor interests will in fact be protected but also that investors—ideally including rival powers who might then be less likely to desperately compete with or sabotage U.S. efforts—are able, *ex ante*, to have a reasonable level of confidence in this fact. A number of measures could be explored to help with this assurance problem. Some of these can be implemented unilaterally by an AGI corporation, while others require the participation of the USG.

Informal signals and pledges. —At a minimum level, corporate leaders and/or government officials could express support for or a commitment to protecting shareholder interests in AGI corporations.

Government relations, PR, and lobbying. —AGI corporations can invest in strengthening their ability to convey their positions to lawmakers and to resist calls for nationalization that might later emerge.

Legal and treaty protections. —Legal constructions can be explored that might make national expropriation more difficult, such as the use of overseas holding companies for some IP. Inspiration might be drawn from how multinationals such as oil majors and mining firms mitigate political risks when operating in regions with unstable governments or weak rule of law. A sufficiently motivated government could also seek to embed support for international investments in American AGI corporations in formal agreements or treaties with other countries. (For more on this, see Appendix 5.)

Geographical dispersal of assets. —An AGI corporation could offshore some assets, such as data centers and IP (such as model weights) to make it harder for its home country to unilaterally seize all of its assets. If the USG knows that an attempt at outright confiscation would fail to capture important overseas assets, it might be less likely to adopt such a policy in the first place.

Technical measures. —For example, remote kill-switches could be installed in data centers that could give foreign countries ways to resist an attempt by the USG to nationalize an AGI company or expropriate their investors. This could be set so as to avoid giving any one foreign country the ability to unilaterally thwart a U.S. AGI project. For instance, a set of *M* countries could each be given a cryptographic key that allows it to periodically send signed messages to U.S. data centers. The data centers continue to operate as long as at least *N* messages are received. This allows a flexible way of allocating fractional veto power (for instance, no country on its own being able to stop the AGI corporation’s data centers, but any two-thirds combination of them being able to halt the data centers until they are satisfied that their contractual rights are respected).

Alignment of interests. —This is an important feature of the OGI model: by allowing political and economic elites to invest in the AGI corporation, it can incentivize influential constituencies to oppose confiscatory government interventions and other hostile actions (both domestically and abroad).

Appendix 3: One or many AGI corporations?

The OGI-N model (many AGI firms) is closer to present reality than the OGI-1 model.

Whether such a competitive commercial landscape is to be regarded as desirable or undesirable depends on where we perceive the main risks from AGI to lie. If we view AGI as being basically just yet another (general-purpose) technology, our default stance would presumably be that competition is beneficial: we tend to get faster progress and greater consumer surplus if there is some competition rather than a single monopolist. The (comparative) global inclusivity of the OGI model, and its ability to somewhat mitigate negative-sum dynamics between states in AGI development, and its potential to discourage extreme concentrations of power (such as might more easily result if AGI development is nationalized or is dominated by a closely held private company), are advantages that apply equally whether there is one or many publicly traded AGI companies. (And investors who wish to protect themselves against the contingency that one of the competitors ultimately “wins” and becomes a monopolist with immense windfall profits could diversify their holdings between all of the plausible contenders.)

An opposing view is that AGI is *sui generis* in part because it presents extraordinary risks such as that of misaligned superintelligence—risks that may require, for their mitigation, a closely coordinated development process.²⁵ It could be important, for instance, that a frontrunner AGI developer has the ability to pause or slow its advances once its AI capabilities reach some critical level, in order to give its alignment team time to implement and test additional safeguards. The scope for such precaution is reduced if other developers are close on its heels and it is evident that spending time on safeguards means forfeiting the race in favor of some less scrupulous and less risk-averse competitor. If we view AGI in this manner, we may not want there to be a competitive market in which multiple companies race to develop and deploy ever more capable AIs.

The OGI model is compatible with either of these views of AGI. If a competitive situation is *desirable*, let there be many AGI corporations. This might be what happens by default. In that case, the OGI-N model could be further supported by encouraging a wide distribution of share ownership (including internationally) along with measures to strengthen assurances that these firms’ property rights will be protected even in scenarios of extreme value gain.

