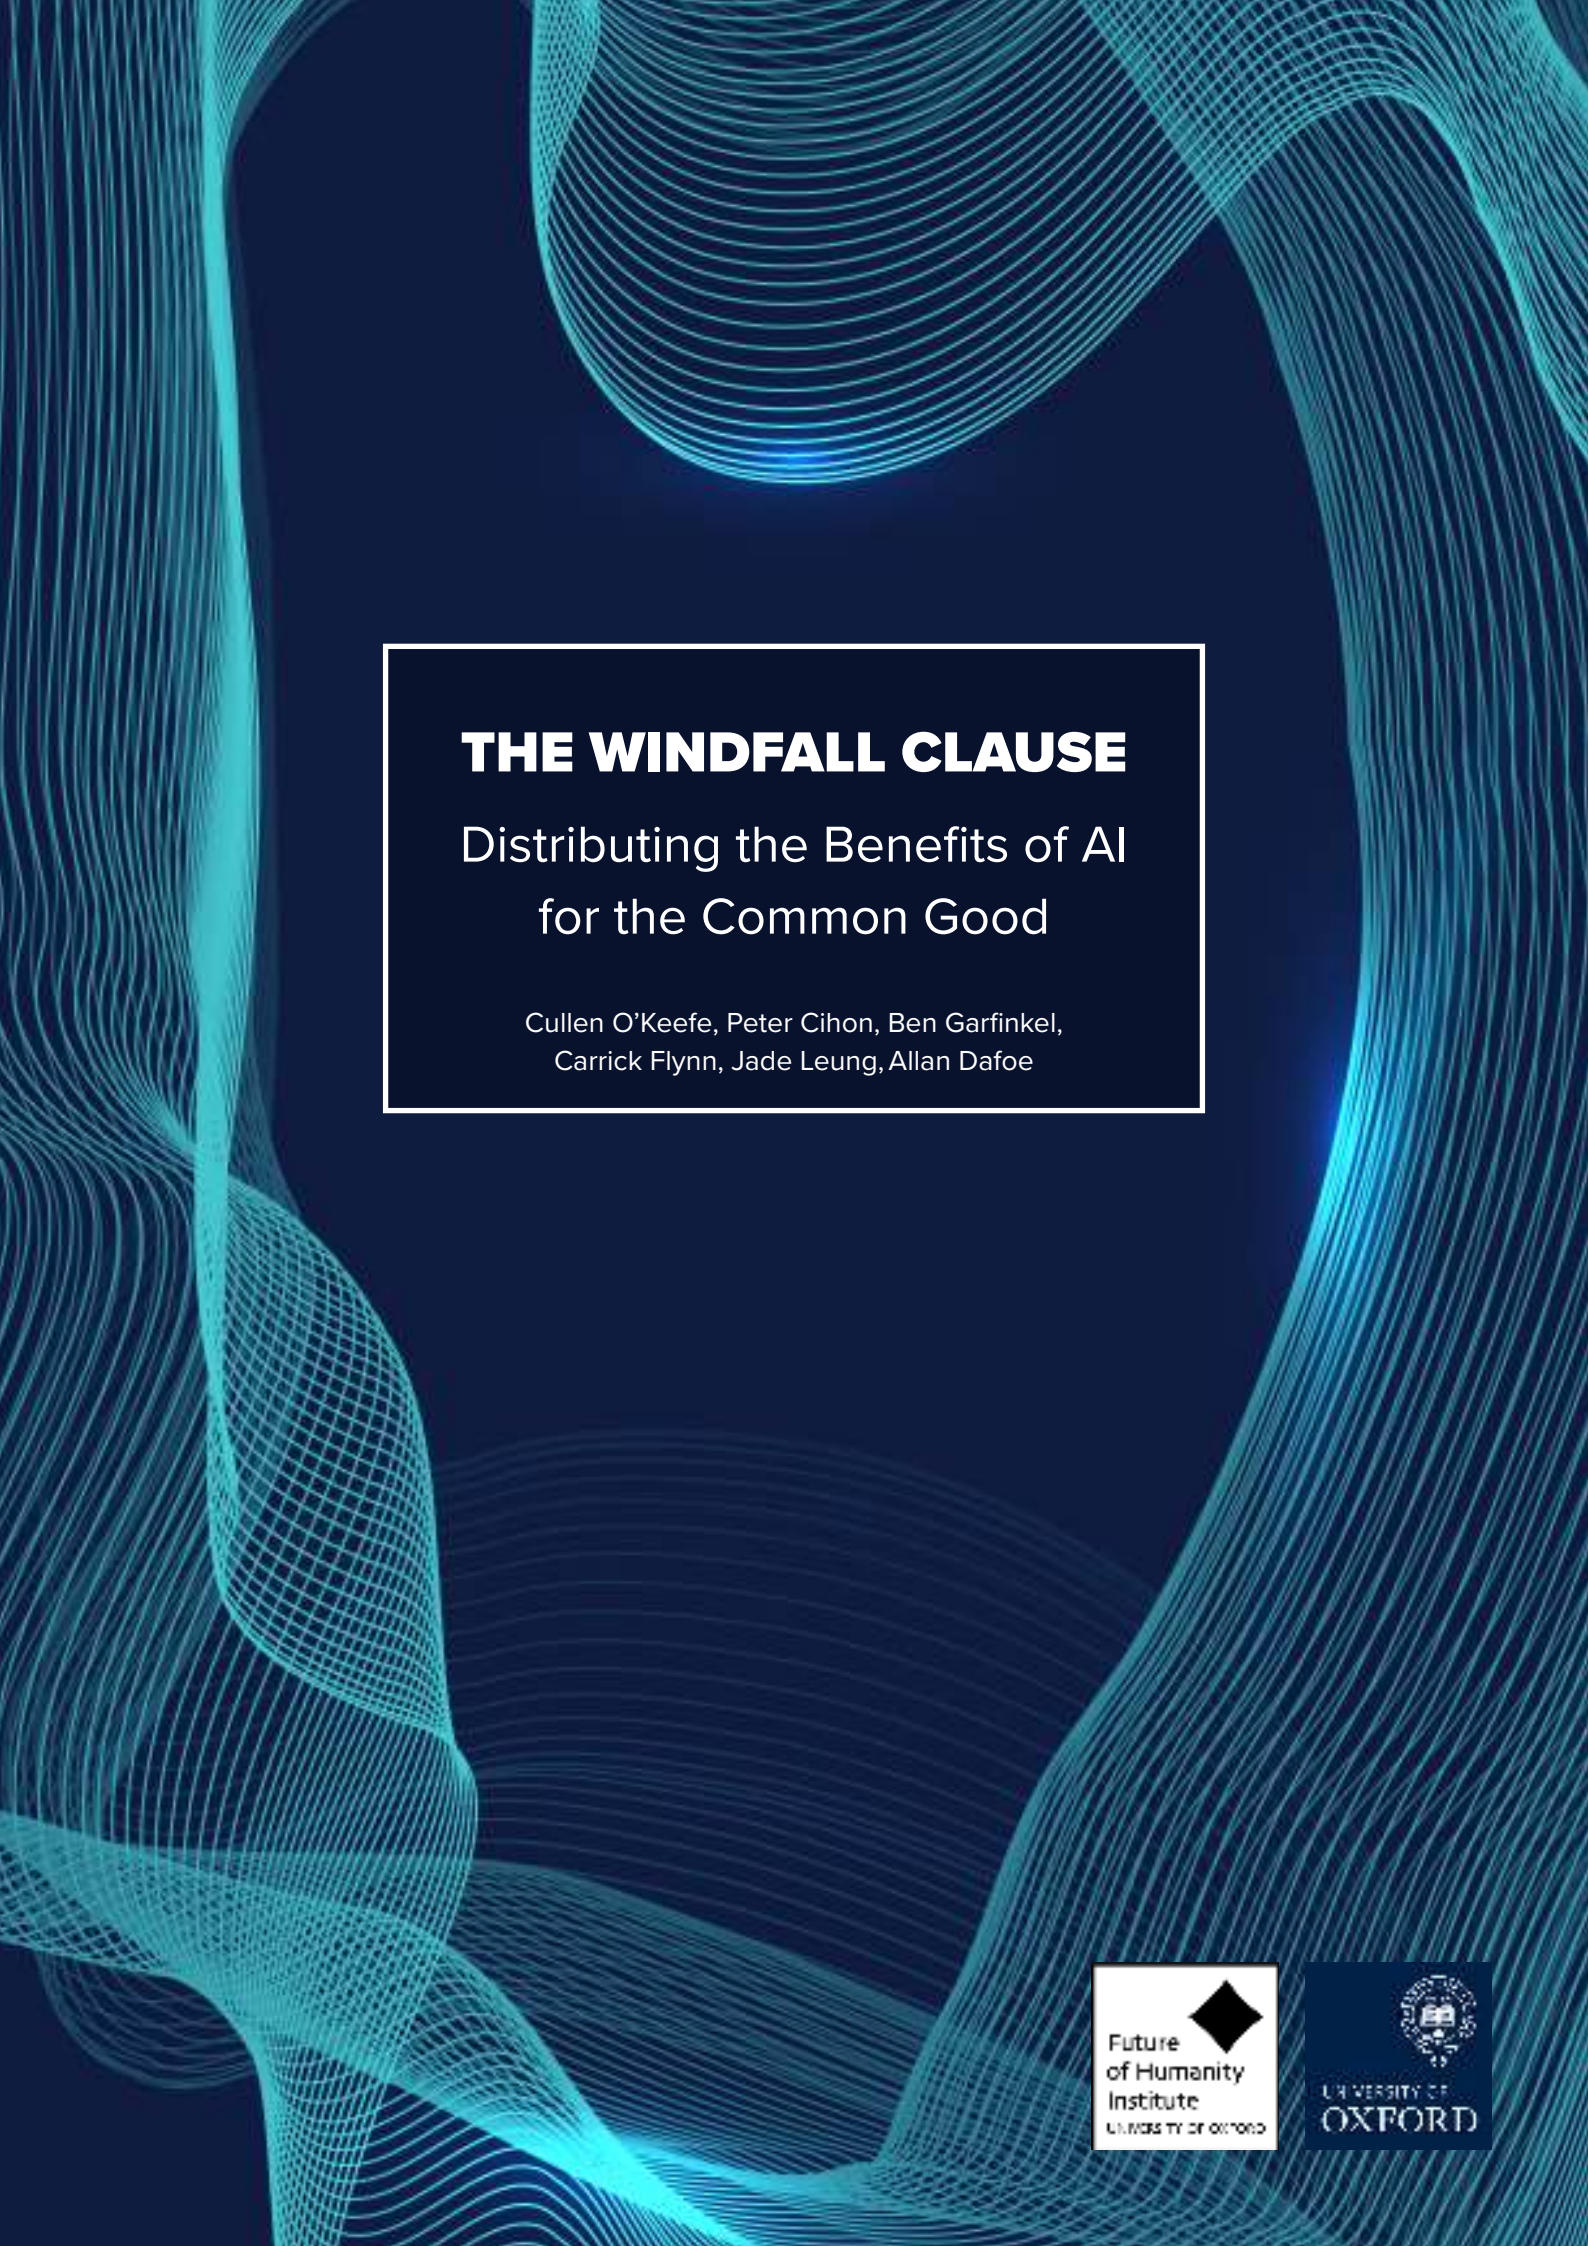

THE WINDFALL CLAUSE

Distributing the Benefits of AI for the Common Good

Cullen O'Keefe, Peter Cihon, Ben Garfinkel,
Carrick Flynn, Jade Leung, Allan Dafoe

Future of Humanity Institute
UNIVERSITY OF OXFORD

UNIVERSITY OF OXFORD

THE WINDFALL CLAUSE

Distributing the Benefits of AI for the Common Good

Cullen O'Keefe,[*] Peter Cihon, Ben Garfinkel, Carrick Flynn, Jade Leung,[†] Allan Dafoe
Centre for the Governance of AI, Future of Humanity Institute, University of Oxford

[*]Corresponding address: cullen@openai.com
[†]Corresponding address: jade.leung@governance.ai

Cite as: O'Keefe, C., Cihon, P., Flynn, C., Garfinkel, B., Leung, J., and Dafoe, A. (2020). *The Windfall Clause: Distributing the Benefits of AI*. Centre for the Governance of AI Research Report. Future of Humanity Institute, University of Oxford. Available at: https://www.fhi.ox.ac.uk/windfallclause/

Future of Humanity Institute

UNIVERSITY OF OXFORD

EXECUTIVE SUMMARY

Over the long run, technology has largely improved the human condition. Most of us live longer, more comfortable, and freer lives than our ancestors because of innovations from the past two centuries. Nevertheless, the economic progress from these innovations has not arrived equitably or smoothly. For example, the enormous wealth generated by the key technologies of the Industrial Revolution—namely factory mechanization, steam power, and railroads—initially accrued to only a few countries and individuals. Even today, industrial wealth remains highly unevenly distributed. In general, while technological innovation often produces great wealth, it has also often been deeply and acutely disruptive to labor, society, and world order.

The world is likely to continue to experience both benefits and disruption from technological innovation. Advances in artificial intelligence (AI) currently warrant special attention in this regard. A number of prominent economists have argued that AI is a rare "general-purpose technology" with the potential to transform nearly every sector of the economy. Some have also suggested, more speculatively, that it may ultimately produce both much higher rates of worker displacement and much higher rates of growth than any other modern technology. In light of ongoing advances in AI, it is imperative that we prepare for the possibility of extreme disruption and act to mitigate its negative impacts.

This report introduces a new policy lever to this discussion: the Windfall Clause. By voluntarily adopting the Windfall Clause, firms would bindingly agree to donate a meaningful portion of their profits if they earn a historically unprecedented economic windfall from the development of advanced AI. We define a "windfall" as a level of income greater than a substantial fraction (e.g., at least 1%) of the world's total economic output. It is unlikely, but not implausible, that such a windfall will occur. As such, the Windfall Clause is designed to address a narrow set of low-probability future scenarios which, if they came to pass, would be unprecedentedly disruptive.

Properly enacted, the Windfall Clause could address several potential problems with AI-driven economic growth. The distribution of profits could compensate those rendered faultlessly unemployed due to advances in technology, mitigate potential increases in inequality, and smooth the economic transition for the most vulnerable. It provides AI labs with a credible, tangible mechanism to demonstrate their commitment to pursuing advanced AI for the common, global good. Finally, it provides a concrete suggestion that may stimulate other proposals and discussion about how best to mitigate AI-driven disruption.

Designing the mechanism by which funds from the Windfall Clause would be distributed would be a significant undertaking. In this report, we enumerate desiderata for an effective and successful distribution mechanism: namely, philanthropic effectiveness, security from improper influences, political legitimacy, and buy-in from AI labs. We preliminarily propose mechanisms for achieving these desiderata and invite further discussion and ideas for how such mechanisms could robustly achieve the moral and societal ambitions of the Windfall Clause.

Among the contributions of this report is demonstrating that the Windfall Clause is legally permissible as a matter of American corporate law, notwithstanding corporate directors' obligation to act in the best interests of the corporation. This is because the Windfall Clause would be low-cost in expectation, given that obligations only vest if a lab earns windfall profits, which is a distinctly low-probability event. Further, a commitment to the Windfall Clause could bring a lab significant benefits through better relations with employees, consumers, citizens, and governments.

The Windfall Clause draws inspiration from precedents in corporate practice, public policy, and personal philanthropy. Notably, the Windfall Clause adds to an increasingly creative global discussion centered around channeling technology-driven economic growth towards robustly equitable and broadly beneficial outcomes. Proposals already represented in these discussions range from reforming international tax law and competition policy to experimenting with radical new models for social welfare such as a Universal Basic Income.

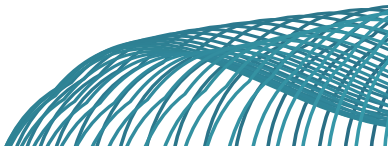

There remain significant unresolved issues regarding the exact content of an eventual Windfall Clause and the way in which it would be implemented. We intend this report to spark a productive discussion and recommend that these uncertainties be explored through public and expert deliberation so that the Clause is well crafted and inclusive. Critically, the Windfall Clause is only one of many possible solutions to the problem of concentrated windfall profits in an era defined by AI-driven growth and disruption. In publishing this report, our hope is not only to encourage constructive criticism of this particular solution, but more importantly to inspire open-minded discussion about the full set of solutions in this vein. In particular, while a potential strength of the Windfall Clause is that it initially does not require governmental intervention, we acknowledge and are thoroughly supportive of public solutions. Thus, in the face of potentially unprecedented prosperity and disruption, we hope to expand the crucial discussion of how to ensure that advanced AI is broadly beneficial for all of humanity.

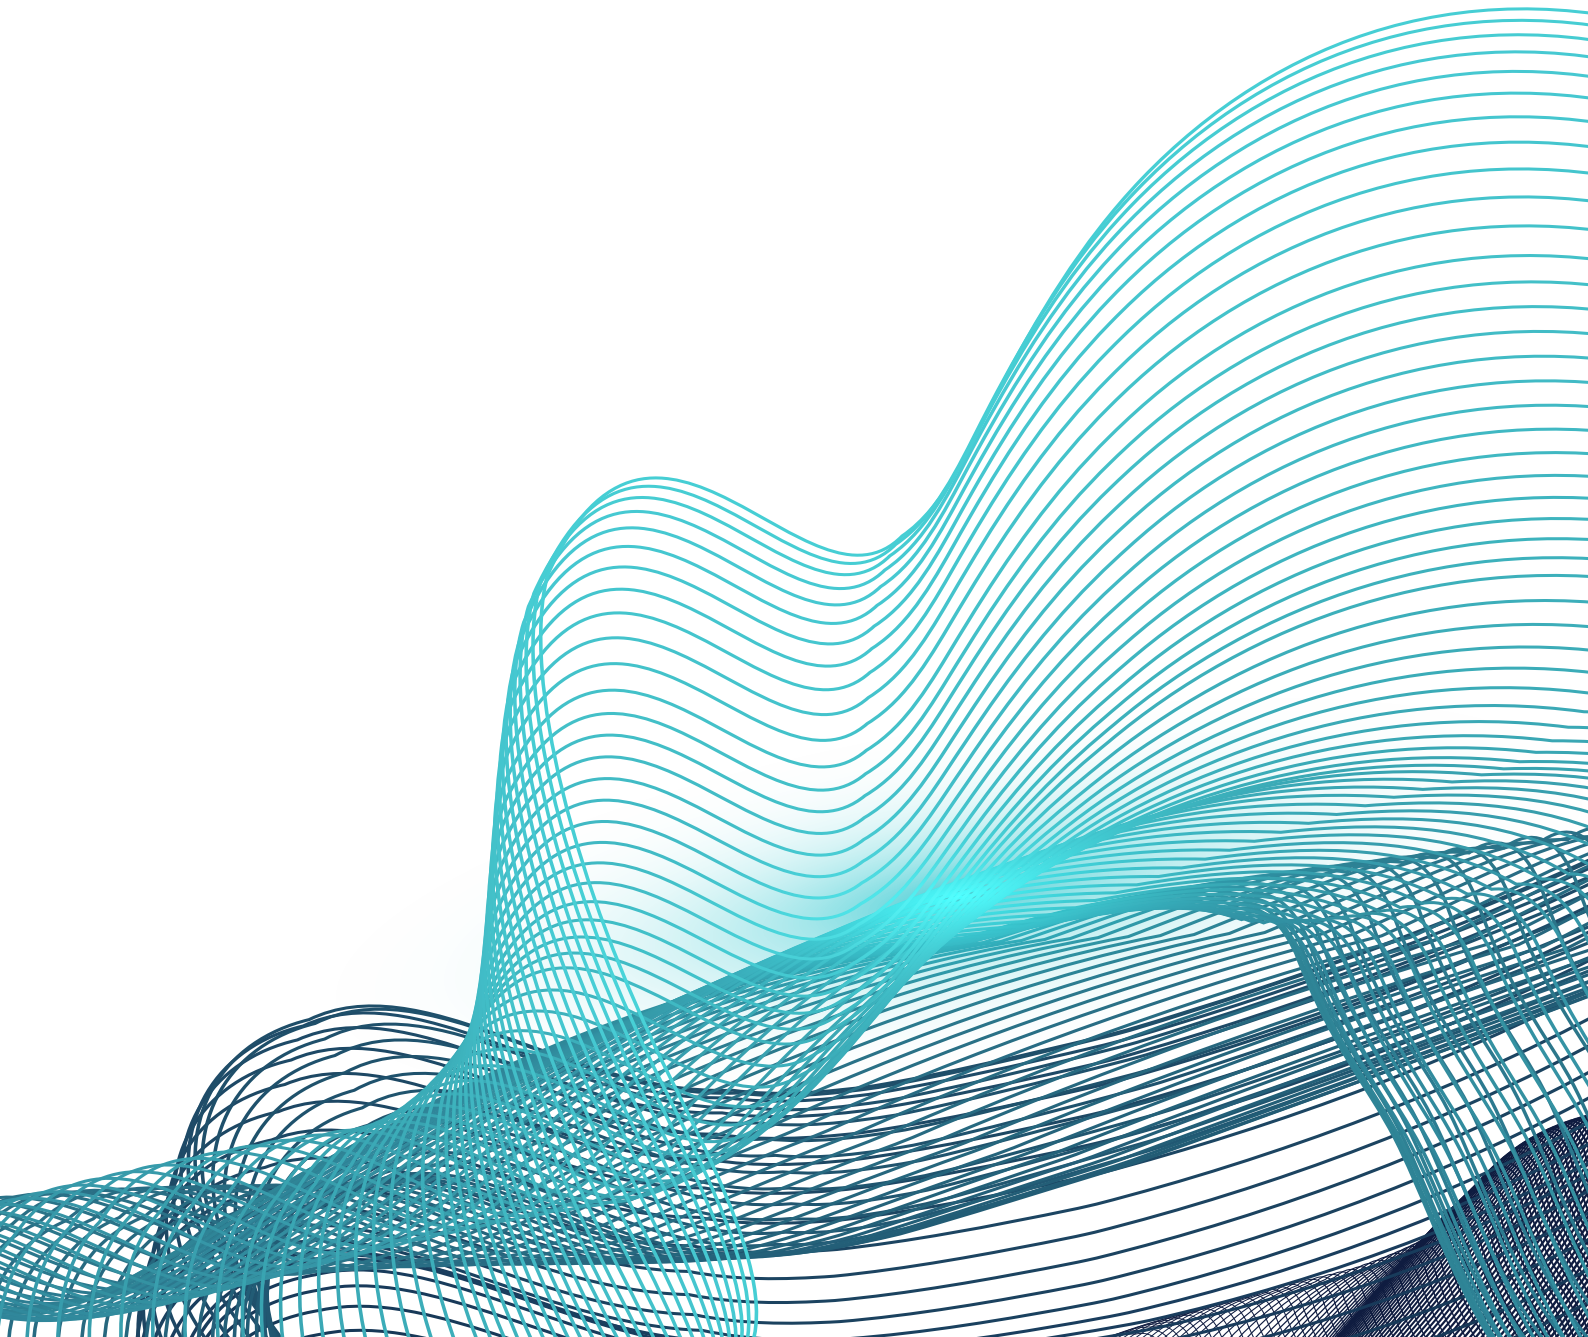

NOTE ON TERMINOLOGY

Throughout this report, we will use the following terms as defined, unless context dictates otherwise:

- *Profits*: income before income taxes.[1]
- *Windfall profits*: extremely high annual profits. We sometimes use ≥1% of gross world product as an indicative threshold.
- *Firm*: an entity that builds or utilizes AI systems capable of generating windfall profits, or attempts to do so. This could be a corporation, nation, university, or group of researchers. However, for this report we focus on private corporate firms for both practical reasons (private firms being the current leaders in AI development) and legal reasons (the fiduciary implications of the Windfall Clause being potentially significant).
- *Signatory*: a firm that has agreed to a Windfall Clause.
- *Distributor*: an entity that receives donated windfall funds as for further beneficial use or distribution. Distributors are the contractual counterparties to signatories in the Windfall Clause.

ACKNOWLEDGEMENTS

For helpful comments and other contributions, the authors thank Emefa Addo Agawu, Philip Trammell, Miles Brundage, Carl Shulman, Anthony Aguirre, Markus Anderljung, Haydn Belfield, Max Daniel, Katya Klinova, Jared Brown, Ulrike Franke, and Nick Bostrom.

This work was financially supported by the Open Philanthropy Project, European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 669751), Future of Life Institute grant RFP2-169, OpenAI, and the Berkeley Existential Risk Initiative.

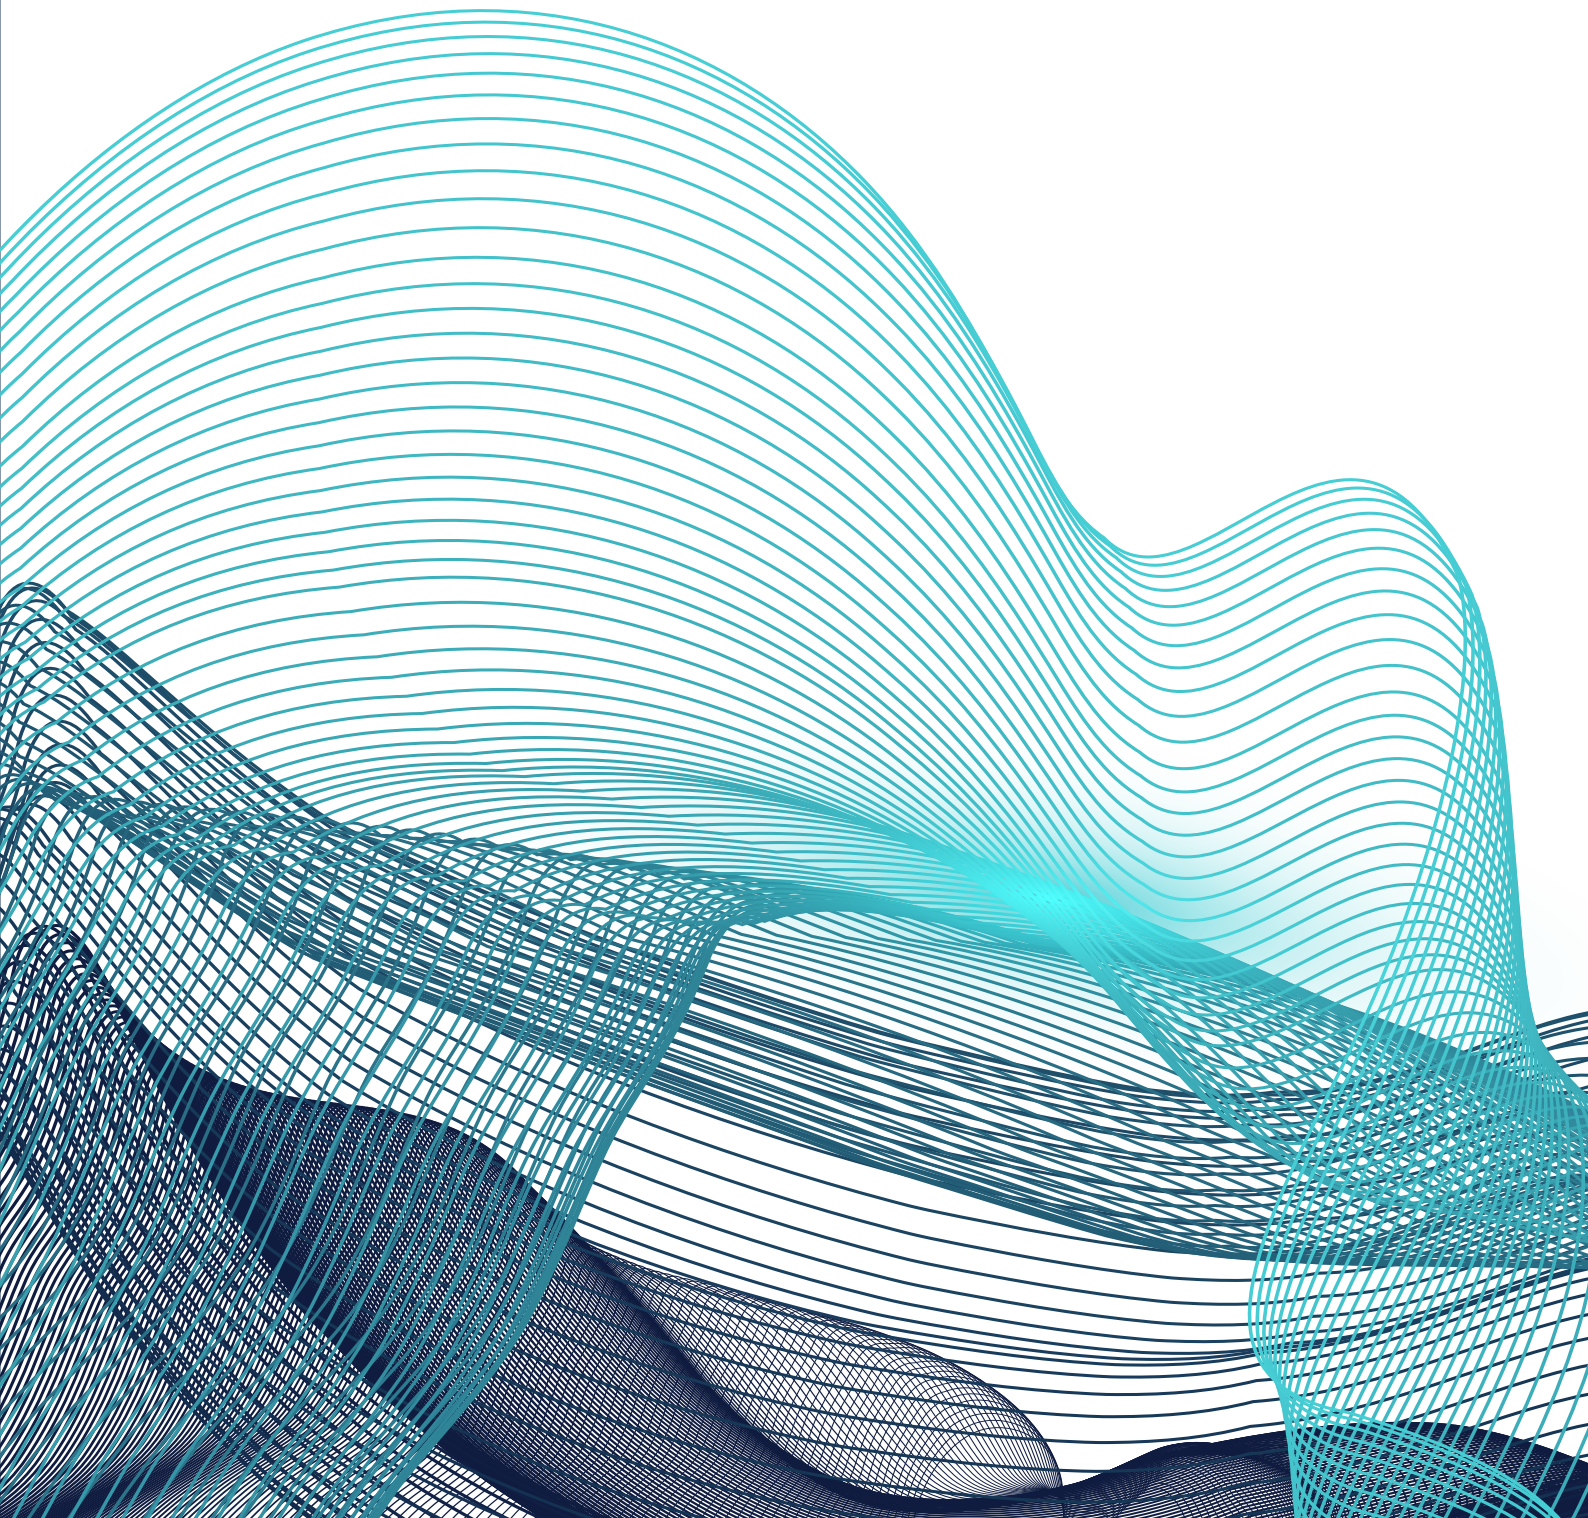

TABLE OF CONTENTS

TABLES AND FIGURES

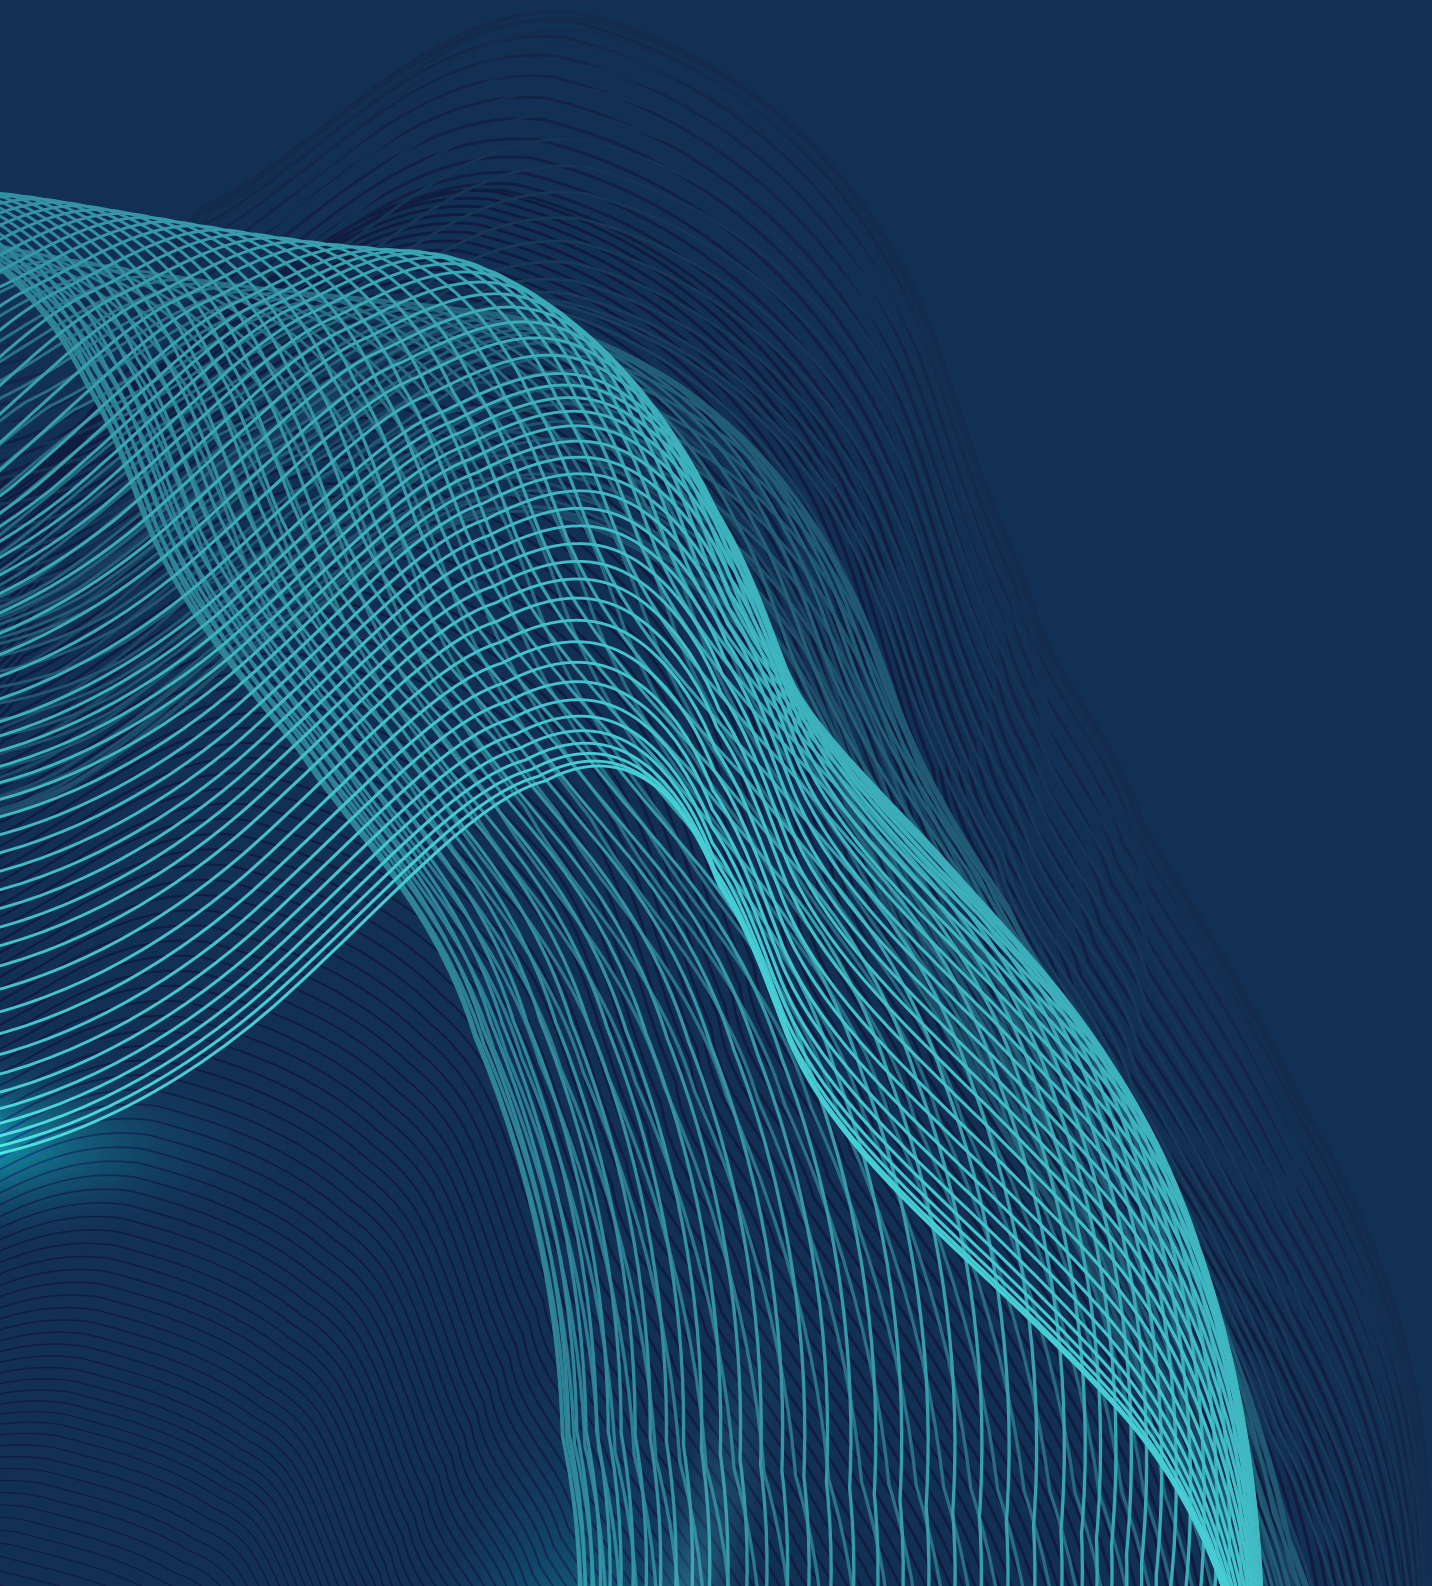

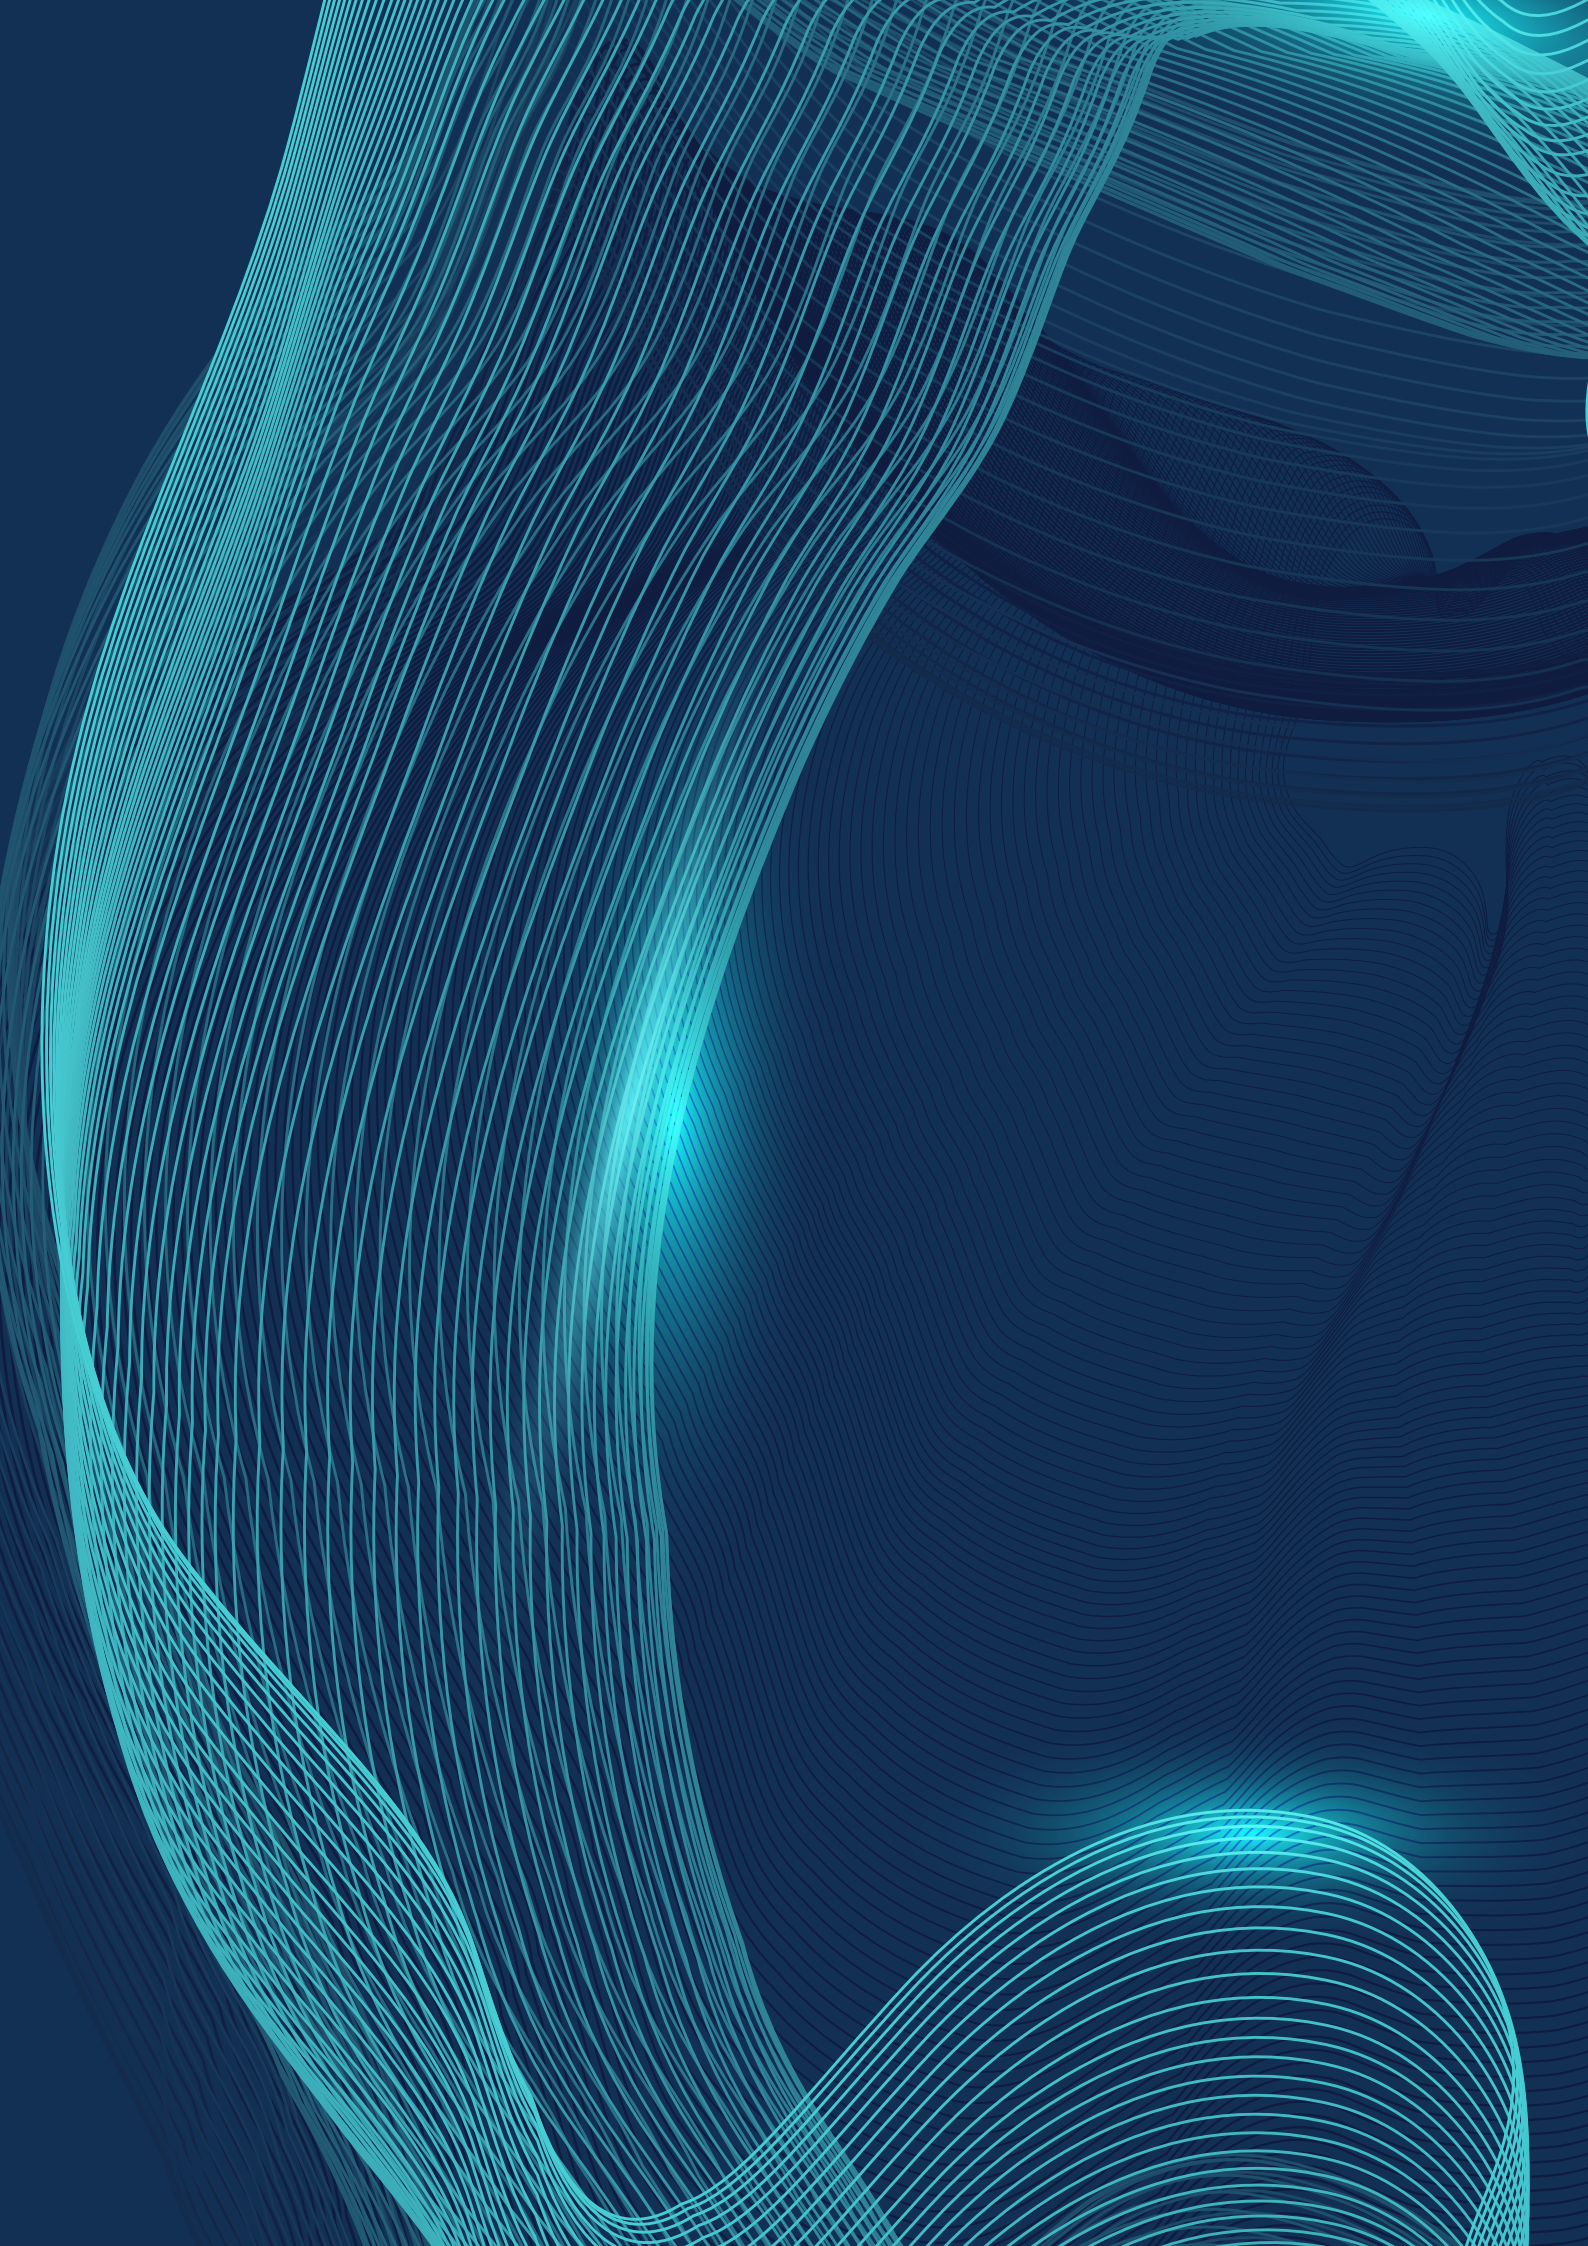

I. INTRODUCTION

A. Context: AI as a Transformative Technology

Some economists have argued that artificial intelligence (AI) is a "general-purpose technology" (GPT), a rare class of technologies with the potential to support broad, cross-sector transformations in the economy.[2] Previous examples of GPTs include the computer, electricity, and the internal combustion engine. If AI follows in the path of previous general-purpose technologies, we should expect a slowly building wave of productivity growth and disruption to existing industries over the next few decades.