If, on the other hand, a competitive situation is *undesirable*, then the basic OGI model could be combined with measures to eliminate or impede any AGI projects besides the anointed one, or to force a merger. For example, the U.S. could make it illegal to operate such projects without a license and give a license only to the officially sanctioned corporation (which in that case might take the form of a public-private partnership, but still with a globally open investment structure). The U.S. could also lean on other countries to ban competing AGI projects within their jurisdictions, or use its clout (such as its influence over the semiconductor supply chain) to make

²⁵ Cf. Bostrom (2014)

it harder for rivals to thrive. To a first approximation, the same options for suppressing competition—should such an approach be desired—would be available with the OGI model as would be available in alternative models (such as a U.S.-led Manhattan Project for AGI).

It is *possible* to imagine a scenario in which it would be desirable for there to exist only one AGI project globally yet in which it would be more difficult to achieve this outcome in the OGI model. For example, maybe China or some other important country would be more open to shutting down domestic AGI efforts in favor of some joint Intelsat-like intergovernmental project than to do so in favor of some U.S.-based corporate project (even if we stipulate that in the latter case they would be allowed to invest as an equal stakeholder and that the project would have assurances of autonomy from the USG, etc.). Currently, however, a situation in which both the U.S. and China and other significant actors in the AI space would be willing to shutter their domestic efforts in favor of a single joint international project appears quite remote. *If* circumstances were to change enough to make that kind of arrangement feasible, then the new circumstances might also make unilateral relinquishment in favor of the US-OGI-1 model where a U.S.-based AGI corporation takes the lead feasible. This might be a scenario where a U.S. company's lead is so large that any competition would be obviously futile and would only heighten existential risk by cutting into the time available for alignment work.

Appendix 4: Comparisons with some other models

This paper has presented the open global investment model in order to facilitate wider discussion of its relative merits or demerits compared to other approaches to AGI governance. This appendix offers a few brief and tentative remarks on some of the benefits of the OGI model. (A comprehensive comparison between all plausible approaches—such as would be needed to reach an all-things-considered judgment about the best path forward—lies beyond the scope of the present contribution.)

“Manhattan project for AGI”

We have compared the OGI with a “Manhattan project for AGI” in a number of places throughout the main text. To recap:

- OGI may be more agreeable to many incumbents, including current AI company leadership, personnel, and investors. Wide investability could increase incentive compatibility among U.S. and international elites.
- OGI promises a wider and more equitable distribution of benefits and influence than a U.S.-only nationalized project.
- OGI obviates the need for massive government funding.
- OGI may reduce the probability of extreme concentrations of power, in part by creating a kind of dual veto or power structure in which both corporation and government have significant power, and in part by embedding the project in civil society where there is more transparency and societal norms and legal structure rather than in a nationalized project run by the security state.
- OGI could somewhat mitigate international negative-sum racing dynamics and potential for conflict by giving many countries and their elites opportunities to participate in a U.S. project.

- The OGI model offers a range of options concerning the extent of government involvement—from no more than in the status quo, to expanded regulation of the AGI sector as risks or societal impacts increase, to informal consultation and government monitoring, to continuous formal oversight arrangements, to a public-private partnership or other forms of soft nationalization.
- The OGI model is consistent with the development of a broader cooperative international framework for responsible AI development, and/or with more unilateral efforts by the U.S. to influence other countries' AI efforts (such as by manipulating the semiconductors supply chain, etc.).

“CERN for AGI”

We can also compare the OGI model to something like a “CERN for AGI”—an international jointly operated and controlled project to develop advanced AI, loosely modelled on the Conseil Européen pour la Recherche Nucléaire, which built and operates the Large Hadron Collider. It could be either a project shared among Western allies or a more fully global cooperation.

A “CERN for AGI” inherits some of the downsides of “Manhattan project for AGI”: it may be less incentive-compatible for current AI incumbents, it would need massive government funding, it might be slow and difficult to set up, and it would probably involve a more bespoke organizational construction that is less vetted and entrenched than standard corporate law and private property norms.