Such change would bring enormous economic wealth. However, it is plausible that this wave will leave a small number of firms with exceptionally large shares of relevant markets,[3] with shareholders in those firms therefore receiving much of the benefit. Current AI research and development activity is already fairly concentrated within a small number of firms, which have exceptional access to key resources such as research talent, data, and computing power. Absent substantial changes in market structure, these firms may well be positioned to capture much of the prospective wealth from advanced AI.

Further, this growth in economic wealth could be unprecedented in magnitude and also in speed. In particular, the magnitude and speed of growth is likely to be substantial if AI systems are able to substitute for human labor for the majority of economic tasks, as has been argued by a number of prominent economists.[4] AI researchers, when asked when "human-level AI" or "artificial general intelligence" will arrive, assign a significant (>25%) chance that it will be developed within the next three decades.[5] If the resulting revenue were to accrue to a narrow set of actors, this could give rise to a level of inequality with no close historical analog. Thus, while such historically unprecedented growth and wealth concentration may be unlikely, it cannot be ruled out given the uncertainties around development timelines for AI.

The Windfall Clause is motivated primarily by these scenarios of starkly pronounced—and perhaps unprecedented—concentration of AI-generated wealth. We view it as one potential tool for mitigating the economic, ethical, and geopolitical concerns associated with the possibility of extreme economic and technological inequality. While the Windfall Clause is far from sufficient to mitigate these concerns on its own, we believe that it might have an important role to play in spurring discussion and action to address these concerns. We remain interested in other mechanisms that could achieve these goals or complement the Windfall Clause.

B. Motivations

As the transformative potential of AI has become increasingly salient as a matter of public and political interest, there has been growing discussion about the need to ensure that AI broadly benefits humanity.[6] This, in turn, has spurred debate on the social responsibilities of large technology companies to serve the interests of society at large. In response, ethical principles and codes of conduct have been proposed to meet the escalating demand for this responsibility to be taken seriously.[7] As yet, however, few institutional innovations have been suggested to translate this responsibility into legal commitments for companies positioned to reap large financial gains from the development and use of AI.

This report offers one potentially attractive tool for addressing such issues: the Windfall Clause.[*] The Windfall Clause would pre-commit AI firms to share profits from AI for the common good[8] in scenarios where these profits are extreme and unexpectedly large (qualifying them as "windfall" profits).[†]

---

* Although we generally refer to "the" Windfall Clause, a Windfall Clause could take on various forms. *See generally infra* § II.

† For legal and strategic reasons, we propose a Windfall Clause that would oblige firms to donate nominal—but nonzero—amounts if they earn *near*-windfall profits *infra* § II(A)(1).

In this section of the report, we first outline the reasons why a Windfall Clause might be desirable from an ethical and public policy standpoint, as well as from the perspective of AI firms. Then, the following sections of the report explain how one might implement a Windfall Clause. Section II discusses the content of the Clause by first defining a preliminary "Windfall Function," which would determine how much money a signatory to the Clause owes as a function of its profits relative to gross world product, and then listing crucial considerations for the charitable distribution of that windfall.[9] Section III establishes the legal permissibility of the Windfall Clause under American corporate law.[10] Finally, the last two sections add important context to the Clause. Section IV compares the Clause to historical philanthropy and public policy efforts. Section V closes by acknowledging the limitations of this proposal and compares it to other proposals aimed at addressing the same problem.

B.1. Ethical and Policy Motivations

Here, we lay out the primary ethical and policy challenges that motivate the development and implementation of the Clause. The challenges below also motivate other policy solutions and indeed will likely require a host of solutions to fully address. We do not suggest that the Windfall Clause is the only, or the best, solution for each particular challenge. Rather, considering the scale and seriousness of these challenges—and the uncertainty around how they will develop—we hope that contributing a policy solution targeted at specific scenarios of vast economic windfall addresses an important, if limited, part of the solution space. Finally, we encourage further research into all of these motivations to test their accuracy.

B.1.i. Addressing Loss of Job Opportunities

We expect the net effects of AI on aggregate wealth to be very positive due to increased productivity.[11] However, many have argued that AI could lead to a substantial lowering of wages,[12] job displacement, and even large-scale elimination of employment opportunities as the structure of the economy changes.[13]

One general consideration that supports these concerns is the strong consensus among AI researchers that most, if not all, human work can in principle be automated.[14] Most AI researchers also believe that, at some point in the future, progress in AI will make near-complete automation feasible, particularly given that the upper bound of AI capabilities likely exceeds human capabilities.[15]

As some economists have argued, a world where AI systems can perform almost any task that a human worker could would likely be a world in which employment opportunities are small and shrinking as AI capabilities progress.[16] This is exacerbated by a related concern that technological acceleration may outpace individuals' ability to retrain, such that by the time a worker has retrained, her new skills are irrelevant.[17]

In the long run, AI could serve primarily as a substitute for, rather than as a complement to, labor.* While many technologies have led to periods of job displacement, few technologies have led to large-scale persistent unemployment as valuable new ways of using human labor have previously emerged. However, if AI makes the near-complete automation of human labor possible,[18] this would make AI a substitute for, rather than a complement to, human labor. This could entail not just the unemployment of many or most humans, but their unemploy*ability*.[19] With an absence of meaningful job opportunities, huge numbers of people could find themselves unable to earn a living for themselves or to support their families. This would not be due to unwillingness to work or failure to invest in skill-building,[20] but rather simply due to the fact that AI systems and complementary technologies can outperform job seekers at any productive task.[21]

* Whether AI complements or substitutes human labor depends on the elasticity of substitution between automated and non-automated labor. If this is low (< 1), then human labor complements automated labor and wages will rise even as automation increases. If this is high (> 1), the wage share will fall, since the economy can just substitute away from human work to automated work. *See, e.g.,* Philippe Aghion et al., Artificial Intelligence and Economic Growth (2017) (unpublished manuscript), https://perma.cc/6BWG-X5BU; William D. Nordhaus, *Are We Approaching An Economic Singularity? Information Technology and the Future of Economic Growth* (Nat'l Bureau Econ. Research, Working Paper No. 21547, 2015), https://perma.cc/F3KS-RNWN.

Further, advances in AI may reduce the opportunity for poorer countries to develop via industrialization and the provision of cheap human labor, entrenching their economic disadvantage relative to industrialized countries:

> [N]ew technologies present a double whammy to low-income countries. First, they are generally biased towards skills and other capabilities. This bias reduces the comparative advantage of developing countries in traditionally labor-intensive manufacturing (and other) activities, and decreases their gains from trade. Second, [global value chains] make it harder for low-income countries to use their labor cost advantage to offset their technological disadvantage, by reducing their ability to substitute unskilled labor for other production inputs. From an economic standpoint, these are two independent shocks that compound each other. In other words, each shock increases the costs of the other. The evidence to date, on the employment and trade fronts, is that the disadvantages may have more than offset the advantages.[22]

If automation primarily eliminates employment opportunities in developed countries before developing countries, then the developed countries may be unable to support demand for goods produced in poorer countries, leading to lower employment there.[23]

Experts currently disagree about the net effects of AI on human employment over various timescales.[24] Nevertheless, the value of the Windfall Clause does not depend on any particular prediction about the effects of AI on human job prospects.* Instead, it is sufficient motivation that there is a nontrivial *probability* that AI will, at some point, have substantial and persistent negative effects on unemployment.[25] The fact that such unemployability would be faultless is of particular note; indeed, people across the political spectrum are generally more supportive of programs that support faultlessly unemployable individuals.[26]

B.1.ii. Improving Allocation of Economic Windfall

Pessimistic predictions about the long-run effects of AI on unemployment remain controversial.[27] A related but distinct concern is that most of the benefits from AI[28] could accrue to a small number of actors[29]—particularly, shareholders of "superstar"[30] technology firms.[31] The distribution of such benefits under existing policies is unlikely to maximize overall welfare,[32] even after taking taxation into account.†

There are several reasons to think that the development of AI could increase concentration of wealth. First, automation could decrease the share of a company's income that goes towards wages, and increase the capital share (the share of income that gets paid to the company's owners, rather than labor).[33] If the capital share increases, then, all else equal, income and wealth inequality will likely increase given that capital ownership is already highly concentrated.

Secondly, the particular nature of AI industries could further drive inequality. AI industries could have features of a natural global oligopoly,[34] meaning that, absent countervailing interventions, they could tend towards having only a single (or a few) firm(s) dominating the industry.[35] This would be a natural extension of more general trends toward oligopoly in digital technology industries[36] in which the primary determinants of market structure—such as economies of scale and network effects due in part to low marginal costs[37]—"tend to point towards higher industry concentration ratios . . . ."[38]

Finally, a unique trait of machine learning (ML) systems is their potential to be "general," meaning that they can apply insights acquired in one domain to tasks in other domains.‡ Develop-

* In part, this is because the Windfall Clause contains a *contingent* commitment: major signatory obligations under the Clause vest only when they achieve windfall profits. *See infra* § II(A). Thus, if windfall does not obtain, no or minor obligations arise.

† This is because taxes primarily benefit citizens of the taxing sovereign and tend to be less efficient at generating social benefits than the best philanthropies. *See infra* § V(D)(1).

‡ Research on such "transfer learning" is ongoing.

ments towards increasing generality of ML technology can thus be expected to raise the value of data in general, given the ability to develop a set of capabilities from training on relevant data and then transfer these capabilities to adjacent domains. This increases the incentives to house adjacent databases and perform adjacent tasks under the same roof.[39] In turn, the potential economies of scope would support the competitive position of larger companies with the resources to focus on several domains rather than an individual domain.

One must be careful in making precise technological and macroeconomic predictions. Various considerations not explored above could weaken the case for a trend towards increasing concentration of market power in AI, and indeed, economists are far from consensus on how the concentration of AI industries will evolve in the future. For example, if machine capabilities prove permanently difficult to transfer across domains (or more efficient when limited to a single domain),[40] AI might give rise to a large number of industry-specific monopolies, without any trend toward global cross-industry monopoly. Further, trends towards increasing data efficiency of ML algorithms, or policy pressure towards the preservation of privacy via the regulation of data, could cap the effect of data on the centralization of market power. Nevertheless, while it is difficult to determine the economic impacts of future technologies with confidence, we must recognize that radical transformation of global market structure is a serious possibility.

The extreme inequality that could result from this transformation is undesirable for many reasons. The most obvious is the inherent badness of inequality. Many moral philosophers[41] and a large portion of the public[42] view extreme inequality as inherently objectionable. That is, all else equal, a more equal society might be better than an unequal one.

Even if inequality is not inherently bad, it might be instrumentally bad. The decreasing marginal utility of money[43] means that, all else equal, increasing the poor's share of income increases overall utility.* In other words, an additional dollar increases the welfare of a poor person more than it increases the welfare of a rich person. Thus, inequality-reducing transfers may increase overall utility.[44,45]

Dissatisfaction with extreme domestic economic inequality might also lead to political instability.[46] If this instability leads to civic unrest, violent conflict, crime, or economic downturn, it would be well worth avoiding by reducing inequality.

Finally, the concentration of immense wealth could in turn concentrate cultural and political power.[47] This could mean that a small number of actors (specifically, AI-firm shareholders) have disproportionate sway over society, potentially for a long time if a stable global AI-services monopoly or oligopoly emerges.[48] This is arguably contrary to democratic values, according to which all citizens' views are supposed to be accorded equal weight.

The Windfall Clause could help address these concerns by committing its signatories to distributing the benefits from AI widely across a number of axes, including geography, nationality, culture, value systems, and socioeconomic situation.[49] Its success in doing so depends in large part on the design of the Clause's distribution mechanism, which is covered in more detail in Section II(B).

B.1.iii. Smoothing the Transition to Advanced AI

The transition to a world with advanced AI might not be smooth and may indeed pose substantial challenges to our ability to build and maintain well-functioning institutions. As articulated by Bostrom, Dafoe, and Flynn:

> The speed and magnitude of change in a machine intelligence revolution would pose challenges to existing institutions. Under highly turbulent conditions, pre-existing agreements might fray and long-range planning becomes

* This is true at least as a first-order matter, though reduction in incentives to earn from redistribution may generate offsetting drawbacks.

> more difficult. This could make it harder to realize the gains from coordination that would otherwise be possible—both at the international level and within nations. At the domestic level, loss could arise from ill-conceived regulation being rushed through in haste, or well-conceived regulation failing to keep pace with rapidly changing technological and social circumstances. At the international level the risks of maladjustment are possibly even greater, as there are weaker governance institutions and less cultural cohesion, and it typically takes years or decades to conceive and implement well-considered norms, policies, and institutions. The resulting efficiency losses could take the form of temporary reductions in welfare or an increased risk of inferior long-term outcomes.[50]

A Windfall Clause could reduce this turbulence. To the extent that AI disrupts prevailing social and economic conditions, distribution of windfall profits would provide all beneficiaries with a financial buffer. For example, if someone loses her job due to rapid advances in automation,[51] the Windfall Clause could provide a sort of "unemployability insurance" that enables her to weather the transition to a new economy.*

This individual stability is desirable in itself, but it could also translate to macroscale stability. Automatic widespread distribution of gains from advanced AI would eliminate a potential catalyst for popular political pressure for drastic responses to the new economic conditions.[52] Individuals who, thanks to windfall dividends, maintain or even see increases in their standard of living, may be less likely to initiate a destabilizing response[53] and thus threaten the path towards a more prosperous society.[54]

Finally, the Clause could also contribute to macroscale stability by providing a concrete focal point for renegotiations of the social contract before the advent of advanced AI. It is more productive and stable to have such discussions prior to reaching levels of AI that would trigger the Windfall Clause, when economic circumstances are "calm," than when some actors are much richer and others are resentful. Such conditions would likely lead to better, more considered decisions being made by actors who have the capacity and willingness to consider the broader context beyond their parochial interests.

B.1.iv. General Norm-Setting

By agreeing to be bound by a Windfall Clause, a signatory would send a convincing, costly signal of the firm's intention to ethically develop AI. Indeed, a number of influential companies, officials, and industry groups have already declared such an intent.[55] These groups are now endeavoring to translate their ethical principles into practices and then, in turn, to demonstrate them publicly. Coupled with an international focus on beneficial AI,[56] these developments represent an emergent norm in AI development.

In practice, however, some observers have criticized these efforts as "ethics washing"[57] or "ethics theater"[58] without substantial impact on AI research, development, or deployment.[59] Certainly, there is substantial variance in current corporate principles and practices for ethical AI. Ultimately a process may emerge that sees companies iterate on principles and practices in response to competition and societal debate. As a part of this process, the credible demonstration of ethical development could reinforce a given company's position. The Windfall Clause can serve as one such concrete commitment to develop AI ethically. This, in turn, would strengthen the emerging international norm of using AI for the common good, making further beneficial uses of AI more likely.

* Of course, non-windfall scenarios may also cause turbulence, including unemployment. The sole benefit of the Clause would be in cases where the Clause is triggered. This might be a small portion of the overall expected turbulence.

B.2. Firm Motivations

The preceding subsection explains why a Windfall Clause could be valuable to society. However, in order for this value to be realized, it is important to understand why and how AI firms would willingly agree to the Clause.

The Windfall Clause is a voluntary commitment. Its impact, then, depends on it being compatible with the incentives and interests of those making the commitment. This can be a challenging requirement to meet. For one, for-profit tech companies often have a range of competing incentives and must also meet certain obligations to their shareholders.[60] In the following section, we thus list reasons why the Windfall Clause could be compatible with the interests of for-profit AI firms, including the opportunity to generate goodwill among consumers and employees and reducing a firm's exposure to political risk. We are eager to engage in further discussion in partnership with these firms in order to understand their perspective and needs.

B.2.i. Generating General Goodwill

Agreeing to a Windfall Clause could generate goodwill for the signatory. Indeed, generating general goodwill[61] is often a justification for corporate social responsibility (CSR) initiatives.* Some studies have found a positive correlation between CSR efforts and corporate financial performance[62] through a "virtuous cycle": companies that spend money on corporate social responsibility benefit from a positive reputation,[63] which in turn enables them to spend more on CSR.[64]

In this vein, investing in CSR initiatives can equate to sound business strategy, particularly for highly visible, customer-facing firms whose positive social actions could directly contribute to customer retention and attraction.[65] The above suggest that a properly designed Windfall Clause could improve business outcomes, particularly for large, high-profile AI firms.

The ability for the Clause to generate such goodwill relies on it being robustly designed such that stakeholders are sufficiently convinced that the signatory intends to adhere to the Clause's principles and obligations. This is a challenging bar to meet and requires thorough analysis beyond this report.

B.2.ii. Improving Employee Relations

A firm might also adopt a Windfall Clause to attract and retain top talent. Recent researcher activism over controversial uses of AI technologies at Google and Microsoft, among others,[66] suggest that many top AI researchers prefer working for socially responsible firms.[67] Adopting a Windfall Clause could help attract socially conscious tech talent, thus lowering recruitment costs.[68] Indeed, firms with a good reputation for CSR or a prosocial mission are able to pay much lower wages.[69] One study found that firms with a strong reputation for CSR pay 38% lower wages than firms with a poor reputation.[70] Experimental evidence has similarly demonstrated a relationship between corporate philanthropy and increased efforts by workers.[71] Given the substantial talent bottleneck and wage premium commanded by machine learning researchers,[72] being seen as socially responsible might prove important or even crucial to remaining competitive in the AI industry. This, too, accords with common justifications for CSR.[73]

B.2.iii. Reducing Political Risk

A Windfall Clause could position private firms to be viewed more favorably by the public and the government, thereby reducing risks of adverse regulation[74] or activist action. Numerous commentators have criticized large Silicon Valley firms for a variety of alleged wrongs, including harming their local economy,[75] exacerbating income inequality,[76] concentrating market power,[77] tax avoidance,[78] and abusing users' personal data.[79] A Windfall Clause could build

* When the signatory is a corporation, the Windfall Clause can be understood as a CSR initiative.

goodwill among the public, dampening harmful public antagonism for a small (expected) cost.[80] Governments may be less likely to excessively tax or expropriate firms committed to providing a public good through the Windfall Clause. In fact, they might even feel political pressure to avoid interfering with these firms by those who expect to benefit if they succeed. Notably, mitigation of political and activist risk is a common justification for corporate social responsibility initiatives.[81]

A Windfall Clause could be particularly good at reducing international political risk—risk of expropriation of foreign firms by host governments. It does so by committing the signatory to compensating the host government and its population fairly in case of windfall profits. Furthermore, nationalizing a firm committed to global profit-sharing would likely draw the ire of other, non-host nations. Thus, as with other CSR efforts,[82] the Windfall Clause could protect firms from political risk outside of their home country.

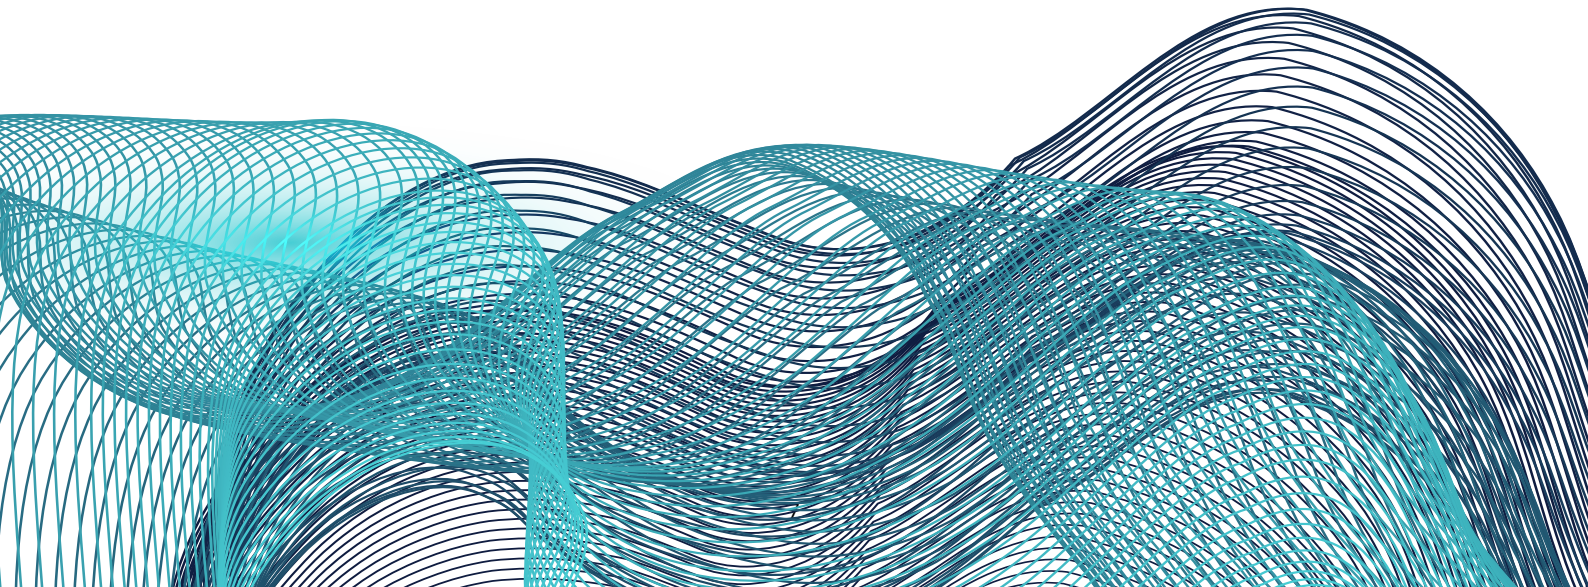

II. DESIGNING THE WINDFALL CLAUSE

A. Defining Contributions: The Windfall Function

The heart of the Windfall Clause is the *Windfall Function*, which defines signatories' obligations under varying profit levels. Although it would be premature to suggest a definitive Windfall Function here, this section discusses several properties we believe the Function should possess. We describe one possible Windfall Function, primarily for the purpose of illustration, and then briefly explore the broader space of possibilities.

A.1. Desiderata

The final design of the Windfall Function should be informed by a range of pragmatic considerations and informed by expertise from a number of different domains. To seed the conversation, we outline here a number of desirable features (desiderata) of this Function and describe one possible Function that fulfills these desiderata for illustrative purposes.

Table 1: Summary of Desiderata

Desideratum	Description	Justification
Transparency	Determining whether a signatory has met their obligations is easy	Avoids costly disputes; provides public relations benefits; increases the enforceability of the Clause
Scale-sensitivity	The amount owed under the Clause depends on the size of the global economy when triggered	Maintains focus on relative windfall
Adequacy	The amount owed under the Clause is proportionate to the challenges it seeks to address	Respects the importance of the policy motivations for the Clause
Pre-Windfall Commitment	Signatories owe small-but-nonzero amounts as they approach windfall profit levels	Demonstrates signatories' intent to fulfill their commitment; tests effectiveness of Clause
Incentive Alignment	Signatories should always have an incentive to earn more profits	Diminishes incentives for evading Clause obligations in order to pursue profit, which in turn mitigates concerns that the Clause could dampen innovation
Competitiveness	Signatories remain competitive with non-signatory firms	Prevents perverse effect of causing non-signatories to out-compete beneficent signatories

First, a signatory's obligations under the Clause should be *transparent*. It should be easy for a wide range of stakeholders to understand the nature of and verify compliance with the signatory's obligations. Such transparency is likely to support both greater accountability and greater public relations benefits for the signatory.* It is also desirable to avoid costly disputes

* Greater public relations benefits would accrue because the signatory is less likely to be accused of making illusory promises.

over whether a given signatory is fulfilling its obligations. Finally, contract terms that are unambiguous are more likely to be enforceable.[83]

Second, a signatory's obligations under the Clause should account for changes in economic conditions (*scale-sensitivity*). We cannot assume that a commitment that is desirable under present economic conditions would also be sensible under all future conditions. One especially important factor to account for is economic growth. An absolute level of profit that would strike one as a "windfall" if it was earned today might appear rather more modest in the context of a much larger overall economy. Since the Windfall Clause is partly motivated by low-probability scenarios in which new innovations generate dramatic economic value,[84] the possibility of faster-than-typical growth should be accounted for.

Third, a signatory's obligations under the Clause should be large enough to significantly mitigate the social concerns that motivate the Clause (*adequacy*),[85] even though efforts from other actors will ultimately be required to ensure that these concerns are addressed.

Fourth, a signatory's obligations under the Clause should be small-but-nonzero under near-windfall profit levels (*pre-windfall commitment*). Despite potentially making it more difficult to secure signatories, it is desirable for both practical and legal reasons for a signatory to take on minimal obligations well before it earns windfall. Practically, compliance with low-level obligations would normalize a signatory's compliance with the Clause and provide some assurance that the Clause functions as desired. Legally, this partial performance can serve as evidence that the signatory intended to be bound by the Clause.[86] Thus, low-level performance would make it more difficult for a signatory to later argue that they were not bound by the Clause. Early compliance would also provide an impetus to the distributors—the entities responsible for distributing windfall gains[87]—to establish processes for doing so. The fact that this would take place when the stakes are low may also be beneficial because the processes would be more likely to be negotiated on the merits, rather than actor self-interest.

Fifth, a signatory's obligations under the Clause should not create an incentive for the signatory to maintain low profit levels or engage in other strategic behavior to undermine the Clause (*incentive alignment*). For example, a signatory should always be better off at higher profit levels—earning more profits should not result in lower retained profits. If the Clause creates incentives for signatories to maintain artificially low profit levels, they might do so by, e.g., making unnecessary capital expenditures. The Function could also fail this criterion if total obligations jump sharply at some point—this would give firms a clear motive for strategic action as they approached that point. The more smoothly total obligations increase, the harder it is for a signatory to decide when to try to evade their Clause obligations.

Sixth, a signatory's obligations under the Clause should not excessively disadvantage the signatory, relative to its competitors (*competitiveness*). A signatory that has signed the Clause should also not be forced to reduce its level of reinvestment to unacceptably low levels after their obligations under the Clause have been triggered. Such a reduction may have negative implications both for the long-term position of the signatory and for economic growth more broadly, given that the signatory would at this point be responsible for a significant portion of the world's economic output.[88]

A.2. An Illustrative Windfall Function

Given these desiderata, we sketch one possible form the Windfall Function below. This sketch should not be taken as a firm proposal. Rather, we intend it to serve as something of a "proof of concept" and as a baseline for more fully developed proposals to build on. It may be the case that an ideal Windfall Function would differ greatly along several dimensions.

In this illustrative Windfall Function, a signatory's obligations under the Clause for a given year depend on its *proportional profits*: its annual consolidated* profits† as a portion of gross world product (GWP, or global GDP) for that year.‡ Proportional profits are divided into four brackets, each associated with a different (and progressively higher) marginal clause obligation, as a portion of marginal profits. Marginal profits are calculated as the difference between the signatory's profit as a portion of GWP minus the lower bound of bracket in which the signatory falls. These brackets are then used to calculate the total financial obligations under the Clause:

Table 2: Illustrative Marginal Clause Obligations

Bracket (Profits as a Percent of GWP)	**Marginal Clause Obligation** (Portion of Marginal Profits)
0%–0.1%	0%
0.1%–1%	1%
1%–10%	20%
10%–100%	50%

As one sample calculation, a signatory whose profits were equivalent to 2% of GWP would be obliged to give 0% for the first 0.1% of proportional profits, 1% for the next 0.9%, and 20% on the remaining 1%. This amounts to approximately 10% of its total profits.§ In general, while marginal obligations increase in steps as per the bracket structure, total obligations increase smoothly.

For perspective on these brackets, the state-owned enterprise Saudi Aramco is the only existing firm whose estimated profits place it in the second bracket or above.[89] The most profitable technology company and the second most profitable firm in 2018, Apple, posted profits equal to approximately 0.06% of GWP, placing it in the first bracket with a 0% Clause obligation.[90] We have found no historical examples of firms whose profits would have placed them in the third bracket, which we associate with windfall profits.

For further context, note that the marginal clause obligation associated with the second bracket roughly equals the portion of total profits that a typical company in the Fortune 500 already donates to philanthropic causes.[91] The marginal clause obligation associated with the third bracket is slightly lower than both the American corporate tax rate and the average corporate tax rate globally.[92] In essence, once a firm achieves windfall profits, its giving obligations gradually rise until they are equivalent to a second layer of taxation.

A.3. Evaluating the Illustrative Windfall Function

The illustrative Windfall Function just described is unlikely to be optimal. However, we can see that it at least roughly fulfills the key desiderata. It therefore suggests that the project of designing an acceptable Windfall Function is likely to be tractable.

First, the Function is easily comprehensible (*transparency*). It is simple to state and can be understood without any advanced knowledge of mathematics, economics, or law. The bracket system should be broadly familiar due to the common use of tax brackets. The Function is also

* Ideally, corporate signatories would be required to report earnings on a consolidated statement once serious obligations vest. This makes determination of compliance and obligations easier. It would also prevent firms from evading Clause obligations by assigning profits to affiliated corporations (e.g., subsidiaries or siblings) that are nominally unbound. One can also prevent this second problem by a contractual provision stating that the Clause applies to any affiliates of the signatory. For further discussion of this issue, *see infra* § V(A)(2).

† For firms that are not for-profit companies, the numerator would have to be in different units.

‡ Proportional profits would be calculated on the basis of the firm's financial statements and the GWP estimates of a designated body such as the World Bank.

§ Calculated as: (0 * 0.1) + (0.01 * 0.9) + (0.2 * 1) = 0.209% of GWP, or (0.209 / 2) ≈ 10% of its total profits.

fairly precise. Of course, one could argue about what the true GWP is.[93] However, it seems to us that the relevant ambiguities can be reduced enough to make the risk of costly disputes quite low.[94] The question of evasion is certainly more complex. Nonetheless, as Section V.A.2. will discuss in depth, it does also appear tractable to put in place sufficiently powerful safeguards. Several strategies commonly employed to evade taxes also do not appear to be directly applicable to this case.[95]

Second, the Function accounts for changes in economic conditions (*scale-sensitivity*). Specifically, it accounts for the possibility of significant economic growth by defining a signatory's obligations in reference to GWP.

Third, the Function imposes significant obligations on windfall-earning firms and is therefore likely to generate revenues appropriate to the task at hand (*adequacy*). It is likely that 20% or 50% marginal obligations will provide significant funds for humanitarian purposes if the Windfall Clause is triggered. For a simplified example, suppose that a signatory earns a $5 trillion (in 2010 dollars) profit* from AI in 2060.[96] Based on OECD estimates, the GWP in 2060 will be $268 trillion (again in 2010 dollars).[97] Thus, the signatory's profits would be 1.8% of GWP.† According to this Function, they would be obligated to give

- 0% of its first $268 billion in profits, for a subtotal of $0;
- 1% of its next $2.412 trillion in profits, for a subtotal of $24.12 billion; and
- 20% of its next $2.32 trillion in profits, for a subtotal of $464 billion

for a grand total of $488.12 billion. This is a substantial amount of money that could accomplish a lot of good in the world.[98]

Fourth, the Function assigns signatories small-but-nonzero obligations under near-windfall profit levels (*pre-windfall commitment*). As discussed, the 1% marginal rate of giving for profits between 0.1% and 1% of GWP is compatible with typical rates of corporate giving.[99] Only once profits rise above 1% of GWP—suggesting a "windfall"—does the marginal rate of giving become extraordinary. While no firm is currently sufficiently profitable to possess any obligations under the Clause, it is unlikely that profits would ever jump quickly from the first bracket to the third. This ensures that the signatory would have significant experience managing small obligations under the Clause before its obligations grow large.

Fifth, because total obligations grow smoothly, there should at no point be a sudden increase in the incentives for a signatory to maintain artificially low profits (*incentive alignment*). While there is indeed a significant jump in marginal rates at 1% of GWP, the jump is in marginal rates, so total obligations just above 1% are not starkly different than those below it.

Finally, the Function implies obligations that are unlikely to be excessively disadvantageous (*competitiveness*). Once a signatory's profits cross the perhaps historically unprecedented threshold of 1% of GWP, its pre-tax giving obligations begin to rise to nearly match a typical corporate tax rate.[100] This is obviously a very significant obligation. However, the obligation also does not seem to be excessively disadvantageous given that any signatory whose profits are consistently greater than 1% of GWP is an extremely successful, and extremely competitive, firm. In general, scaling obligations with the signatory's profits as a portion of GWP ensures that only sufficiently competitive firms will hold significant obligations. Concerns about adverse effects of the Clause on the rate of investment remain valid. Nonetheless, we can also assume that any signatory whose profits exceed 1% of GWP would still be in a position to make investments that are very large in an absolute sense. Section V discusses concerns about the rate of investment more fully.[101]

---

* This expected profit level is picked only for ease of illustration; we make no claims as to its plausibility over other possible expected profit levels.

† Of course, GWP could be higher if the above GWP predictions do not account for the advent of advanced AI by 2060. For the ease of illustration, we will nevertheless assume that these GWP forecasts are accurate even when conditioned on the advent of advanced AI.

It is difficult to imagine that any signatory's profits will ever cross the threshold of 10% of GWP. Even if we assume an unusually high profit margin of 20%, this scenario appears to be one where at least half of all revenue generated across the world is accruing to a single firm. Nonetheless, if we choose to consider such a scenario, it also seems plausible that a marginal rate of giving as high as 50% would be warranted and not excessively disadvantageous to the signatory since there would be little room for competition.

A.4. Directions for Further Investigation

There are many possible versions of the Windfall Function. Here we discuss several possible variations in less detail, to seed further investigations in the future.

One way to vary the Function is to alter the thresholds and marginal giving obligations associated with the various brackets. The number of brackets might also be increased or decreased. Alternatively, brackets might not be used at all. The marginal giving rate associated with the Function could instead increase smoothly, though this might be less scrutable.

A second way to vary the Function is to base obligations on a measure other than proportional profits. Absolute profits might instead be used, although, as noted above, this would not directly account for economic growth. A Windfall Function based on market capitalization appears especially promising.*

A third way to vary the Function is to adjust obligations on the basis of a signatory's financial position at the point where it agrees to the Clause. For instance, it may be appropriate for a firm with a higher probability of eventually achieving windfall profits to commit to smaller conditional obligations. The justification here is that the expected cost of agreeing to a given conditional obligation is smaller the less likely it is that the conditional obligation will ultimately need to be fulfilled.

A fourth way to vary the Function is to define obligations in terms of shares of stock, rather than dollar amounts. For instance, a firm may be obligated to create a certain number of shares of stock to be held by a nonprofit entity if the firm's profits or market capitalization rise above a certain level. This would result in an effective transfer of future profits from other shareholders to the nonprofit entity, due to stock dilution. One way to implement this model would be to grant contingent convertible (CoCo) bonds to the relevant nonprofit entity, with the stipulation that the bonds can be converted into shares of stock only when the firm achieves windfall profits.[102,103]

A fifth way to vary the Function is to associate obligations with strict caps on profit. For instance, a cap might be placed on the total quantity of profit that can be paid out as dividends associated with an individual share of stock. Once the cap is reached, the profit that would be paid out as a dividend to the relevant shareholder is instead given as a Windfall Clause payment. This model is similar to the "capped-profit" model adopted by OpenAI LP, which prevents first-round investors from earning more than one hundred times their initial investment.[104] One potential benefit of the model is that it more directly addresses concerns about individual inequality in cases where unexpected windfall profits would primarily accrue to a small number of fortunate shareholders. However, this model also introduces a number of additional practical and legal complexities. For instance, the model would seemingly result in a class of shareholders who are incapable of receiving further dividends.

A final way to vary the Function would be to account for the effects of taxation on the signatory. We do not propose such accommodation here due to the legal and economic complexity of the topic. However, such accommodation might be necessary to minimize undesirable distortions.

* Market capitalization may be desirable since it is harder to manipulate downwards without upsetting shareholders. However, tying the Clause to market cap could oblige a firm to pay even when it posts no profits.

Ultimately, there is a need for significant further work exploring the space of possible Windfall Functions and evaluating particular options within it. It is important that this work involves the input of a wide range of stakeholders and experts in relevant disciplines, such as corporate finance and optimal tax theory.

B. Distributing the Windfall

The Windfall Clause is only effective to the extent that the windfall is distributed to fulfill the Clause's aims.[105] We expect that designing and implementing distribution systems for the Clause would take substantial work and accordingly will not offer a full blueprint for distribution here. Instead, we here briefly list some significant considerations for such a system and in Appendix II offer some outlines of potential windfall distribution systems.

The considerations we describe below are:

1. Effectiveness
2. Accountability
3. Legitimacy
4. Firm Buy-In

These considerations may sometimes pull in different directions. Thus, a distribution system will sometimes need to trade off these desiderata against each other. Nevertheless, we are hopeful that subsequent work will yield distribution plans that score well across all these dimensions. As with the Windfall Function, conversations with a wide variety of stakeholders and experts are surely warranted here.

B.1. Effectiveness

The distribution scheme should be effective at achieving the desired goals. The broadest definition of this would be that the distribution system is aimed at maximizing expected good produced per dollar received.[106] Of course, judgments about what is "good" are philosophically and politically contentious, and a Windfall Clause will have to balance competing ethical frameworks.

For many measures of the good, however, charitable opportunities vary,[107] and the opportunity costs of failing to maximize expected good produced per dollar received are high.[108] Failing to use windfall profits effectively would both frustrate the Clause's main purpose and squander an amazing opportunity to leverage the fruits of AI for the good of humanity.

Another type of effectiveness would aim not at maximizing the good generally, but specifically at minimizing any harms from advanced AI, such as unemployment and inequality.[109]

Finally, the distribution systems will need a geographical scope. Given that many of the problems posed by AI are global in nature, a global distribution system seems like a sensible default assumption. However, political realities might hinder this. Allocation within that scope also needs to be defined. For example, in a global scheme, do all states get equal shares of windfall? Should windfall be allocated per capita? Should poorer states get more or quicker aid?

At this stage, these goals are under-defined and open to debate. As such, throughout this report we hope that readers will consider them a starting point for further discussion on what an appropriate set of goals could and should be for a windfall distribution scheme.

B.2. Accountability

Part of the motivation of the Windfall Clause is to moderate powerful actors' influence over how windfall profits are distributed.[110] Thus, an effective distribution system would contain mechanisms to prevent distributions being used in a way that is overly parochial. An obligation to use donations effectively would hopefully help, but further measures are warranted. Options could include:[111]

- Transparency measures, such as written rationales for key decisions[112] and reporting obligations to third parties;
- Compliance measures like the "four-eyes principle" (that "transaction[s] . . . must be approved by at least two people")[113] or rewarding whistleblowers for exposing mismanagement;[114]
- Enforceable fiduciary duties imposed on decision-makers[115] such that they are legally obligated to act for the benefit of others rather than their own self-interest;
- Incentive structures that reward effective use of funds;[116] or
- External review and enforcement by public or private[117] entities.

B.3. Legitimacy

Since the Clause is intended to benefit all people, all people should have some voice in ensuring their needs are heard and addressed. Aggregating preferences across so many people would be a serious challenge, but a worthwhile one. Possible mechanisms for doing this include:

- Allowing popular election of some or all windfall decision-makers;
- Allowing international or intergovernmental organizations to influence windfall spending;
- Giving some or all individuals standing to legally challenge the management of the windfall;* or
- Using existing and new consensus mechanisms to guide windfall spending.[118]

B.4. Firm Buy-in

The Windfall Clause depends on firms' buy-in and, to some extent, cooperation. Distribution mechanisms that adequately consider and protect firms' interests—both pre- and post-windfall—are therefore more attractive and effective. Ways of promoting firms' interests without compromising the above values could include:

- Publicizing and celebrating firms' commitments to the Clause; or
- Allowing firms to have some influence over how the windfall is spent.

A good distribution system would combine these and other desirable characteristics to advance the Windfall Clause's goals. Designing such a system would undoubtedly consume a significant portion of the negotiations before and after implementation of the Windfall Clause. We hope that the general public, civil society, and AI firms will engage in productive deliberation on the many important considerations in pursuing this part of the greater Windfall Clause project.

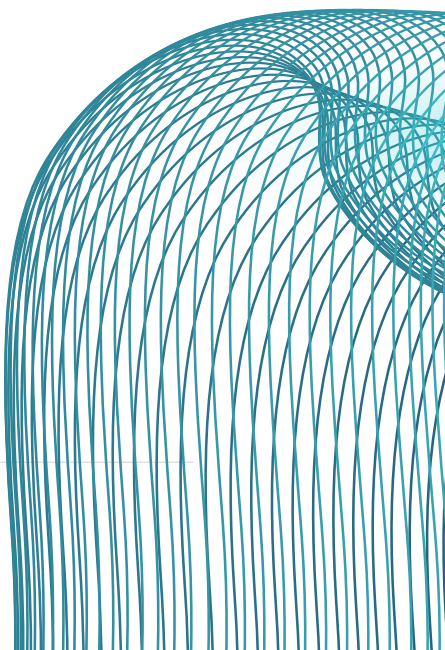

* This would be the flipside of giving windfall decision-makers fiduciary obligations. *See supra* § II(B)(2).

III. LEGAL PERMISSIBILITY

This section demonstrates that the Windfall Clause is permissible as a matter of corporate law. Although some firms may not be for-profit corporations, we focus on corporations in particular given their leading role in the current AI ecosystem and particular restraints on their ability to engage in philanthropy. In Appendix I, we also explain why the Clause could be binding as a matter of contract law and difficult to challenge long after signing as a matter of civil procedure.

In this section, we pay the most attention to the corporate law of the state of Delaware, since a large plurality of major American corporations are incorporated there. We also analyze the legal permissibility of the Windfall Clause in the People's Republic of China,[119,120] though we stress that none of the authors are experts on Chinese law. Future analysis will hopefully clarify the permissibility of the Clause in these and other relevant jurisdictions.

Our key argument is that, notwithstanding corporate executives' obligations to their shareholders, the Windfall Clause is permissible as a matter of corporate law.

A. Permissibility in the US

A.1. Delaware Law of Corporate Donations

Delaware corporations are generally empowered to make charitable donations.* However, fiduciary duties still require for-profit corporate management to exercise that power in the best interests of the corporation.

"Corporate expenditures today are judged under the business judgment rule, a standard that accords substantial deference to management's judgment. The fact that a perceived benefit is intangible, noneconomic, or uncertain will not invalidate a corporate expenditure."[121] Thus, "courts [have] upheld discretionary corporate giving on the theory that donating to charity benefits the corporation,"[122] such as through generating goodwill.

Beyond the requirements of the traditional business judgment rule (BJR),[123] Delaware courts have apparently imposed an additional restriction on corporate donations: the gifts must be "reasonable."[124] To determine reasonableness, Delaware courts explicitly consider the charitable deduction allowance in the Internal Revenue Code (I.R.C.) a "helpful guide."[125]

I.R.C. § 170 governs the charitable deduction. The general rule is: "[t]here shall be allowed as a deduction any charitable contribution . . . payment of which is made within the taxable year."[126] For corporate taxpayers, such deductions cannot exceed 10% of the taxpayer's annual taxable income.[127] In the only Delaware case on this point, *Sullivan v. Hammer*,[128] the Delaware Chancery Court, in an unpublished opinion, approved a proposed shareholder class action settlement that included corporate donations to a museum. Parties disagreed about whether to value the gift at 7%[129] or between 15% and 17%[130] of the defendant-corporation's net income. The Chancery Court dodged the question of the proper evaluation of the gift but nevertheless upheld the gift as "within the range of reasonableness."[131] The Supreme Court of Delaware upheld that decision on appeal, also without definitively evaluating the gift.[132]

Despite early hostility to corporate philanthropy,[133] in the handful of modern challenges thereto[134] no American court has invalidated a corporation's philanthropic act. However, in other contexts Delaware courts have found that a corporate manager who is employed by a nonprofit institution lacks independence from a major donor to that institution.[135] This suggests that a corporate manager who made a corporate donation to a nonprofit that employed that manag-

* Every Delaware corporation has the power to "[m]ake donations for the public welfare or for charitable, scientific or educational purposes, and in time of war or other national emergency in aid thereof . . . ." DEL. CODE tit. 8, § 122(9) [hereinafter DGCL]. Note, however, that the corporate management can exercise the powers of a corporation in ways that breach their fiduciary duties to shareholders. This subsection explains the relevant duties. The next subsection explains why the Windfall Clause, as a corporate donation, is not a breach of those duties.

er would *not* receive the protection of the BJR because such a transaction would not be disinterested. However, courts have not found that managers benefit from donations to charities on whose board the manager sits.[136] Finally, despite influential statements to the contrary,[137] courts apparently have no qualms with managers directing donations to their "pet charities."[138]

A.2. Delaware Law as Applied to the Windfall Clause

At first glance, a Windfall Clause might seem vulnerable to challenge on the grounds that it could, in the case where a firm exceeds the threshold, commit a corporation to donate in excess of the allowable I.R.C. deduction. However, this argument commits the mistake of evaluating the Clause in hindsight, only in the event that it is triggered. The proper approach is to evaluate the *expected value* of the Clause at the time of commitment.

The BJR maximizes shareholder wealth by insulating directors from second-guessing by courts: in the absence of BJR protections, directors would prefer low-risk investments, leading to suboptimal expected returns.[139] A related justification for BJR protection is the "risk that hindsight bias will color our assessments of what an acceptably good [decision] process would have been or would have produced."[140]

Thus, the proper way to evaluate a Windfall Clause, as with any corporate decision, is by its expected value.[141] Even though the cost of a Windfall Clause might exceed Delaware donation guidelines if triggered, the relevant inquiry is the *expected* cost of the Clause: the cost if triggered multiplied by the probability of a triggering event.* Furthermore, since the cost of the Windfall Clause would be borne, if at all, in the future, its expected costs should be temporally discounted to yield a present expected cost. Due to both the improbability of any given firm achieving windfall profits and the exponential discounting of any profits they did earn, the gross expected present cost of the Windfall Clause would probably be very low for any signatory.[142]

For the reasons stated above,[143] the benefits of the Clause could be significant. This forms the basis for the case that corporate directors would not be grossly negligent in believing that the Clause is in the corporation's best interests.[144] The Clause should therefore survive a challenge under the BJR.[145]

A useful analogy can be drawn between the Windfall Clause and stock option compensation, which is incontrovertibly permissible. Like the Windfall Clause, stock option payments have a permissibly low present expected value, but can have a much higher value once exercised.† For example, in 2005, Facebook issued a large block of stock options to its CEO Mark Zuckerberg.[146] Zuckerberg exercised these options in 2012[147] and 2013,[148] for a net pre-tax value of over $5.3 billion.[149] If evaluated ex post (i.e., after Zuckerberg exercised the options), the 2005 stock options were wildly excessive CEO compensation for a company that earlier that year was evaluated at only $100 million.[150] However, when one evaluates the stock options ex ante—as is proper—then they were likely appropriate.[151] Analogously, the Windfall Clause might appear excessive as a corporate donation if evaluated ex post, but it is clearly reasonable if evaluated ex ante.