To these shared disadvantages, a “CERN for AGI” adds some further disadvantages of its own. An international civilian project would confront formidable challenges in terms of information security: with staff drawn from all over the world, and perhaps operating under the umbrella of various diplomatic protections, how would it be possible to prevent espionage or theft of insights, code, or models before they have been made safe for deployment? This is not an issue in fundamental physics research, but it would be critically important in some AGI scenarios. (This difficulty is magnified in the version that is fully globally encompassing, but even in a version that is restricted to some broad coalition of allied countries it could prove quite daunting.) Another potential disadvantage is that it is unclear how competitive “CERN for AGI” could be with the breakneck development speed that is achievable within a corporate setting or an AGI Manhattan project.

A “CERN for AGI” would also have important advantages compared to a “Manhattan project for AGI”. Most notably, it would potentially be more globally equitable and potentially more agreeable to various great powers.

One might think that a “CERN for AGI” would be especially attractive in scenarios in which it is desirable that there exists only one leading AGI effort globally. In theory, all the world's AGI efforts could be concentrated in one globally cooperative joint project. However, we must note that the mere existence of a “CERN for AGI” would not by itself eliminate competing AI corporations or remove the incentives for powerful countries to set up their own national AGI projects. While a truly massive international project would soak up a significant fraction of the world's compute and talent, it is doubtful that any currently feasible version would be big enough to give it a commanding lead. So it might be only if we imagine that the “CERN for AGI” would be coupled with a binding international agreement between all capable actors to refrain from

pursuing their own national projects (and to prohibit companies and organizations within their own jurisdiction from developing frontier AGI) that we would have a model that would achieve a de facto global AGI monopoly. At present, it is very questionable whether the political will for such an arrangement is forthcoming. And if we do imagine that circumstances change such that this does become feasible, then we should also ask whether—in those stipulated circumstances—a similar arrangement might not also be available for the OGI model. (It might not be: it is conceivable that states would be more willing to subordinate their own AI efforts under a “CERN for AGI” than they would be willing to do so under a U.S. AGI Corp, even if they had the option of becoming shareholders of the latter.)

“Intelsat for AGI”

We could also compare OGI to an “Intelsat for AGI”. Intelsat was an intergovernmental consortium and treaty organization, headquartered in Washington, DC, which was tasked with building out a global satellite communications network. It operated under a governance structure in which membership countries were shareholders and had voting power proportional to their investments and usage of the system (albeit with several complicating provisions designed to balance the interests of various stakeholders—including a two-tier board system where one chamber had one vote per country regardless of investment; regional quotas; supermajorities required for certain decisions; an initial U.S.-dominated phase that gradually evolved into one where other countries had increased power; etc.).

Like OGI and the “Manhattan project” and “CERN” models, the “Intelsat” model would not on its own prevent the emergence of competing projects. (In the historical case, the Soviet Union developed a rival “Intersputnik” system together with other socialist countries, and several countries deployed their own national or regional satellite alternatives.) Any of these models would need to be coupled with additional measures to achieve a global monopoly, should such a thing be desirable.

The Intelsat model is in some ways a step closer to the OGI model than the CERN model is, in that Intelsat had an important business component. Commercial motivations played a role alongside geopolitical considerations, and participants were expecting (and received) a financial return on their investments. Influence over the project’s decision-making was also allocated more or less in accordance with each member’s contributions.

One key difference between the OGI and the “Intelsat” models is that OGI (at least in its purer forms) allows for investments and participation by private individuals and corporations. This has several advantages:

- Having ownership be at least partially private could give participating countries greater confidence that the U.S. (or a host country) will not seize the project’s assets when the stakes become high, inasmuch as economic and political elites who are personally invested in the project would have incentives to protect its autonomy and financial interests. (We can make a comparison with how multinational corporations sometimes strive to mitigate political risk when building expensive infrastructure in countries with uncertain rule of law by bringing on board prominent local business families as co-investors.)

One might raise the concern that an AGI corporation that is *entirely* privately owned could be at increased risk of becoming resented by publics who feel left out. For this reason among others, it may be advantageous if some portion of shares is owned by states and/or sovereign wealth funds, pension funds, etc. (or perhaps, in an ideal case, even by some UN-controlled vehicle set up for the purpose). It is worth noting that this risk of popular antagonism could give an AGI firm's directors a business reason—even in the case of a pure for-profit corporation—for providing some public goods (such as making some tools or models available for free or at cost, contributing to scientific research, creating education or access programs for disadvantaged groups, helping advance the UN's Sustainable Development Goals, and so forth).

- The OGI model allows well-entrenched property right norms and existing corporate legal structures to be utilized. This may be both more robust and trustworthy than novel ad hoc legal structures, and faster to set up.

Treaty arrangements for an international AGI project (such as an AGI Intelsat or CERN) would be slower and more difficult to establish. It is possible that a formal treaty could in some ways be more robust and trustworthy than regular commercial property rights law (although this is not entirely obvious). In any case, the two options are not mutually exclusive. An AGI corporation could initially be established under a regular corporate law regime; it could subsequently be bolstered by informal or unilateral pledges from or agreements between government officials. If and when a formal treaty framework between many nations for the governance of AGI becomes politically feasible, it could be used to further solidify the status of the AGI Corp as well as to specify further aspects of an international regime for responsible AI deployment.

- The OGI model is much closer to being on the path of current developments in the AI sector and would plausibly be favored by AI incumbents over alternatives that would involve either the nationalization of the AI sector or the creation of some new international project endowed with superior resources and other advantages designed to make it infeasible for private actors to compete against it. (Although an Intelsat-for-AGI project could contract out work to for-profit corporations, as Intelsat itself did, getting some contract work doled out from an international project may be less attractive to existing AI firms than autonomously pursuing their own complete AGI efforts.)
- The OGI model alleviates the need for state funding, since some or all the funding could come from private investors and from already funded state accounts (such as pension funds). This is significant because the funding required to establish a clearly dominant global AGI leader would be sizable, and many governments might struggle to come up with the funds required for participation in a viable “Intelsat for AGI”. (The original Intelsat was eventually privatized, and it operated as a publicly traded company between 2013 and 2020.)

However, we can also point to certain respects in which an “Intelsat for AGI” could be more attractive than the OGI model to some countries. One might be an intangible sense of greater legitimacy: an official collaboration between governments might be perceived as a more dignified or appropriate vehicle for ushering humanity into a new era than a regular corporation—especially one that is (from the point of view of non-host nations) domiciled in a foreign country.

A more tangible consideration is that in the default scenario, tax revenues from a firm accrue principally to the nation in which it is incorporated. This could make the US-OGI especially

attractive to the United States, but it would be a pro tanto reason for other nations to prefer an Intelsat arrangement in which their shares of profits are not subject to U.S. taxation. In theory, the U.S. could agree to offer tax concessions for foreign investors in AGI Corp; but in practice, other countries may be more likely to attain a favorable tax treatment conditional on the “Intelsat for AGI” than conditional on US-OGI.

Appendix 5: Legal foundations: commercial property rights vs. international agreements

To further compare the OGI model with international alternatives like a “CERN for AGI” or an “Intelsat for AGI”, it is worth examining the relative robustness of their legal foundations.

The historical record shows that international agreements entered into by the United States have sometimes experienced abbreviated lifespans. For example, the Manhattan Project, which was a collaboration between close allies—the U.S. and the UK, and to a more limited extent Canada—was governed by written agreements between Roosevelt and Churchill (the Quebec Agreement in 1943 and the Hyde Park Aide-Mémoire in 1944), which included a commitment that “full cooperation” in developing atomic energy for military and commercial purposes should continue after the war unless and until terminated by joint agreement.²⁶ Nevertheless, once the project succeeded, the U.S. unilaterally ended the agreement by cutting off its allies from the nuclear technology they had jointly developed (McMahon Act in 1946).²⁷ The U.K. subsequently launched its own nuclear weapons program to independently recreate the technology at great expense.