Of course, legal conclusions are always subject to uncertainty due to the possibility of departure from precedent. Such departure might be more likely than usual with as novel a contract as the Windfall Clause. Still, assuming the state of the law is stable and applied in this case, the Windfall Clause should survive legal scrutiny.

B. Chinese Corporate Philanthropy Law‡

China houses a substantial number of important AI firms. The parallel question of whether the Windfall Clause would hold under Chinese law thus warrants investigation.

* More precisely, the sum of the products of the values from all possible windfall contingencies and their respective probabilities.
† Indeed, since there is no upper bound on stock price, stock options could have arbitrarily high value ex post.
‡ None of this Report's authors are experts in Chinese law, so this analysis should be considered preliminary and cursory. If this project advances further, we anticipate needing to engage much more substantially with Chinese law.

The corporate law of the People's Republic of China (hereinafter, "China") is fundamentally different from Delaware law in that it instructs corporations to "observe social morality and business ethics, act in good faith, accept supervision by the government and the public, and bear social responsibilities."[152] Thus, in the past decade China's national and subnational governments have actively *promoted* CSR, including corporate philanthropy initiatives.

China's current CSR regime originates from the mid-2000s:

> In 2005, President Hu Jintao articulated China's new policy vision of building a "harmonious society," which was adopted as a policy mandate by the Central Committee of the Communist Party in 2006. This policy directive included a mandate for all governments to "strengthen CSR," paving the way for governments at the central and subnational level to move into this policy space. Beginning in 2006, President Hu Jintao and other leaders within the Party and the National People's Congress began to emphasize that businesses operating in China must not place profit seeking above morality and the broader social welfare, but rather should adopt responsible and sustainable business practices.[153]

This policy culminated in revisions to China's Company Law,[154] including the incorporation of Article 5, which reads: "When engaging in business activities, a company shall abide by laws and administrative regulations, observe social morality and business ethics, act in good faith, accept supervision by the government and the public, and bear social responsibilities."[155]

The CSR elements of Article 5 are aspirational, not enforceable.[156] Nevertheless, "[t]he provision encourages firms to consider the social wellbeing of all the constituencies. According to this rule, companies would adopt government and public interest as its guiding value, which is in line with the overall desire to achieve a more harmonious society."[157]

"Governments at all levels have been active in raising awareness of CSR principles."[158] However, the primary focus has been encouraging businesses to meet and exceed laws[159]—especially with environmental and sustainability regulations[160]—not philanthropy. Nevertheless, China's CSR regime does include encouragement of corporate philanthropy. At the national level, China's new Charity Law[161] provides for tax deductions on corporate donations.[162] Some subnational governments reward corporate philanthropists with "extra points towards local CSR certification."[163] "Others directly solicit corporate contributions to development or charitable projects sponsored by state-approved organizations."[164]

To summarize, unlike the Delaware corporate law system, Chinese corporate law actively encourages some degree of corporate philanthropy. This provides a promising starting point for further legal analysis on the nuances of corporate law in China, and how they interact with policy proposals like the Windfall Clause.

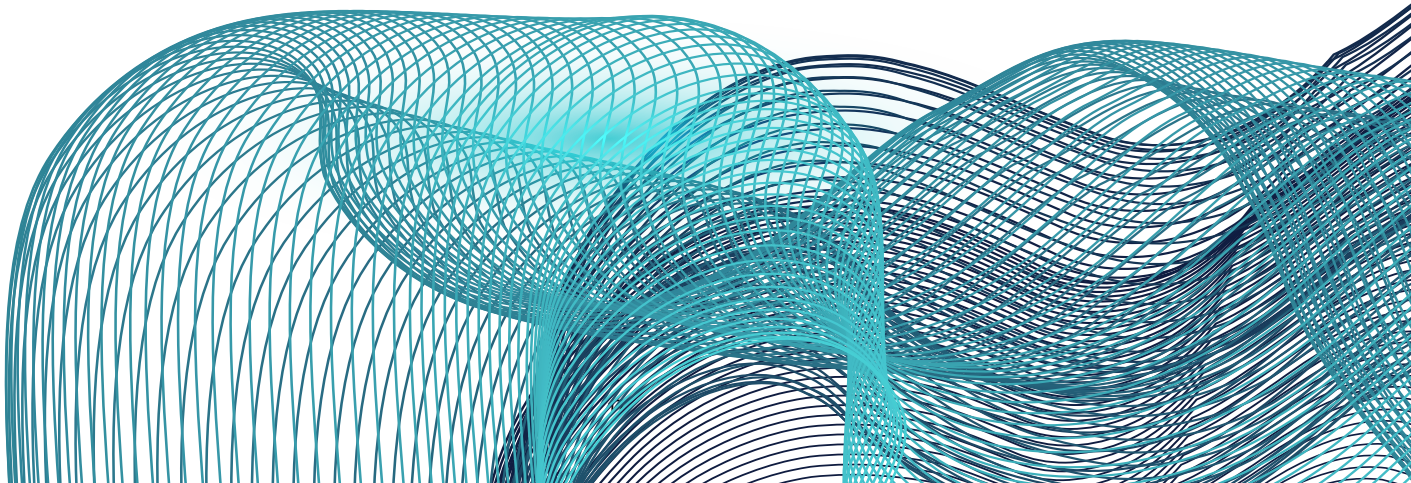

IV. HISTORICAL PRECEDENTS

While the Windfall Clause, if triggered,* would involve redistributing an unprecedented magnitude of funds, the Clause represents a continuation of innovations in philanthropy. Existing CSR initiatives already match the scale of the Windfall Clause in present expected cost, scope, and breadth. National policies are constructed to transfer vast sums to support the welfare of future generations based on realized windfalls. Personal philanthropy—in both theory and practice—often uses a threshold model. The Windfall Clause draws inspiration from all three. Recognition of these precedents is helpful for understanding the viability of such a plan. Although ambitious, therefore, the Windfall Clause is not as unprecedented as it may first appear.

A. Precedents from CSR Practices

A.1. Overall Cost

If triggered, the Windfall Clause could be the largest CSR program in history. Yet, the likelihood of any particular project triggering its commitments under the Clause is both very low and likely far off in the future; thus, its expected present cost resembles philanthropic efforts of corporations today.

As a very simplified illustration, recall the example from *supra* § II(A)(3) wherein a signatory earns $5 trillion (in 2010 dollars) profit from AI in 2060 when the GWP is $268 trillion (again in 2010 dollars). Under the proposed Windfall Function *supra* § II(A)(2), such a corporation would be obligated to give $488.12 billion (again in 2010 dollars).[165]

Suppose, optimistically, that a firm in 2019 had a 1% chance of earning such profits by 2060 (and otherwise a 99% chance of earning profits that would not trigger the Windfall Clause at all). Further suppose, for extreme simplicity, that if the firm achieves windfall profits, under the Clause it will continue to owe $488.12 billion per year forever.†

If we discount annually at 10% (the approximate cost of capital for internet software firms),[166] the present cost of such a commitment, if realized, is equal to:

$$\sum_{t=(2060-2019)}^{\infty} \$488.12 * 10^9 * (1 - 0.1)^t$$

This series converges to $64.934 billion (in 2010 dollars). If we further discount that amount by 99% (to account for the probability of not achieving windfall profits and therefore bearing no costs under the Clause), we arrive at an *expected* present cost of $649.34 million (in 2010 dollars).

To compare this to past corporate philanthropy, 2015's largest donor was Gilead Sciences with $446.7 million.[167] In 2010 dollars, this would be $414.17 million.[168] Note that this historical comparison only includes outright cash donations, not the (opportunity) cost of other forms of CSR.[169] Thus, even under very optimistic assumptions about a signatory's prospects for achieving and maintaining windfall-generating AI, the expected present costs to a signatory are only 60% greater than leading corporate philanthropy efforts today.

* Some Windfall Functions might be continuous, *see supra* § II(A)(4), so there is no single point at which the Windfall Clause is "triggered." For the ease of discussion, we can say that a Windfall Clause using such a function is triggered when the signatory's marginal obligation exceeds 10%.

† In reality, of course, this would depend on how GWP rises relative to the firm's profits. The projections and calculations for such a scenario are far too speculative and beyond the authors' capacity to estimate here. Thus, we use the very simple assumption of perpetual and constant obligations to avoid this speculation. We hope that future work will allow for better estimates of the expected costs of the Windfall Clause.

A.2. Other Factors

The Windfall Clause would be a global effort. Firms have long engaged in global CSR, especially in areas touched by their supply chains.[170] The Responsible Business Alliance, for example, is a group of over 140 companies—including leading computer hardware manufacturers—that administers rules on issues in supply chain responsibility.[171]

There is precedent for a firm distributing massive, unexpected profits to support social welfare. During the development of the atomic bomb, the Belgian Congo-based Union Minière du Haut-Katanga reaped huge profits from uranium mining[172]—so much so that they became a political liability.[173] Concern over criticism and nationalization efforts led the company in 1946 to offer the colonial government a check for the equivalent of between 22 to 109 million 2015 dollars[174] to support social welfare in the colony.[175] (The colonial governor declined because he did not want to draw attention to the sheer size of the company's profits.)[176]

There is also precedent for the multi-firm nature of the Windfall Clause. One historical example is the informal "5 percent club" in the 1960s and '70s, whose members were leading US firms that donated 5% of pre-tax earnings to charities.[177] Although collaboration among competitors can sometimes trigger antitrust scrutiny, the Federal Trade Commission has acknowledged that collaborative CSR efforts can add social value without harming competition.[178]

Increasingly, leading firms have entered into social responsibility efforts with civil society organizations, not simply industry consortia. Within this institutional form, a civil society organization administers rules for prosocial behaviors such as sustainable forestry, fair pay, and charitable donations. Compliant firms receive exclusive benefits, often through consumer or employee demand for a certification mark on products.[179] The Forest Stewardship Council,[180] Rainforest Alliance,[181] Fair Trade,[182] and (Red)[183] are all well-known examples of such partnerships. The structure of the Windfall Clause follows this lineage of civil society initiatives to promote CSR.

Corporate social responsibility efforts are also regularly contingent on some threshold, albeit one lower than that of the Windfall Clause. Commonly, this threshold is whether the firm turns a profit; if so, then the firm will pursue CSR efforts. This has sometimes become a central focus of the business. Newman's Own, Inc. is a well-known example that markets its food products based on its charitable efforts. All profits from the corporation are passed to its sole owner, the Newman's Own Foundation, which supports various charitable causes.[184]

Another model does not focus the business solely on charity, but rather passes some profits to charity. Industrial foundations—nonprofit organizations that own a controlling interest in a for-profit business—are one such form.[185] Although there is no legal requirement for such foundations to serve charitable purposes, profit-dependent dividend revenue from foundation-owned stock commonly funds charitable giving.[186] When Google filed for its IPO in 2004, they announced their intention to "contribute significant resources to the [Google F]oundation, including employee time and approximately 1% of Google's equity and profits in some form."[187] They followed through on this commitment in 2005.[188]

Finally, there is precedent specifically among AI firms. One notable AI firm, OpenAI, recently restructured from a nonprofit to a "capped-profit" company.[189] They state that this is designed to allow them to raise the capital necessary to build powerful AI systems while also maintaining their philanthropic mission.[190] The Windfall Clause follows from these precedents of profit-threshold contingent charitable commitments.

B. Precedents from Public Windfall Governance

A number of public policy precedents address the unique challenges posed by windfall profits in other industries.

Wealthy countries use sovereign wealth funds to manage their windfalls for the benefit of their present and future citizens. Governments endow these funds with profits from natural resources or currency reserves from international trade surpluses. Their scale is massive, surpassing

that of even the Windfall Clause: sovereign wealth fund assets around the world amount to well over $8 trillion, with the Norwegian Government Pension Fund Global alone amounting to over $1 trillion.[191] These funds often aim to benefit future generations through either wealth preservation and transfer (in developed countries) or development (in developing countries).[192] The Norwegian fund, for instance, states that it acts "as a financial reserve and as a long-term savings plan so that both current and future generations get to benefit from our [current] oil wealth."[193] Although there are key differences between sovereign wealth funds and the Windfall Clause*—the scope and scale of distribution for one—they both serve to improve the economic allocation of windfall resources and do so at a massive scale.[194]

While at a much smaller scale than the proposed Windfall Clause, local American governments sometimes use a "windfall clause" when subsidizing private sports arena construction. Contracts between the subsidizing government and the franchise sometimes contain a clause that ensures public benefit if a windfall sale of the "home team" occurs, as compensation for the public subsidization.[195] Other existing policies subject especially profitable companies[196] and natural monopolies[197] to stricter governmental control for the public good.

Thus, the Windfall Clause has precedents in existing public policies for dealing with various previous windfalls.

C. Precedents from Personal Philanthropy

Many expect very wealthy individuals to share much of their financial windfall, with those who hold onto most of their wealth facing persistent criticism.[198] Bill and Melinda Gates and their foundation epitomize—and indeed helped shape—this norm. They and Warren Buffett launched the Giving Pledge: a nonbinding commitment by billionaires to donate more than half of their wealth.[199] A similar project—the Founders Pledge—binds technology entrepreneurs to give some percentage of their exit proceeds to charity upon selling their shares.[200] The Windfall Clause applies this expectation of wealthy people to especially profitable firms.[201]

In practice, studies show that ex ante commitments encourage greater charitable giving due to loss aversion (among other reasons).[202] Although the relevant psychological studies may not apply to firms as such, they can apply to decision-makers within firms.[203] Thus, the ex ante nature of the Clause makes philanthropic actions more likely to materialize, compared to philanthropic acts after the acquisition of a windfall.

In conclusion, though the Windfall Clause is a unique policy proposal, its various aspects draw heavily from existing corporate, policy, and philanthropic practices.

* Notably, in sovereign wealth funds, the payments into the fund may occur over time and do not follow a threshold model. Such funds are initiatives of national governments, not private industry.

V. OBJECTIONS, ALTERNATIVES, AND LIMITATIONS

In this section, we identify—and where appropriate, respond to—possible objections and limitations to the Windfall Clause as proposed, and analyze alternative mechanisms[204] for achieving the Clause's aims. The points that we address here are those we consider most worthy of careful attention; we certainly invite further discussion on these points, as well as those not raised in this section.

To summarize, we do not claim that the Windfall Clause is a perfect mechanism, let alone the only mechanism of its kind that should be pursued. Indeed, we are not very confident that, all things considered, it ought to be adopted. However, we do believe that the Clause is promising enough to warrant further investigation, notwithstanding even its greatest shortcomings. To the extent that the Windfall Clause needs to be revised, or not pursued, because of its shortcomings, we believe that the public conversation needed to recognize this will be valuable itself. We look forward to discussing and collaborating with others in the construction of the Windfall Clause, and other socially beneficial legal commitments.

A. Reasons Why the Windfall Clause Might Not Work

A.1. "The Windfall Clause will never be triggered."

Some might doubt that the Clause will ever be triggered—i.e., they may be skeptical that any firm will attain profits sufficient to trigger nontrivial obligations under the Clause. This is not so much a reason why the Clause may not work—indeed, by design, the Clause is only intended to trigger under certain circumstances. Rather, this could be a reason for questioning the relevance of the Windfall Clause as a tool for addressing the negative impacts of AI development.

To be sure, the Windfall Function we defined in Section II(A) only results in significant obligations under the Clause if profit reaches very high, currently difficult-to-imagine amounts. However, we view these profit levels as possible in the future given the unprecedentedly transformative nature of AI. We also hope that further refinements to the Windfall Function can yield obligations that are neither too likely nor too unlikely; the Windfall Function offered above is a first attempt, not the commitment-ready form.

Further, it is also worth noting that the Clause is not intended to be the sole policy mechanism to address potentially negative outcomes resulting from an AI-driven economy. Rather, it is best suited for the particular scenario in which windfall profits are achieved. Other measures are likely desirable to address other (i.e., non-windfall) issues which may arise in scenarios of AI development which do not reach levels of windfall profits, but nevertheless warrant attention and intervention. We certainly hope that such mechanisms are explored as complements to the Windfall Clause.

A.2. "Firms will find a way to circumvent their commitments under the Clause."

One might worry that, as the Clause begins to impose substantial obligations, firms will seek to exploit some "loophole."[205] This is a reasonable concern; indeed, as the obligations of the Clause increase there could be trillions of dollars on the line, creating strong incentives for firms to avoid their obligations. Great care should be taken to design a clause and supporting mechanisms that increase the likelihood that commitments cannot be circumvented. Here, we address specific ways firms might attempt to sidestep their commitments, and suggestions for how we may mitigate these risks.*

* As with other legal analysis in this report, the focus here is on American law. If other law is applicable, additional legal analysis might be warranted.

A.2.i. "Firms will evade the Clause by nominally assigning profits to subsidiary, parent, or sibling corporations."

The worry here is that signatories will structure their earnings in such a way that the signatory itself technically does not earn windfall profits, but its subsidiary, parent, or sibling corporation (which did not sign the Clause) does.[206] Such a move could be analogous to the "corporate inversion" tax avoidance strategy that many American corporations use.[207] Thus, the worry goes, shareholders of the signatory would still benefit from the windfall (since the windfall-earning corporation remains under their control) without incurring obligations under the Clause.

We think that the Clause can mitigate much of this risk. First, the Clause could be designed to bind the parent company and stipulate that it applies not only to the signatory proper, but also to the signatory's subsidiaries.[208] Thus, any reallocation of profits to or between subsidiaries would have no effect on windfall obligations.* Second, majority-owned subsidiaries' earnings should be reflected in the parent corporation's income statement,[209] so the increase in the subsidiary's profits from such a transfer would count towards the parent's income for accounting purposes.† Finally, such actions by a corporation could constitute a number of legal infractions, such as fraudulent conveyance[210] or breach of the duty to perform contracts in good faith.[211]

A.2.ii. "Firms will evade the Clause by paying out profits in dividends."

The worry here is that a corporation would distribute windfall to shareholders before meeting its obligations under the Clause. If it were to do so, it would be breaking the law. Dividends must be paid out of a corporation's surplus,[212,213] which should account for liabilities such as a legally binding Windfall Clause obligation.[214] This concern would, however, stand if the Windfall Clause is a non-legally binding commitment at the point at which the obligations vest.

A.2.iii. "Firms will sell windfall-generating AI assets to a firm that is not bound by the Clause."

One might also worry that a firm will sell its windfall-generating AI assets to another technology firm that has not signed the Clause, thereby effectively shifting the windfall from a bound firm to an unbound one. However, unless the buyer was offering the firm more for the assets than the assets' net present value to the firm, this is unlikely to be a good business decision for a firm. As such, taking such a decision to sell would be considered a failure to fulfill one's fiduciary duties.

To see why, imagine that XYZ Corp. is a successful firm with windfall-generating AI assets. For the purposes of simplified illustration, further imagine that those assets have a net present value to XYZ of $X trillion. Suppose further that the net present value to XYZ's shareholders (i.e., after accounting for Windfall Clause obligations) was $X/3 trillion (i.e., one-third of the value goes to shareholders, while two-thirds goes to the Windfall Clause). In considering whether to sell to an unbound firm, XYZ Corp., as the owner of the assets, should not accept any payment less than $X trillion. This is because the sale price of the assets will count towards XYZ's profit, thus the Windfall Clause will apply to its proceeds from the sale. So, XYZ's shareholders would be worse-off for any sale price of less than the net present value of windfall-generating AI assets. If there is a sale, then, the Windfall Clause would receive as much or more present value than it would if the signatory-firm held onto the assets.

Even given our hopes that these specific means of circumvention could be prevented, others probably remain. Thus, we do not wish to project confidence that there will be no way for a

* To avoid the problem of the parent potentially splitting profits between subsidiaries and thereby avoiding triggering the Clause, the Clause could specify that the proper measure of the parent's profits is the sum of its profits and the profits of all its subsidiaries.

† The reason this is not a barrier to corporate tax avoidance is that "[corporate e]arnings are taxable wherever they occur (worldwide), but taxes are not due until or unless the earnings are brought back, or repatriated, to the U.S. Further, U.S. corporations receive credit for foreign taxes paid to avoid double taxation." Neely & Sherrer, *supra* note 207, at 3 n.1. Thus, as long as the earnings are kept offshore, taxes are not due on them until repatriated. The Windfall Clause need not suffer from this problem, since it could stipulate that it applies to profits of all subsidiaries, wherever earned.

signatory to circumvent the Clause. We expect that much future work will be aimed at drafting and testing various potential Windfall Clauses to improve its reliability as a commitment mechanism.

A.3. "No firm with a realistic chance of developing windfall-generating AI would sign the Clause."

Would AI firms actually sign the Clause? While there remains significant uncertainty around this question, there are several reasons to believe that it might be in firms' interests to become signatories to the Clause—i.e., they assess that the benefits of doing so outweigh the costs.

First, as mentioned above,[215] the counterfactual case for a firm (i.e., what would have happened if they developed windfall-returning AI but had not signed the Clause) is plausibly *not* business-as-usual. Instead, dramatic increases in concentration of wealth and power[216]—combined with mass unemployment[217]—could yield immense political pressure to tax or expropriate windfall-generating assets. Firms would have significant liability in this scenario as well—plausibly even greater than that imposed by the Clause. However, from a public and employee relations perspective, the Clause may be more appealing than taxation because the Clause is a cooperative, proactive, and supererogatory action. So, to the extent that the Windfall Clause merely replaces taxation, the Windfall Clause confers reputational benefits onto the signatory at no additional cost.

Even if it were the case that the total expected benefits and costs from the Windfall Clause sum to zero, the benefits from improved employee and public relations would accrue earlier (ideally, continuously from the signing date) than the costs. Thus, the firm would be able to capitalize on the benefits (through increased revenues and decreased costs) earlier than it would bear any costs. If those benefits and costs are equivalent atemporally, the earlier benefits would still outweigh the costs due to the time value of money.*

For these reasons, we expect that the Clause could be positive-sum in expectation for signatories. However, this is ultimately an empirical prediction, and we are mindful of many other factors that affect firms' calculations which we may have failed to account for.

This objection does make an important point, though: the clock is ticking. Key leaders at major AI firms have acknowledged the importance of examining the distributional effects of AI.[218] The Asilomar AI principles, which includes "Shared Prosperity: [t]he economic prosperity created by AI should be shared broadly, to benefit all of humanity,"[219] have been endorsed by leading AI researchers at Google, OpenAI, Facebook, and many others.[220] While these are all positive signs, this sentiment could change in the future. If a firm gains a clear prospect of capturing windfall profits, then the Clause will become more costly to them in expectation, making it less likely that they will sign it. Thus, we view this as a time-sensitive project.

A.4. "If the public benefits of the Windfall Clause are supposed to be large, that is inconsistent with stating that the cost to firms will be small enough that they would be willing to sign the Clause."

The argument being put forward here is that the following are inconsistent: (1) the Windfall Clause is a valuable project,[221] and (2) the Windfall Clause has low gross cost to firms.[222] This argument rests on the assumption that that the public policy benefits from the Windfall Clause will be approximately equal to the costs to signatories. Thus, the claim that the Windfall Clause has simultaneously large public benefits and small private costs is problematic.

We disagree, but the reasons for this disagreement are nuanced. First, some of the Windfall Clause's public benefits do not depend on the Clause actually ever being triggered. For example, the Clause could stabilize important societal and international relations[223] and promote

* However, to the extent the Windfall Clause replaces ex post taxation, there may be public policy costs from making this tradeoff ex ante, without the benefit of knowledge of future economic conditions.

beneficial norms regarding the development of AI,[224] all while incurring no or minimal monetary costs to its signatories. Knowledge that signatories have adopted the Clause could, in itself, be sufficient to advance many of these goals, even if windfall profits are unlikely. As an example of this first effect, knowledge that the Windfall Clause is in place could have societally stabilizing effects by diminishing popular concerns about harms from windfall scenarios, even if those scenarios are very unlikely. Individuals that expect to be better situated to deal with adverse consequences from AI will likely worry less about its negative effects.[225] The Windfall Clause could give individuals a secure future in a world with windfall-generating AI, thus somewhat reducing their anxiety about its development.

Second, the time value of the Clause is different from shareholders' perspective than from a public policy perspective. Shareholders should exponentially discount the expected cost of the Windfall Clause, as they would for other future debts.[226] However, it is inappropriate to use the same discount rate for many of the non-monetary goods that the Windfall Clause could bring.[227] For example, entrenched extreme concentrations of wealth and power could cause the long-term future to be worse than it could be, and the badness of this is almost entirely independent of how quickly such concentration occurs in the near-term. Thus, the private costs of the Windfall Clause should be more heavily discounted than its public benefits, leading to these benefits accrued to the public being greater than the private costs to signatories.

B. Unintended Consequences of the Windfall Clause

B.1. "The Windfall Clause reduces incentives to innovate."

Some might worry that the Clause will diminish signatories' incentives to innovate because it caps the amount of private benefit that a given firm can gain in the future. Indeed, it is difficult to predict the overall effects of something like the Windfall Clause on the incentives of signatories in general, much less of specific AI firms composed of different company cultures and business models.

Nevertheless, we find several reasons to believe that the Clause would not significantly reduce incentives to innovate. For one, the Clause only targets extremely unlikely profit levels[228] that therefore make up very little of signatories' present expected value.[229] Thus, diverting some of those profits through the Clause should not seriously affect firms' present incentives. Further, by capping firm obligations at 50% of marginal profits,[230] the Clause leaves room for innovation to be invested in even at incredibly high profit levels.[231]

However, if the Clause did reduce incentives to innovate, this might be an acceptable outcome for several reasons. Firstly, we expect firms to agree to the Clause only if it is largely in their self-interest.[232] Thus, the Clause is arguably in firms' self-interest *overall* (i.e., even accounting for any reduction in incentives to innovate) because it builds consumer goodwill,[233] improves employee relations,[234] and reduces political risk.[235] Secondly, unbridled incentives to innovate are not necessarily always good, particularly when many of the potential downsides of that innovation are externalized in the form of public harms.[236] The Windfall Clause attempts to internalize some of these externalities to the signatory, which hopefully contributes to steering innovation incentives in ways that minimize these negative externalities and compensate their bearers.[237] Finally, one must consider that under windfall scenarios the gains from innovation are already substantial, suggesting that globally it is more important to focus on distribution of gains than incentivizing additional innovation. Indeed, the public has and will continue to heavily subsidize the development of AI through, e.g., basic research, educating engineers, developing many of the infrastructural necessities for AI, etc., so it seems appropriate that they recoup some of that investment, even if that means dampening incentives to innovate on the margin.

B.2. "The Windfall Clause will shift investment to competitive non-signatory firms."

The concern here is that, when multiple firms are competing for windfall profits, a firm bound by the Clause will be at a competitive disadvantage because unbound firms could offer higher returns on new capital. That is, investors would prefer firms that are not subject to a "tax" on their profits in the form of the Windfall Clause. This is especially bad because it could mean that more prosocial firms (i.e., ones that have signed the Clause) would be at a disadvantage to non-signatory firms, making a prosocial "winner" of an AI development race less likely.[238]

This is a valid concern which warrants careful consideration. Our current best model for how to address this is that the Clause could commit (or at least allow for the option of) distributions of equity,* instead of cash. This could either take the form of stock options or contingent convertible bonds. This avoids the concern identified by allowing firms to, for example, issue new, preferred shares which would have superior claim to windfall profits compared to donees. This significantly diminishes the concern that the Clause would dilute the value of new shares issued in the company and allows the bound firm to raise capital unencumbered by debt owed under the Clause.† Notably, firm management would still have fiduciary duties towards stockholding windfall donees.

Even with this as a potential solution, the Windfall Clause could still undesirably alter firms' and investors' incentives. Therefore, careful microeconomic analysis of various Windfall Functions and other Clause provisions should constitute a large portion of future investigations. If a Windfall Clause would unavoidably create perverse incentives, it may not be a project worth pursuing.

B.3. "The Windfall Clause draws attention to signatories in undesirable ways"

One might worry that the Clause could draw undesirable attention to its signatories. Specifically, one might worry that agreement to the Clause implies that the signatory believes that "there is a realistic chance that in the next few decades [that signatory] will be a colossal [i.e., windfall-generating] company . . . ."[239] Call this the "Colossus Claim."[240] Endorsement of the Colossus Claim might generate criticism from various stakeholders (who would perceive the Claim as "hype and bluster")[241] and politicians (who would subject signatories to higher scrutiny if they perceive them to have endorsed the Claim).[242]

As a general note, one should not infer endorsement of the Colossus Claim from agreement to the Clause. Indeed, due to management's fiduciary duties to shareholders,[243] a firm's willingness to agree to the Clause should be *negatively* correlated with its probability of achieving windfall profits in the near future. Thus, perhaps counterintuitively, agreement to the Clause is evidence *against* the Colossus Claim (under the reasonable assumption that managers act within their fiduciary obligations). However, we acknowledge that the public perception attached to firms may be informed by less rational, more instinctive reasoning which may lead to signatories being tagged as more in support of the Claim than is accurate. How the Windfall Clause is framed in the public eye is thus an important element to think through.

For other important stakeholders—for example, investors, employees, competitors—the Clause does not add much new public information. Prominent tech entrepreneurs have already openly acknowledged the possibility of developing AI which could deliver extreme benefits, as have several academic institutions and investors.[244] Agreement to the Clause thus supplies little new information on the firm's AI prospects, except perhaps a downwards revision for the reasons just mentioned.

* To avoid problems associated with selling large stakes in companies, the equity should be nonvoting. *See supra* note 106.

† This approach is admittedly more difficult, however, because it could require amending a corporate charter to create new classes of stock.

For policymakers, the fact that the Clause acknowledges the (possibly very large) upsides from AI is not necessarily a reason for opposition. Indeed, political support has remained consistent for (domestic) firms achieving major technological breakthroughs.[245] There remains a risk that if policymakers infer support for the Colossus Claim from the signing of the Clause, signatories will become central targets if the state were to turn to tools such as nationalization in an attempt to capture these benefits.

B.4. "The Windfall Clause will lead to moral licensing."

Another objection might be that the Windfall Clause will have a moral licensing effect—i.e., it will make signatories feel as though they have done something good, thereby making them feel more comfortable acting unethically in other ways.[246,*]

One should always be wary of possible licensing effects. However, at the same time, many policies are worth pursuing notwithstanding the possibility for licensing. The proper question, then, is whether the Clause is worthwhile despite the possibility for moral licensing. For the reasons given throughout this report, we believe that it is.

None of this, however, should dampen constructive engagement between civil society groups and AI firms. Indeed, part of the success of the Clause depends on governmental and popular pressure on signatories to abide by its terms and spirit. At the same time, adoption of the Clause would be a significant and credible commitment for signatories, and therefore should be celebrated.

C. Problems with the Design of the Windfall Clause

C.1. "Rule of law might not hold if windfall profits are achieved."

The success of the Windfall Clause depends on its enforceability as a matter of law. No contract, however, is ever entirely safe from successful legal challenge. Thus, the Windfall Clause would need to be drafted very carefully to maximize the probability of enforceability; this is a complex task which requires more thorough treatment beyond what we have presented in this report.

A related set of concerns comes from the fact that a firm who would have obligations under the Clause would, by definition, be one of the most powerful agents in the world. Thus, the Clause might fail if the signatory uses its financial and technological advantages to undermine enforcement of the Clause. Examples of such behaviors could include:

- Defeating any legal challenges to its actions;
- Finding unanticipated "loopholes" in the Clause, or ways to evade it; or
- Bribing, threatening, or otherwise manipulating lawmakers, judges, witnesses, or law enforcement.

By diminishing successful firms' power, the Clause mildly mitigates some of these risks, but it cannot completely prevent them. Some additional mechanisms that could help reduce these issues include:

- Selecting a venue or venues and governing law or laws with wide jurisdictional reach;
- Selecting a venue or venues with stable and well-functioning enforcement mechanisms;
- Giving windfall distributors an ownership stake in the signatory;
- Stipulating that firms will pay for distributors' attorney fees; or
- Giving distributors access to any AI that could give signatories an advantage in litigation.

* A related objection, also addressed in this subsection, is that the Windfall Clause will buy signatories political cover to get away with unethical behavior.

Still, the rule of law depends on factors far beyond the purview of this single project. Further, the fact that the rule of law may not hold with the development of transformative AI is a broader concern beyond the Windfall Clause. We look forward to working with others looking to ensure that the rule of law persists even if truly transformative AI is achieved.

C.2. "The Windfall Clause operates like a progressive corporate income tax, and the ideal corporate income tax rate is 0%."

The Windfall Clause could be designed to operate like a progressive corporate income tax.[247] Some commentators argue that the ideal corporate tax rate is 0%.[248] One common argument for this is that corporate income tax is not as progressive as its proponents think because corporate income is ultimately destined for shareholders, some of whom are wealthy, but many of whom are not.[249] Better, then, to tax those wealthy shareholders more directly and let corporate profits flow less impeded to poorer ones.[250] Additionally, current corporate taxes appear to burden both shareholders *and*, to a lesser extent, workers.[251]

The Windfall Clause admittedly shares this flaw: it is an imperfect tool for differentially targeting wealthy shareholders of successful firms. However, the Windfall Clause should not be regarded as a substitute for individual taxation or giving commitments. Rather, the Windfall Clause is a firm-level complement to individual-level obligations.

Giving commitments targeted to firms have some desirable properties. Firms may have greater incentive to appear socially responsible, and thus are more amenable to the kinds of social pressure that a robust discussion about the Windfall Clause would bring. The pool of potentially windfall-generating firms is much smaller and more stable than the number of potential windfall-generating individuals, meaning that securing commitments from firms would probably capture more of the potential windfall than securing commitments from individuals. Thus, targeting firms as such seems reasonable. However, we encourage economic analysis of the Clause to determine whether its incidence would unfairly burden workers or have other undesirable effects not foreseen here.

C.3. "The Windfall Clause undesirably leaves control of advanced AI in private hands."

One might reasonably propose that all humans ought to have a say in the use of such a powerful tool as windfall-generating AI. The Windfall Clause distributes windfall profits from AI but does not distribute control of it.* This risks exacerbating the current, plausibly undesirable distribution of power between firms with the technology and everyone else, leading to imbalanced capacities to influence the development trajectory of transformative AI. This may raise objections, particularly from those to whom the concentration of control of such a publicly consequential technology in private actors' hands is unacceptable.

We are sympathetic to these concerns, and indeed, acknowledge that we largely avoid discussing who ought to control advanced AI, and have instead focused on who ought to benefit from it. The question of control is a more complex one, invoking different legal, economic, and philosophical issues which deserve thorough deliberation and analysis beyond the scope of this report. We hope that focus on profit here will serve as a useful starting point for large future conversations about control of this powerful technology. It may be that such discussion reveals that full distribution of control of advanced AI is desirable. But this report neither makes nor examines that claim closely. Instead, we rest on what we hope is a modest premise that the fruits of one of the most transformative technologies ever—the culmination of human ingenuity—ought to be shared widely.

We do, however, hold firmly that the mechanism by which windfall profits is distributed should be legitimate and accountable, as outlined in Section II. This is distinct from the question of

* To avoid certain legal issues, even if the windfall is distributed as equity rather than debt, such equity would be nonvoting. *See supra* note 103. Therefore, even such a Clause would not distribute control of the AI.

who should control the technology. Thus, while the question of who should control the technology is left for further discussion, the question of how decisions are made as to who should benefit from its bounty is affirmatively responded to in support of the appropriate distribution of decision-making power beyond private hands.

D. Alternatives to and Variants of the Windfall Clause

D.1. "Windfall profits should just be taxed."

Some might object to—or be uneasy with—extra-governmental nature of windfall profits, perhaps arguing that the windfall should instead be taxed and spent by democratically accountable governments.[252]

Taxation certainly has many merits relative to the Windfall Clause. Distributive measures that arise from democratic institutions are often more legitimate than extra-governmental ones.[253] Additionally, given states' substantial resources and powers, compliance with taxation might be higher than compliance with the Clause. Given these merits, taxation remains a critical part of the larger AI policy conversation, and the Windfall Clause is not intended to be a substitute for taxation schemes. We also note that, as a private contract, the Windfall Clause cannot supersede taxation.* Thus, if a state wants to tax the windfall, the Clause is not intended to stop it.† Indeed, taxation efforts that broadly align with the goals and design principles of the Windfall Clause are highly desirable.‡

However, there are also reasons to prefer the Clause to taxation under some circumstances. Particularly in the near-term, an important consideration is political tractability. Given current realities, we do not anticipate that legally enforceable taxation and global distribution of AI windfall is politically feasible, whereas beginning a conversation around a voluntary commitment such as the Windfall Clause may be more so. Furthermore, all else equal, it is possible that signatories would be more receptive to donations rather than taxation in order to reap the benefits of consumers, employees, and governments perceiving signatories to be socially responsible.[254] In other words, committing to distribute windfall through an exceptional act of philanthropy might benefit signatories in a way that merely complying with tax law does not. This makes the Clause more aligned with the incentives of firms and thus plausibly more tractable compared to taxation.

The Clause is also plausibly more likely to be robust to firm lobbying efforts. In particular, if preemptive taxation proves intractable, attempts to implement taxation post the development of windfall-generating AI will be much more difficult given the scale of the resources that a successful firm will have at their disposal.[255] An ex ante commitment such as the Windfall Clause hopes to circumvent this issue, although there are certainly still comparable risks of well-resourced firms attempting to undermine enforcement of the Clause.§

Further, a primary advantage of the Clause over most current forms of taxation is that the

* And, conversely, companies cannot choose the corporate tax rate, but they can choose their own philanthropic efforts—i.e., corporate philanthropy is more tractable than tax reform.

† We leave aside a more complicated discussion of whether and how a government might be able to tax windfall profits without allowing the taxpaying corporation to deduct Clause obligations from its taxable income.

‡ If a government taxes the windfall above current levels—either before or after deducting Clause obligations—signatories might be excessively burdened. We therefore tentatively recommend that the Clause obligation can be superseded by an *appropriate* tax, to prevent such excessive taxation of the windfall. However, it is important to think through what would be an appropriate tax. For example, we do not think that a poorly thought through national tax should supersede a well-designed and functioning global Windfall Clause, but we do think that a well-designed and broadly legitimate global tax should supersede the Clause. This will be a very difficult distinction to encode in the Clause itself. Rather than attempt to do so ex ante, we recommend specifying an ex post mechanism through which a signatory can petition to have the Clause superseded. *See generally* O'Keefe, *supra* note 93. An example of this could be appointing a global independent legal-ethical body that would judge whether a given tax should supersede (part of) the Windfall Clause. *Cf. id.* at 25–29 (discussing ex post third-party resolution).

§ Readers might object that this claim is inconsistent with predictions made elsewhere in the report that there could be significant public backlash to windfall profits if unfairly distributed, and that this backlash might result in policy adverse to firms. Of course, both backlash and firm resistance thereto are possible. However, the objection is well-founded insofar as one of these forces is likely to win out in the end. It is difficult to predict which "side" would likely "win." However, the conflict process *itself* is likely undesirable for all parties, since it will sink resources into a largely zero-sum fight. Thus, insofar as the Windfall Clause mitigates this conflict altogether, it is positive-sum.

Clause's distribution mechanism is international rather than national. From a public policy perspective, global philanthropy can better promote international cooperation and benefit, whereas national taxation efforts would tend to overwhelmingly benefit the taxing nation's citizens rather than being globally (and therefore more effectively)[256] distributed.[257,*] To ensure that their citizens benefit, then, non-host nations must find a way to tax (or otherwise appropriate) the benefits from AI. Not all nations will be able to do this, however, and hence the host nation will reap most of the rewards of hosting a successful firm. The Windfall Clause could avoid this by distributing benefits globally, not just overwhelmingly to the citizens of the firm's nation.

Finally, the most effective philanthropic initiatives can often accomplish more good per dollar than national governments,[258] in part given the existing inequality between nations paired with the strength of parochial interests.[259] Indeed, it appears difficult to convince national governments to fund more-cost-effective international projects over less-cost-effective domestic ones (as evidenced by the fact that they often choose not to under current conditions).[260] As such, if one were concerned about the efficacy of windfall distribution, an extra-governmental initiative such as the Windfall Clause may in some circumstances be more successful on this measure. Further, if it were the case that national governments could spend windfall profits more effectively than private philanthropy, then distributors could simply remit payment to those national governments.

To summarize, we are sympathetic to concerns about charitable accountability, legitimacy, egalitarianism, and effectiveness. It is somewhat unclear whether the Windfall Clause is preferable to taxation on these measures, although there are reasons to believe that it could be under current circumstances. Nevertheless, we are certainly open to taxation policies that accomplish the Clause's goals.

D.2. "We should rely on antitrust enforcement instead."

One might hope that antitrust law will remedy any potential AI-services monopoly.[261] Indeed, past corporate behemoths faced intense antitrust scrutiny, leading to several large dissolutions.[262] If that did happen, the necessity of the Clause would indeed be diminished[263] because it would address some issues associated with concentration of economic power.

However, it is not obvious that antitrust would be an appropriate or available tool for addressing a windfall-generating firm. The most important reason for this is that mere size or acquisition of monopoly power is not actionable under current antitrust law:

> The mere possession of monopoly power, and the concomitant charging of monopoly prices, is not only not unlawful; it is an important element of the free-market system. The opportunity to charge monopoly prices—at least for a short period—is what attracts "business acumen" in the first place; it induces risk taking that produces innovation and economic growth. To safeguard the incentive to innovate, the possession of monopoly power will not be found unlawful unless it is accompanied by an element of anticompetitive *conduct*.[264]

Thus, current[265] American antitrust law could not address a mere AI monopoly unless it had acquired that monopoly via prohibited anticompetitive conduct. If, as we hypothesize,[266] AI services are a natural monopoly, then an AI-services monopoly could result without prohibited anticompetitive conduct. The Windfall Clause is most useful in such a scenario.

Ex post regulation of that monopoly might be desirable.[267] However, depending on the circumstances, the Windfall Clause could make such regulation either:

1. Less necessary, because harms from an AI monopoly would be mitigated[268] and benefits from monopoly profits widely shared;[†] or

* Furthermore, since already-wealthy nations will probably be the homes of windfall-generating companies, taxation would most likely redistribute windfall almost entirely *within* the wealthiest nations.

† Citizens would then have to weigh the costs of regulation (in the form of decreased windfall) against the benefits thereof (the erosion of monopoly).

2. Easier, because the wealth of the monopoly and its shareholders (and therefore their political power) would be diminished.*

Relatedly, part of the appeal of the Windfall Clause (as opposed to new laws to deal with potential AI market power) is that it could solve some problems associated therewith[269] with minimal government intervention. Thus, under some circumstances, the Clause could be a more tractable intervention than promoting new monopoly laws.

D.3. "We should establish a Sovereign Wealth Fund for AI."

Another approach is to redistribute capital. One such proposed mechanism has been called the "Sovereign Automation Fund (SAF)",[270] analogizing to the sovereign wealth funds we discuss above.[271] The idea is that some public-benefiting group (either a government or a nonprofit) would acquire shares in AI companies.[272] That group would then distribute inalienable voting rights to some set of constituents (e.g., all people in the world) to provide democratic control over the pooled shares.[273] The constituents would also receive dividends from the SAF.[274]

The SAF has advantages and disadvantages relative to the Windfall Clause. One advantage is that the SAF immediately affects all signatories.[275] This is an advantage because it ensures that signatories are habituated to complying with the SAF and does not distort new investors' returns, whereas the Windfall Clause does not affect signatories until they reach very high profit levels.[276] Thus, as a signatory approaches such high profits, it might "more strongly resist[]" the Clause than the SAF.[277]

The SAF has many other desirable features, some of which it shares with the Windfall Clause. First, like the Windfall Clause, the SAF would be "based on an asset that should appreciate as automation and wealth in AI companies grows . . . ."[278] Second, the SAF as described above would give everyone actual and inalienable partial control over AI, not just money.[279]

However, the SAF also has disadvantages. The immediate nature of the SAF raises its costs. Earlier obligations are more probable obligations and are therefore a higher expected cost to the signatory. The fact that Windfall Clause obligations vest at only very high profit levels dramatically reduces the expected cost of those obligations, and therefore might make firms more amenable to the Clause.

Furthermore, corporate directors are likely to react negatively to an SAF attempting to acquire a significant voting share of the corporation, and could be incentivized to use a "poison pill" to prevent such an acquisition.[280] Poison pills commonly prevent acquisition of shares as low as 10–15%[281]—far lower than an SAF would need to globally distribute meaningful control over AI firms. Furthermore, acquiring and actively managing significant stakes in multiple competing AI firms could trigger antitrust scrutiny of the SAF.[282]

A related downside to the SAF is that buying meaningful stakes in relevant technology companies would be extremely expensive. Indeed, for some relevant corporations, such as Alphabet, this would be impossible due to their corporate structure.[283]

Governmental acquisition of firm shares might avoid these problems.[284,285] However, this would face intense political opposition as a perceived violation of liberal international norms against expropriation.[286] Further, such an acquisition could be subject to legal challenge as a taking requiring just compensation to the firm (at least if done in the US).[287]

A Windfall Clause that uses an SAF-like structure for distributing the windfall might ultimately be desirable. However, as with other specific distribution proposals, we leave that discussion for a later date.

* Public-benefit organizations funded by the Windfall Clause could also operate as an effective countermeasure to lobbying by the monopoly.

D.4. "We should implement a Universal Basic Income instead."

One possible variant* of the Windfall Clause is a firm-funded Universal Basic Income (UBI) program.[288] Certainly, one of the Windfall Clause's major advantages is that it retains the flexibility to distribute funds in whatever way is globally welfare-maximizing. While it is plausible that a UBI could be a suitable distribution mechanism in this vein, it is unlikely to be sufficient.

A UBI program could be a great boon for humanity. For example, evidence-based charity evaluator GiveWell[289] finds that, commonsensically, cash transfers are very effective at ameliorating poverty.[290] Distributing cash also has intuitive appeal and might avoid some of the accountability and oversight problems associated with other distribution proposals.

However, UBI has some shortcomings. One is that some programs appear to do more good than UBI-like cash transfers. GiveWell's recent cost-effectiveness analyses conclude that several health interventions—such as deworming and malaria prevention—are more cost-effective than cash grants.[291] Thus, implementing a UBI-only distribution plan could fail to maximize total welfare.

A related reason why a UBI-only windfall distribution plan would be suboptimal is that UBI is unlikely to adequately spur investment in certain goods, such as public goods.[292] Windfall distributors charged with maximizing total welfare might therefore be able to achieve better results through investing in public goods rather than through cash grants. However, new innovations in public goods funding like "Liberal Radicalism" could capture many of the benefits of UBI while also incentivizing public goods funding.[293] We are enthusiastic about the possibility of using the Windfall Clause to experiment with this and other modes of distributing benefits both widely and collectively.

Finally, UBI seems necessarily (as a "universal" program) poorly targeted towards the particular harms that advanced AI could cause.[294]

Thus, if UBI is to be a part of the distribution plan, it should not be its entirety.

* Strictly speaking, this could be a "permutation" rather than a true "alternative" because the Windfall Clause could simply implement a UBI program. Indeed, our Windfall Trust proposal, *infra* app. II(B), approximates this.

CONCLUSION

Artificial Intelligence may be poised to fundamentally change the nature of the global economy. While these changes may bring unprecedented bounty and improve quality of life across the world, they also involve risks and potential turbulence. One challenge, then, is the substantial uncertainty in predictions about the timing, nature, and effects of these changes. In the face of such uncertainty, sound preparation demands that we invest in solutions for multiple scenarios.