Formal treaties, which have been ratified by the U.S. Senate, give counterparties a somewhat higher level of assurance that the U.S. will honor its undertakings. Even so, some breaches have occurred—most notoriously in the U.S.’s dealings with Native American tribes in the nineteenth century. A well-known example is the violation of the Treaty of Fort Laramie (1868), which had guaranteed the Sioux Nation exclusive rights to the Black Hills of South Dakota. After gold was discovered there, the U.S. military seized the land in 1876–77. Only in 1980 did the Supreme Court (*United States v. Sioux Nation of Indians*) rule this to have been an illegal taking and awarded the tribe \$100 million in compensation (which the tribe has steadfastly refused to accept on grounds that it seeks the return of the land itself rather than monetary compensation; the unclaimed funds, meanwhile, are held in an interest-bearing federal trust account, and are currently estimated at around \$2 billion.²⁸)

Although a treaty has the status of “supreme Law of the Land”, if subsequent to entering a treaty Congress passes a statute demanding its abrogation, then counterparties have no legal recourse in the United States (due to the “last-in-time” rule in the U.S. legal system). Furthermore, treaties that are not “self-executing” (in the legal sense) are ineffectual if Congress declines to pass the further enabling legislation necessary for their implementation, again leaving counterparties without legal remedy in the U.S.

²⁶ U.S. Department of State (1943, 1944)

²⁷ U.S. Government (1868)

²⁸ U.S. Supreme Court (1980)

This structural vulnerability of international commitments stands in contrast to the legal framework protecting domestic property and commercial contracts. While treaty violations can invoke doctrines like the “last-in-time” rule to prioritize later statutes, domestic property rights are anchored in the Constitution’s Fifth Amendment, which demands due process and just compensation for any takings. This safeguard extends to non-U.S. persons, such as foreign states or individuals holding stocks in American companies, who can bring suit in U.S. courts to enforce their claims.

The resilience of these protections is reflected in the rarity and controversy of overrides. *Arguable* cases include the government’s abrogation of gold clauses in contracts during the Great Depression (upheld in 1935 in a 5-4 decision in *Perry v. United States* on the narrow grounds of Congress’s plenary power over currency).²⁹ Another well-known case is the property losses that occurred amid Japanese American internment in World War II pursuant to Executive Order 9066 (1942).³⁰ The Supreme Court initially upheld the exclusion orders in *Korematsu v. United States* (1944) on grounds of military necessity to prevent espionage and sabotage during wartime, though property deprivations were largely a side effect of the internment (owing to the evacuees’ inability to maintain or protect their assets while detained) rather than direct government seizures.³¹ The *Korematsu* decision was later repudiated by a federal district court in 1983, and Congress provided reparations of \$20,000 per survivor through the Civil Liberties Act of 1988.³²

The exceptional nature of these cases underscores the relatively strong protections of private property rights that exist in the U.S. legal system. For foreign owners, political pressures—e.g. during national security crises—have enabled asset freezes or seizures under laws like the Trading with the Enemy Act or International Emergency Economic Powers Act, though these have often been subject to judicial review, post hoc compensation, or international claims processes.³³ Of course, no legal framework can be entirely immune to sufficiently extraordinary political forces. In theory, given enough political support, the U.S. Constitution could be amended to enable the retroactive large-scale seizure of the property of foreign individuals—or of U.S. citizens for that matter.

In summary, the constitutional anchoring of property rights leveraged in the OGI model may offer a more stable and enforceable legal basis for broad international participation in AGI development compared to the treaty-dependent structures underlying models like a CERN or Intelsat for AGI. For a maximal level of protection, multiple mechanisms might be stacked and perhaps combined with physical measures, such as shared control of data centers (either by geographically distributing them across the territories of multiple participants or by other means such as cryptographic protocols).

²⁹ U.S. Supreme Court (1935)

³⁰ Executive Office of the President (1942)

³¹ U.S. District Court (1984); U.S. Congress (1988)

³² U.S. Supreme Court (1944)

³³ U.S. Congress (1917, 1977)