One possible future is a scenario in which advanced AI services result in unprecedented "windfall" profits that accrue to a very small number of actors. This scenario is highly undesirable, not only because humanity at large will have borne the risk of innovations along the way, but also because the creation of such wealth might be accompanied by mass unemployment and other occurrences that increase human suffering.

The Windfall Clause aims to address the downsides of such a scenario, if the development of truly windfall-generating AI were to occur. With further refinement, the Windfall Clause could be a significant, credible, and tractable way for AI firms to direct their inventions towards the enrichment of humanity generally and still reap substantial rewards for doing so.

The distributive implications of AI remain both pressing and unclear. We hope that this report contributes an ambitious and novel policy proposal to an already rich discussion on this subject. We look forward to engaging with respondents, whose wide-ranging priorities and concerns will strengthen this effort. More important than this policy itself, though, we look forward to continuously contributing to broader conversations on the economic promises and challenges of AI, and how to ensure AI benefits humanity as a whole.

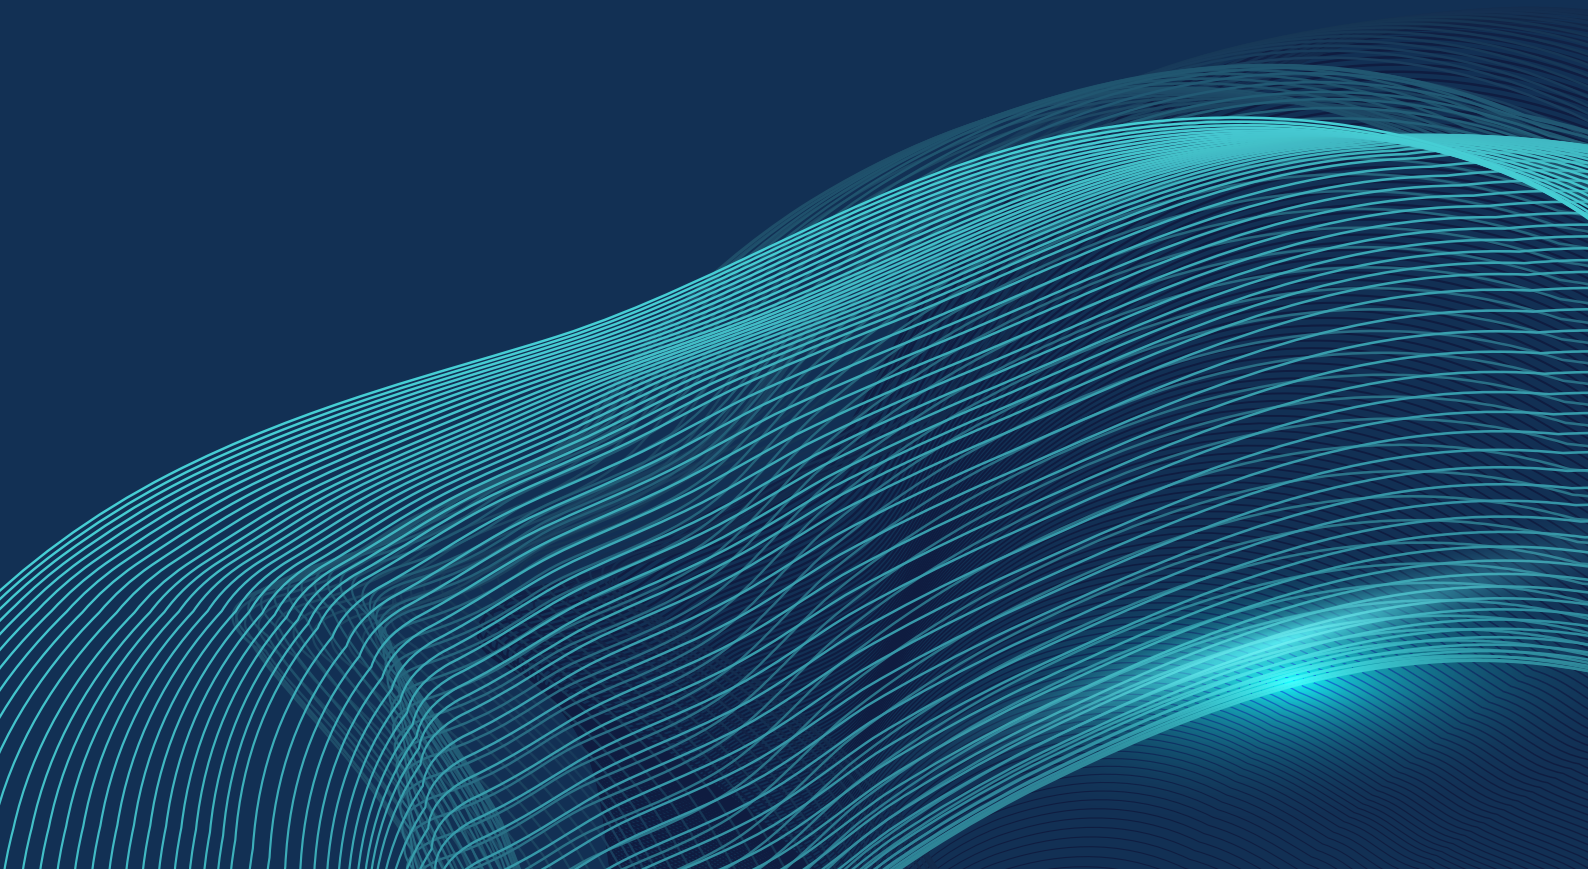

APPENDIX I: OTHER LEGAL ISSUES

In this Appendix, we outline further legal issues that warrant attention in considering the design and implementation of the Windfall Clause.

A. Bindingness

Signatories' obligations under the Clause would be binding. First, as discussed above, the Clause could impose obligations on distributors, such as publicizing the philanthropic efforts of the signatories and setting up philanthropic programs.[295] In such a case, the Clause would be a regular binding contract.[296]

Even if the Clause is purely donative (i.e., unsupported by return consideration by the distributors), it would be binding under the doctrine of promissory estoppel. The Restatement (Second) of Contracts § 90 contains the doctrine of promissory estoppel:

> (1) A promise which the promisor should reasonably expect to induce action or forbearance on the part of the promisee or a third person and which does induce such action or forbearance is binding if injustice can be avoided only by enforcement of the promise. The remedy granted for breach may be limited as justice requires.
>
> (2) A charitable subscription [i.e., promise to donate] . . . is binding under Subsection (1) without proof that the promise induced action or forbearance.[297]

In comments thereto, the Restatement explains:

> One of the functions of the doctrine of consideration is to deny enforcement to a promise to make a gift. Such a promise is ordinarily enforced by virtue of the promisee's reliance only if his conduct is foreseeable and reasonable and involves a definite and substantial change of position which would not have occurred if the promise had not been made. In some cases, however, other policies reinforce the promisee's claim. Thus the promisor might be unjustly enriched if he could reclaim the subject of the promised gift after the promisee has improved it.
>
> Subsection (2) identifies two other classes of cases in which the promisee's claim is similarly reinforced. American courts have traditionally favored charitable subscriptions and marriage settlements, and have found consideration in many cases where the element of exchange was doubtful or nonexistent. Where recovery is rested on reliance in such cases, a probability of reliance is enough, and no effort is made to sort out mixed motives or to consider whether partial enforcement would be appropriate.[298]

The following illustrate application of this rule:

> 16. A orally promises to give her son B a tract of land to live on. As A intended, B gives up a homestead elsewhere, takes possession of the land, lives there for a year and makes substantial improvements. A's promise is binding.[299]
>
> 17. A orally promises to pay B, a university, $100,000 in five annual installments for the purposes of its fund-raising campaign then in progress. The promise is confirmed in writing by A's agent, and two annual installments are paid before A dies. The continuance of the fund-raising campaign by B is sufficient reliance to make the promise binding on A and his estate.[300]

Since Signatories to the Windfall Clause "should reasonably expect [the Clause] to induce action or forbearance on the part of the promisee or a third person" (such as a distributor), the Clause would be binding under this provision of the Restatement.[301]

Some major jurisdictions—such as California[302] and New York[303]—deviate slightly from the Restatement approach to charitable subscriptions by still requiring actual "action or forbearance" by the promisee or a third party. The Clause would be binding even under these regimes, however, since actions like setting up distributors to eventually receive windfall or inducing other potential firms to sign onto the Clause would constitute adequate consideration even under these more restrictive regimes.

B. Injunctive Relief

One potential failure mode for the Windfall Clause would be a court enjoining[304] the donations. The most obvious way this could happen would be that a firm obtains windfall profits, but a disgruntled shareholder files a derivative action to try to enjoin the obliged donations as a violation of the corporation's fiduciary duties. We explain here why we think such a challenge would be unlikely to succeed under Delaware law.

"To succeed in a request for a permanent injunction, a party must show [1] actual success on the merits; [2] that it would suffer irreparable harm if the injunction is not granted; and [3] that the balance of the equities favors it."[305] Injunction is an equitable remedy,[306] meaning that the Chancery Court retains wide discretion over whether to grant injunctive relief.[307] As the third prong of the injunction test suggests, courts in equity weigh general fairness when deciding whether to grant an injunction.

B.1. Explication of Injunction Test

As we argued above,[308] we believe that a shareholder challenge to the Windfall Clause would fail on the merits. We therefore focus on the second (irreparable harm) and third (balance of equities) prongs here, to show why such a challenge would not warrant injunctive relief *even if* it was successful on the merits.

B.1.i. Irreparable Harm

A court will issue an injunction only if the plaintiff can demonstrate that she will suffer "irreparable harm" unless the court issues an injunction. "Irreparable harm" is a legal term of art, defined by the Delaware courts. Irreparable harm exists only when an award of monetary damages could not fairly and reasonably compensate the plaintiff,[309] including when such an award would rely upon speculation as to the extent of the harm suffered.[310] For example, loss of a unique strategic business opportunity can constitute irreparable harm.[311] However, neither the possibility of great injury alone[312] nor the "mere apprehension" of uncertain future damages[313] constitute irreparable harm. To receive an injunction, the plaintiff must show a "reasonable apprehension" of a future wrong.[314]

Importantly, there is a "well-recognized legal principle that equity will enjoin a threatened breach of fiduciary duty."[315] Delaware courts have not explicated this basis for injunctive relief much,[316] but it appears to be a *per se* rule: any threatened breach of fiduciary duty will support injunctive relief, even if the consequences of the breach would not otherwise constitute irreparable harm.[317]

B.1.ii. Balance of Equities

A court will issue an injunction only when the harms from an injunction do not outweigh the harms to the plaintiff that an injunction will prevent.[318] In performing this balancing test, a court may consider harms to the public and innocent third parties.[319] A court in equity will also refuse to issue an injunction when it would affect the rights of parties not before the court.[320]

B.2. Application to Windfall Clause

A Delaware court will not enjoin a contract containing a Windfall Clause. Although executing of such a contract in violation of fiduciary obligations is arguably a *per se* irreparable harm, a court of equity will not abrogate the contractual rights of innocent third parties.

B.2.i. Irreparable Harm

The harm of losing windfall profits is not itself irreparable. Since the loss is purely monetary and easily calculable, it is classically inappropriate for injunctive relief.

The more likely basis for enjoining a Windfall Clause is that windfall payments would constitute a "threatened breach of fiduciary duty."[321] Analytically, the probability of injunctive relief on this basis depends on what, exactly, the "threatened breach" is. Two possibilities exist: agreeing to a Windfall Clause and performing a Windfall Clause to which the corporation already agreed.

A court will not enjoin mere agreement to a Windfall Clause. This is for two reasons. First, such an agreement would not, on its own, constitute a breach of fiduciary duty unless windfall profits were imminent, for reasons explained above.[322]

Relatedly, given the current improbability of any specific firm generating windfall profits, mere agreement to a Windfall Clause would support, at most, "mere apprehension" of a future harm. Thus, no "reasonable apprehension" of irreparable harm would exist at the time of the agreement itself.

Performing a Windfall Clause in violation of fiduciary obligations, if possible,[*] would constitute an irreparable harm under the *per se* rule of *duPont*. Thus, whether a court would enjoin performance of a Windfall Clause turns on whether the balance of equities favors an injunction.

B.2.ii. Balance of Equities

A court will not enjoin performance of a binding Windfall Clause, even if such performance violates fiduciary duties. This is because a court in equity considers how an injunction will affect innocent third parties. Enjoining performance of a Windfall Clause would severely harm the innocent donee and public, so a court of equity is very unlikely to do so.

More specifically, courts in equity will virtually never abrogate contractual rights of an innocent third party. "Equity follows the law . . . ."[323] This means that "courts of equity will recognize and give effect to all legal rules in their proper sphere . . . and that, beyond that, the policies reflected in rules of law will be extended to equitable estates by analogy where appropriate."[324] Thus, "[e]quity can not . . . abrogate a legal right contracted for by a party."[325] This principle is well-recognized outside of Delaware too.[326]

Consider, for example, In re *El Paso Corp. Shareholder Litigation*[327] That case addressed a proposed merger between El Paso and Kinder Morgan:[328]

> [T]he parties entered into [a] "Merger Agreement." The Merger Agreement contain[ed] a commitment from El Paso to assist Kinder Morgan in the sale of [El Paso's fossil fuel extraction and production] business, which Kinder Morgan hoped could be accomplished before the closing of the Merger. The Merger Agreement also contain[ed] a "no-shop" provision preventing El Paso from affirmatively soliciting higher bids, but g[ave] the Board a fiduciary out in the event it receive[d] a "Superior Proposal" from a third party for more than 50% of El Paso's equity securities or consolidated assets.[329]

Plaintiffs moved for a mandatory injunction that contravened the terms of the Merger Agreement.[330] The court found probable success on the merits[331] and likely irreparable harm.[332] How-

---

[*] Note the peculiar result that follows if this is indeed possible: agreeing to a Windfall Clause should be permissible, see supra § III, but following through on that commitment might not be. This seems contradictory: either a fiduciary cannot contract to do something that might violate their fiduciary duties (no matter how unlikely such a contingency is), or a transaction that would otherwise violate fiduciary obligations is permissible if it is the result of a business judgment that was reasonable in expectation. Given that corporate fiduciaries must often make decisions on the basis of expected value, and that some such decisions will inevitably turn out to have been poor deals for the corporation ex post, the latter interpretation should be correct. Thus, we expect that, where agreement to a Windfall Clause is protected by the Business Judgment Rule, Delaware law does not require fiduciaries to breach the Windfall agreement when performance thereof would commit the corporation to donate what would otherwise be "unreasonable" amounts of money.

ever, the court refused an injunction on the balancing prong of the equity test.[333] This was because of the form of the requested injunction:

> [P]laintiffs want an odd mixture of mandatory injunctive relief whereby I affirmatively permit El Paso to shop itself in parts or in whole during the period between now and June 30, 2012, in contravention of the no-shop provision of the Merger Agreement, and allow El Paso to terminate the Merger Agreement on grounds not permitted by the Merger Agreement and without paying the termination fee set forth in the Merger Agreement, but then to lift the injunction and then force Kinder Morgan to consummate the Merger "if no superior transactions emerge."[334]

In finding such a proposal inequitable, the court held that "that sort of injunction would pose serious inequity to Kinder Morgan, which did not agree to be bound by such a bargain."[335]

Courts do sometimes consider enjoining fundamental corporate transactions that have already begun, but almost always before closing, when parties' legal rights have fully vested.

A notable exception to this is the Chancery Court's decision in the famous *Revlon* case.[336] There, Revlon's board tried to defend against a hostile takeover by granting to a friendly bidder a "lock-up" option to purchase valuable assets and a breakup fee in case the friendly deal failed.[337] The court preliminarily enjoined escrowing and transfer of the lock-up assets and cancellation fee,[338] despite the friendly bidder's contractual rights thereto.[339] The court held that "[i]n terms of relative hardship to the parties the need for both bidders to compete in the marketplace far outweighs the limiting of [the friendly bidder]'s contractual rights."[340] However, other language from the case suggests that the court did not view the friendly bidder as a truly innocent third party: "[the friendly bidder] and Revlon considered the [asset lock-up and breakup fee] as combined security to [impermissibly] secure the exclusion of [the hostile bidder] from further participation."[341]

A plaintiff unsuccessfully asked the Chancery Court to rescind[342] a consummated merger in *Arnold v. Society for Savings Bancorp, Inc.*[343] Earlier in the same case, the Supreme Court held that directors of the target corporation impermissibly made misleading statements in the merger proxy statement.[344] However, the directors bore no personal liability due to a §102(b)(7) waiver.[345] On remand, plaintiffs sought to rescind the merger. In an unpublished opinion, the Chancery Court roundly rejected this as inequitable:

> The equities in this case do not justify the extraordinary remedy of rescission. Rescission of the merger would unduly harm the acquiror, BoB, who has relied on the finality of the merger. Plaintiff has yet to demonstrate that BoB engaged in any unlawful or improper activities in carrying out the merger. BoB did not ignore Plaintiff's claims and pursue an improper merger. Before BoB completed the merger, this Court had ruled that Plaintiff's claims lacked a reasonable probability of success on the merits. Until Plaintiff demonstrates otherwise, I consider BoB an innocent party who acquired Bancorp in good faith in an arms-length transaction. A court of equity should protect the rights of an innocent acquiror, just as commercial law protects the rights of a good faith purchaser. On the present record, I believe it would be manifestly unfair to BoB for the Court to rescind the merger. If Plaintiff can prove his aiding and abetting claim against BoB, then the balance of the equities may change. Nevertheless, practical obstacles would still prevent rescission of the merger.[346]

Note, however, that the court cites its prior rejection of plaintiff's injunction request[347] as evidence of its good faith. This suggests that absent such a ruling, its claim to good faith would be weaker, though probably still sufficient.

Another Delaware case to explicitly address the issue of vested merger rights is *Jedwab v. MGM Grand Hotels, Inc.*[348] In that case, Kirk Kerkorian was a controlling shareholder of MGM

Grand Hotels.[349] At Kerkorian's direction, MGM "entered into an agreement with Bally Manufacturing Corporation . . . contemplating a merger between a Bally subsidiary and [MGM]."[350] The MGM board approved the proposed merger.[351]

A preferred stockholder asked the court to enjoin the transaction on the grounds that it "apportion[ed] the merger consideration [un]fairly among classes of the company's stock."[352] The court refused the injunction on the grounds that the plaintiff failed to satisfy the first prong of the injunction test.[353] In dicta on the balancing prong, however, the court wrote that "Bally's contract rights—*while not dispositive*—present an additional circumstance supporting the denial of the pending motion."[354] It paid similar deference to "the interests of the public common stockholder . . . to have the proposed merger effectuated without judicial interference."[355] This presumably refers to the payment such shareholders were to receive from the sale.[356] In an accompanying footnote, the court explained that "[h]istorically courts of equity have accorded great deference to the rights of bona fide purchasers from trustees who have no notice of a breach of trust."[357]

In light of all of the above evidence, it is very unlikely that a court of equity will enjoin performance to which a Windfall Clause distributor has a contractual right. Since the distributor would have such a right, the Delaware Chancery Court would not enjoin performance of the Clause.

C. Timing of a Challenge

For equitable claims, such as breach of fiduciary duty, the statute of limitations is not controlling.[358] Instead, courts in equity apply the doctrine of "laches," which "'operates to prevent the enforcement of a claim in equity if the plaintiff delayed unreasonably in asserting the claim, thereby causing the defendants to change their position to their detriment.'"[359] "This doctrine 'is rooted in the maxim that equity aids the vigilant, not those who slumber on their rights.'"[360]

Nevertheless, since "equity follows the law," "absent unusual circumstances, the analogous statute of limitations will be given great weight in deciding whether the plaintiff's claim is barred by laches."[361] "The analogous limitations period for a breach of fiduciary duty claim is three years."[362] The three-year limitations period also "applies [by analogy] to shareholder derivative actions which seek recovery of damages or other essentially legal relief."[363] The limitations period starts (i.e., the cause of action "accrues") "at the time of the wrongful act, even if the plaintiff is ignorant of the cause of action."[364]

Importantly for the purposes of a Windfall Clause, the continued performance of a challenged contract does *not* count as a continuing wrong. The case *Kahn v. Seaboard Corp.*[365] addressed an allegedly self-dealing contract entered into in 1986.[366] The contract chartered seven Seaboard vessels for ten years.[367] Plaintiffs filed suit in 1990 (i.e., more than three years after the alleged improper transaction)[368] under the theory that continued performance of the conflicted transaction constituted a continuing wrong, rendering the action timely.[369] The court rejected this argument:

> The wrong attempted to be alleged is the use of control over Seaboard to require it to enter into a contract that was detrimental to it and beneficial, indirectly, to the defendants. Any such wrong occurred at the time that enforceable legal rights against Seaboard were created. Suit could have been brought immediately thereafter to rescind the contract and for nominal damages which are traditionally available in contract actions. Complete and adequate relief, if justified, could be shaped immediately or at any point thereafter.
>
> This type of case is unlike a continuing tort where the defendant continues, without right, an action injurious to plaintiff. Where a continuing wrong acts as an answer to the defense of limitations it is typically the case that plaintiff can prove her claim by reference only to actions within the limitations period. Thus, for example, if plaintiff is complaining about a nuisance (a noise for ex-

> ample) emitted by a neighboring plant for, say five years, plaintiff can prove her claim by proving the elements of the claim which occur daily or weekly or whenever, within the limitation period. It is irrelevant for limitations purposes that these daily or weekly invasions have been going on for years. Here, however, the "continuing wrong" is performance of a contract. It is implicitly admitted that payments were made by Seaboard as provided in the contract. There is no claim that payments in excess of those contemplated by the time charter have been made. *So long as the time charter is not rescinded, the payments it calls for are legal obligations, not wrongs*. Thus, unlike a continuing wrong the only liability matter to be litigated involves defendants' 1986 actions in authorizing the creation of these contract rights and liabilities.[370]

Thus, potential derivative suit plaintiffs would have three years *from the time the Clause is signed* to challenge it. This means that late-stage challenges to the Clause are very unlikely to succeed under now-existing law.

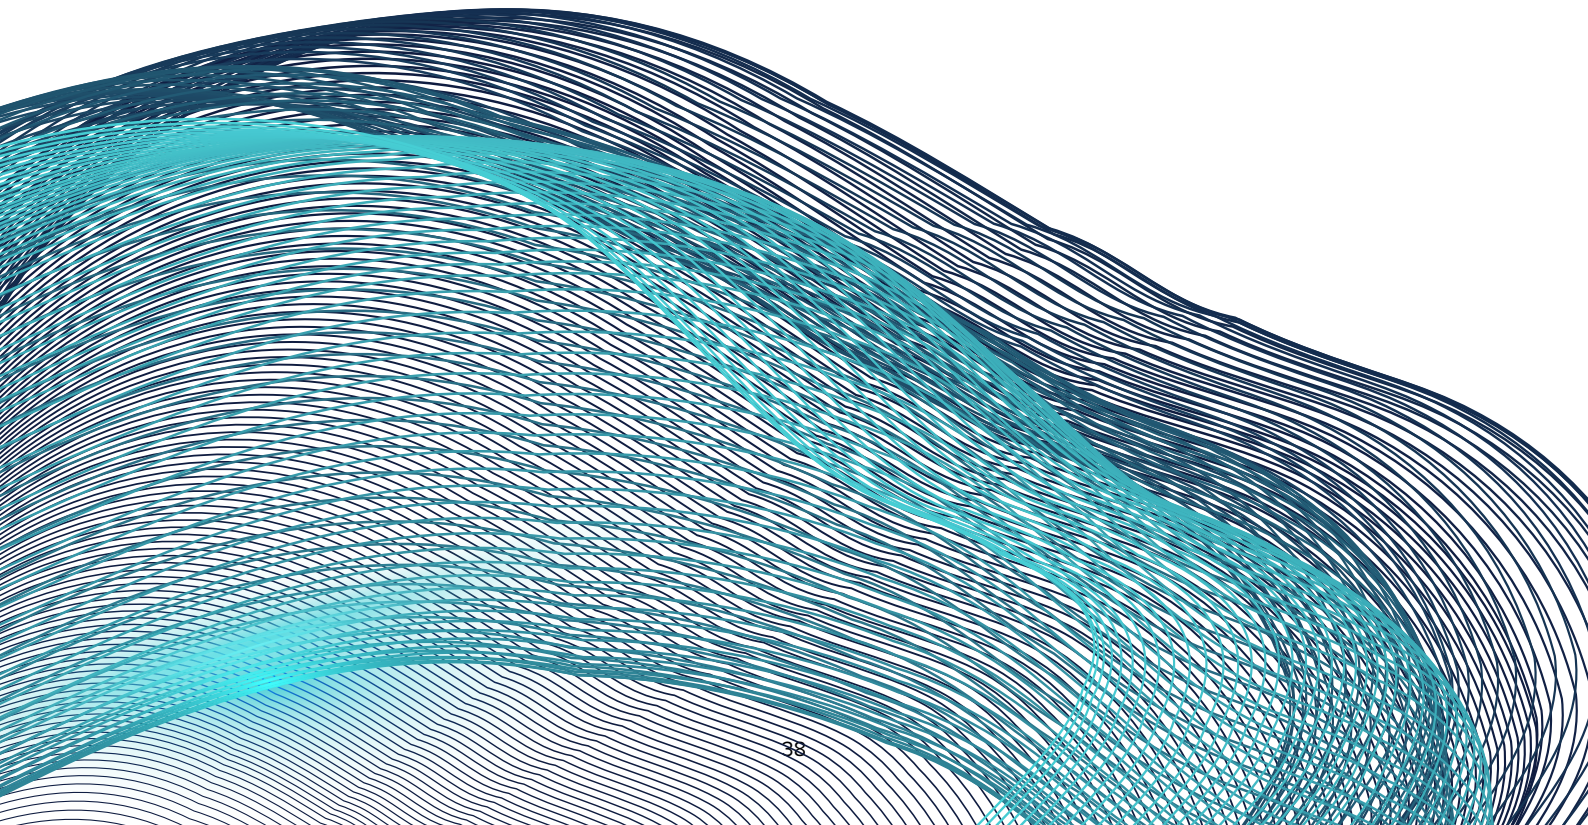

APPENDIX II: THREE SKETCHES OF POSSIBLE WINDFALL DISTRIBUTION SYSTEMS

This Appendix briefly lists three different forms a windfall distribution system could take in light of the desirable characteristics as outlined in Section II. These three proposals for distributors are outlined along with explanation of each proposal's various features, especially as they relate to the public policy and firm motivations for the Clause.[371]

We stress that each of these is, at this stage, a mere sketch of a proposal. Once firms agree to the Windfall Clause, distributors might include any of these, each of these, some hybrid of each, or something else entirely.* We look forward to further public engagement on how to optimally distribute proceeds from the Windfall Clause. We specifically examine four features of each plan:

1. Its structure;
2. Its functioning: what distributors' goals will be and how they will accomplish them;
3. Its accountability and legitimacy: how the plan's distributors will be kept safe from improper interference with its mission and have broad approval from ultimate beneficiaries;† and
4. How the plan will advance firms' interests.

A. Global Wellbeing Windfall Fund

Under this plan (the Global Wellbeing Windfall Fund, or GWWF), the distributor is a multimember body with the goal of finding and funding effective welfare-maximizing projects.

A.1. Structure

The GWWF is a multimember funding body.‡ GWWF members—which we will call "managers"—collaboratively identify and give grants to promising projects to improve global welfare:[372]

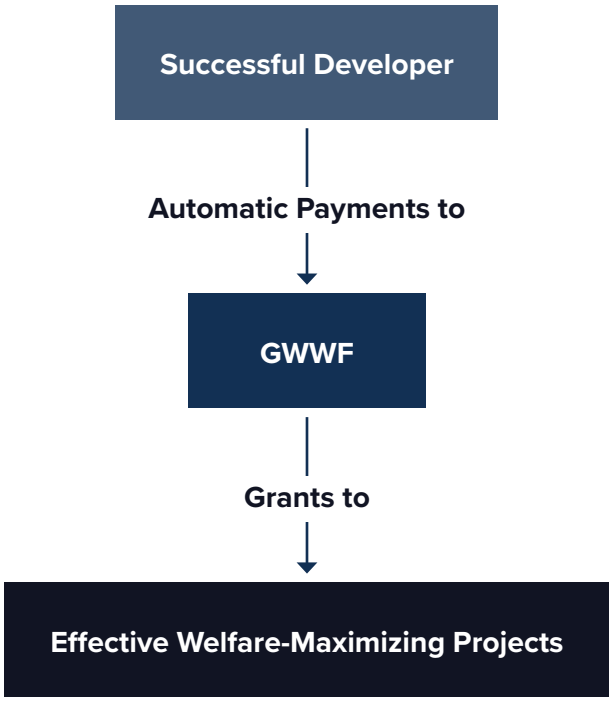

Figure 1: Global Wellbeing Windfall Fund Payment Flow

* One might draw further inspiration from the alternatives we examine *infra* § V(D).
† Given the large amounts of money potentially at stake, designing accountable distribution systems is extremely important.
‡ The reasons for having multiple members are discussed in the next subsection.

A.2. Functioning

The GWWF would have the general goal of making grants based on empirical evidence and with the aim of maximizing global welfare. This goal could be constrained or broadened as politically feasible,* but at minimum should also explicitly entail commitments to the observation of human rights in pursuit of its mission.

GWWF managers would be experts in welfare-relevant fields such as public health, public policy, economics, and psychology.† They could function individually or in interdisciplinary teams.‡ The experts would only release funding to highly cost-effective programs.[373]

A.3. Accountability and Legitimacy

The GWWF would control a lot of money. Our aim is to enable them to use that money to maximize welfare, free from inappropriate outside interference that would direct the funds towards parochial or corrupt uses. The GWWF could implement a number of features to accomplish this. Security measures the GWWF could use include:

- Requiring that grants be approved by multiple[374] or a majority of GWWF managers;
- Enabling other GWWF managers to veto suspicious grants (if a majority is not required);
- Recruiting managers from and domiciling managers in different countries, to prevent a single country from exercising undue influence over grantmaking decisions;§
- Preventing managers from approving grants that primarily benefit their own nation or their host nation;
- Anonymizing voting on whether to fund certain projects;[375]
- Limiting the amount of funding that a single nation can receive, or which a single grant can convey;¶ and
- Establishing procedures for removing and replacing corrupted GWWF managers.

Apart from these accountability features, the GWWF could make use of a novel accountability structure proposed in a law review article by Geoffrey A. Manne: "contract plaintiffs."[376] The problem that Manne seeks to address is limited legal oversight of charities: usually, only state attorneys general have standing to sue nonprofit managers for misuse of funds, but they have poor incentives to do so.[377] Thus, Manne proposes a contractual solution wherein nonprofits are required by charter to give a third party (the "contract plaintiffs") the right to monitor and sue their managers for failure to advance the purposes of the nonprofit.[378] Including such a requirement in the GWWF's charter(s) would ensure their strict adherence to the charitable goals of the GWWF.[379]

A.4. Firm Interests

A successful firm would benefit from the GWWF by having its brand directly associated with extremely effective global wellbeing projects. To help the firm promote this connection, the GWWF could be required to release information (e.g., videos, photographs, reports) about

---

* For example, it could or could not include animal welfare and the welfare of future human generations. Grounding GWWF-funded projects in existing consensus commitments (for example, the United Nations' Sustainable Development Goals, *cf. infra* app. II(C)) could garner greater consensus for the GWWF and its activities.
National diversity could also be important. For example, the most effective philanthropic opportunities are concentrated in the world's poorest nations. *See Your Dollar Goes Further Overseas*, *supra* note 256. However, residents of other nations might object to receiving no aid, even if such a scheme is welfare-maximizing. To ensure the political and geopolitical feasibility of the scheme, therefore, the Clause could specify that projects in every country will receive a set amount of money per year, possibly in proportion to their population.

† Where appropriate (e.g., for large, controversial, or speculative grants), the managers could be permitted or even required to hire domain-area experts for further consultation.

‡ If the functional unit of the GWWF was teams, rather than individuals, then the GWWF security measures would apply to teams.

§ For maximum security, these should ideally be countries that have a strong rule of law; that respect contract rights; that have adequate self-defense forces; and that are unlikely to coordinate with many other hosts. Of course, diversity of membership could provide additional legitimacy and decision-making benefits.

¶ Much as ATMs limit depositors' daily withdrawals. The idea is that if a country successfully corrupts a single grant or GWWF manager, the damage will be limited. Since in some (indeed, perhaps many) cases grants will be urgently needed, this limitation could perhaps be waived pending approval by a supermajority of other managers.

its projects that would help the firm publicize its charitable activities. Grants could also be conditional on publicizing the firm's connection to the project, to directly garner goodwill from project beneficiaries.

B. Windfall Trust

This distributor proposal would democratically distribute the windfall through a "Windfall Trust."[*,380]

B.1. Structure

This plan would distribute windfall money to a trust[†] over which all living people would be equal beneficiaries. Once the Windfall Clause is triggered, the successful firm will distribute windfall payments to Windfall Trust. All living individuals (or perhaps all adults) would be beneficiaries of the Trust:

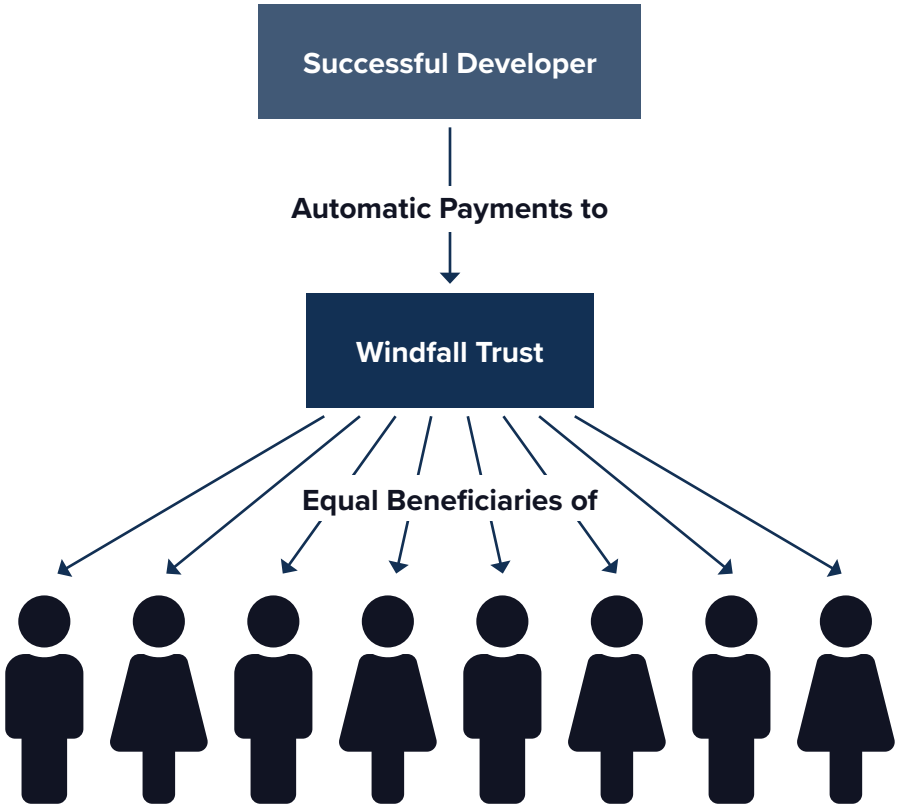

Figure 2: Windfall Trust Payment Flow

B.2. Functioning

As its name suggests, Windfall Trust will function much as a regular trust, operating for the financial benefit of its beneficiaries—i.e., everyone. As such, it will have a wide variety of tools at its disposal to improve its beneficiaries welfare, including distributing a windfall as required by the terms of the Trust and investing for the benefit of future beneficiaries.[381] The result, therefore, would be a trust working to enrich all humans through means consistent with any public benefit restrictions in its charter.

---

* We do not have a specific single jurisdiction in mind when describing Windfall Trust. Some of what follows is therefore necessarily stylized, though on occasion we note how aspects of this proposal interact with particular relevant legal regimes. However, we are optimistic that what follows is generally feasible as a matter of law.

† We are intentionally leaving for a later date more complicated discussions about implementing such a trust. Open issues would include whether such a trust could qualify as charitable and the issue of creating and administering a trust for such a large number of people.

B.3. Accountability and Legitimacy

As with the GWWF, the Windfall Trust should be domiciled in a country with low risk of nationalization or expropriation. Furthermore, again like the GWWF, the Clause could hedge against nationalization and expropriation by designating multiple Windfall Trusts in different nations, with a mechanism for cutting windfall flows off from corrupted Trusts and creating new ones if necessary.

Beneficiaries would also be able to hold trustees of the Trust accountable through traditional means of fiduciary accountability.[382] Thus, they would hopefully feel that the Trust's management is legitimately serving their interests.

B.4. Firm Interests

All Windfall Trust beneficiaries—i.e., everyone—would have a reason to have goodwill towards a successful firm (or at least to not feel animus towards it). Additionally, per its terms (and perhaps subject to the assent of a majority of living beneficiaries), the Windfall Trust could potentially offer a lucrative buyout of the firm or its assets, thus allowing the firm to recoup some of their commitments under the Clause. This makes the Windfall Trust potentially more appealing than the other distribution proposals here, since it would allow some of the pledged windfall to flow back to the firm's investors.[*]

C. AI for Good Fund

The AI for Good Fund (AI4G) is inspired by existing proposals (also under the name "AI for Good") to use AI to advance humanitarian causes like the United Nations' Sustainable Development Goals (SDGs).[383,†] The UN Secretary General called for such an approach in his 2018 Strategy on New Technologies.[384]

[*] Of course, ex post those investors would have been better off if the firm had never signed the Clause. But, all else equal, they are arguably better-off under the Windfall Trust proposal than under, e.g., the GWWF, since the Windfall Trust proposal could redirect some of the pledged windfall back to the firm and therefore its shareholders through a buyout or asset acquisition, whereas the GWWF could not.

[†] As the SDGs are accomplished (either due to the efforts of the AI4G or because we accomplish them before the Windfall Clause is triggered), the set of challenges addressed by the AI4G could expand to include advancement of other global consensuses, such as human rights treaties, nonproliferation treaties, and environmental treaties.

C.1. Structure

AI4G would consist of several "pods," each of which is tasked with using AI to advance the SDGs.[385] An Oversight Body would have the authority to allocate funding to create, monitor, combine, split, and retire pods based on performance and changing needs. To avoid possible redundancy,* pods would not be targeted at a specific SDG but would rather focus on a *domain* of AI applications relevant to the SDGs. For example:

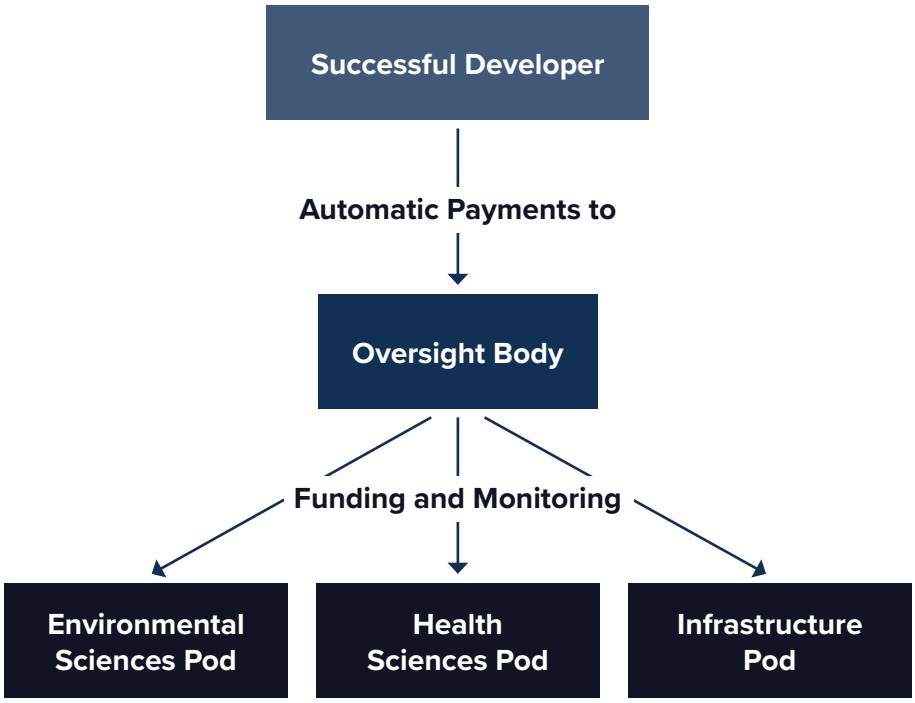

Figure 3: AI for Good Fund Payment Flow

Each pod would be authorized to use its funding to pursue its mission by, e.g., hiring domain experts, hiring AI researchers, creating training data (e.g., hiring people to demonstrate certain tasks in robotics), paying for access to compute, paying for access to existing AI systems, paying for overhead, building national and local capacity to implement identified solutions, and paying for expert advisors.†

C.2. Functioning

Pods would have broad latitude in how they choose to pursue their mandate, provided that they consult appropriate experts, use or address‡ AI in their work, and tether their work to the advancement of the SDGs or similar consensus goals (as determined by the Oversight Body).

* For example, the first three SDGs ("[e]nd poverty in all its forms everywhere," "[e]nd hunger, achieve food security and improved nutrition, and promote sustainable agriculture," and "[e]nsure healthy lives and promote well-being for all at all ages," *see id.*) are related problems and presumably have related solutions. It would therefore be counterproductively duplicative to have a distinct pod for each.

† Some of this (e.g., existing AI systems) could come from the firm. However, the public might react negatively to pods paying the firm for help in accomplishing the SDGs. The Clause could therefore allow the firm to deduct in-kind donations to AI4G from its Clause obligations. Although this would be economically equivalent to pods paying for services from the firm, this altered scheme might be received better by the public. To avoid corruption of the altruistic ends of AI4G, however, in-kind donations from the firm should be allowed only if the pod requests them—i.e., if they are the most effective means for advancing the pod's mission.

‡ For example, one pod might be dedicated to alleviating automation-induced unemployment. Such a pod could—but need not—use AI in its work. For example, it could instead pay people to perform socially valuable tasks for which no natural market exists (and which automation cannot easily replace), such as becoming more educated or engaging with their community. (Similar schemes have been suggested in, *e.g.*, FORD, *supra* note 6, at 172–93; POSNER & WEYL, *supra* note 204, at 205–49; RIFKIN, *supra* note 6, at 236–74.)

C.3. Accountability and Legitimacy

The primary source of security for AI4G would be that it is pursuing consensus-driven work. As with the other distributor proposals, AI4G pods could be distributed among diverse, politically stable states to minimize the possibility of external undesired interference.

Furthermore, pods' mandates could instruct that the outputs, if applied responsibly, are stabilizing—that is, pods' work will hopefully improve the political environment by incrementally creating a world that is, e.g., richer, more peaceful, less discriminatory, more highly educated, and more equitable.*

If, instead of directing windfall directly to pods, the Clause provides for an Oversight Body to fund and actively manage the pods (as we suggest above),† then security of the Oversight Body will be crucial to the security of this proposal. We imagine that the Oversight Body could be made most secure by, first, providing that multiple stakeholders govern it. For example, its board could be divided between national governments, AI firms,‡ development experts, and democratically elected representatives. The Oversight Body would also be secure because it could not release money other than to pods, which would spend money free from any improper influence by the Oversight Body.§

C.4. Firm Interests

As with the other distributor proposals, the fact that AI4G will bring universal (or at least very broad) benefit will earn the firm goodwill. Furthermore, since the projects will use AI, the firm will be able to credibly demonstrate that AI is a universal boon. The fact that AI4G would need AI expertise both in the Oversight Body and implementation creates an opportunity for the lab (and even other signatories) to be involved in AI4G and therefore closely associate themselves with it.

* Indeed, one pod might be tasked with promoting international stability.

† We anticipate that the active approach is preferable to a fixed approach due to both the potential for progress on SDGs necessitating a shift in priorities and due to the emergence of new and unanticipated global challenges with which AI4G could help.

‡ The reason to include AI firms would be that they would have expertise to judge potential altruistic applications of AI.

§ A contract plaintiff, *see generally supra* app. II(A)(3), could have standing and incentive to challenge any improper Oversight Body actions, thus maintaining the pods' independence and the general effectiveness of AI4G.

ENDNOTES

1 *See generally* ROBERT N. ANTHONY ET AL., ACCOUNTING: TEXT & CASES 66 (11th ed. 2004). We anticipate the need for future discussion on signatories' accounting methods, especially to allow firms to prudently reinvest profits without the Clause punishing them for doing so. For the purposes of this introductory report, however, we think this definition is sufficient.

2 *See, e.g.*, Manuel Trajtenberg, *Artificial Intelligence as the Next GPT: A Political-Economy Perspective*, *in* THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA 175 (Ajay K. Agrawal et al. eds, 2019).

3 *See infra* § I(B)(1)(ii).

4 *See infra* §§ I(B)(1)(i)–(ii).

5 *See* Katja Grace et al., When Will AI Exceed Human Performance?: Evidence from AI Experts (2018), https://perma.cc/2K2D-LE3A ("Researchers believe there is a 50% chance of AI outperforming humans in all tasks in 45 years and of automating all human jobs in 120 years . . . .").

6 This has centered most visibly on the issue of AI applications beginning to replace human labor in many sectors. *See, e.g.*, JONATHAN P. ALLEN, TECHNOLOGY AND INEQUALITY (2017); ERIK BRYNJOLFSSON & ANDREW MCAFEE, THE SECOND MACHINE AGE (2014); JOHN DANAHER, AUTOMATION AND UTOPIA (2019); ALLAN DAFOE, AI GOVERNANCE: A RESEARCH AGENDA 32 (Aug. 27, 2018), https://perma.cc/4KG2-MTL3; MARTIN FORD, RISE OF THE ROBOTS (2015); MARTIN FORD, THE LIGHTS IN THE TUNNEL: AUTOMATION, ACCELERATING TECHNOLOGY AND THE ECONOMY OF THE FUTURE (2009); MICHAEL C. HOROWITZ ET AL., ARTIFICIAL INTELLIGENCE AND NATIONAL SECURITY 14–19 (July 2018), https://perma.cc/4GCY-LUVL; JERRY KAPLAN, HUMANS NEED NOT APPLY (2015); JAMES MANYIKA ET AL., JOBS LOST, JOBS GAINED: WORKFORCE TRANSITIONS IN A TIME OF AUTOMATION (2017), https://perma.cc/XRU5-2A77; JEREMY RIFKIN, THE END OF WORK (1995); AARON SMITH & JANNA ANDERSON, AI, ROBOTICS, AND THE FUTURE OF JOBS (Aug. 6, 2014), https://perma.cc/S8PN-UKJY; THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA, *supra* note 2; DARRELL M. WEST, THE FUTURE OF WORK (2018); Carl Benedikt Frey & Michael A. Osborne, *The Future of Employment: How Susceptible Are Jobs to Computerisation?*, 114 TECHNOLOGICAL FORECASTING & SOC. CHANGE 254 (2017); Jeffrey D. Sachs & Laurence J. Kotlikoff, *Smart Machines and Long-Term Misery* (Nat'l Bureau Econ. Research, Working Paper No. 18629, 2012), https://perma.cc/MS34-B3EK; *Asilomar AI Principles*, FUTURE OF LIFE INST., https://perma.cc/AU9J-X239 (archived Oct. 31, 2018); Sergey Brin, *2017 Founders' Letter*, ALPHABET: INVESTOR RELATIONS (2017), https://perma.cc/59PF-WB4R; Catherine Clifford, *Bill Gates: 'A.I. Can Be Our Friend'*, CNBC (Feb. 4, 2018, 2:34 PM), https://perma.cc/L2U6-JL7M; Catherine Clifford, *Elon Musk: Robots Will Take Your Jobs, Government Will Have To Pay Your Wage*, CNBC (Nov. 4, 2016, 2:19 PM), https://perma.cc/QY7V-PSZD; David Rotman, *Technology and Inequality*, MIT TECH. REV. (Oct. 21, 2014), https://perma.cc/AV7L-63QE; Ian Sample, *Joseph Stiglitz On Artificial Intelligence: 'We're Going Towards A More Divided Society'*, GUARDIAN (Sept. 8, 2018), https://perma.cc/9PAC-TUL8; Richard Wike & Bruce Stokes, *In Advanced and Emerging Economies Alike, Worries About Job Automation*, PEW RES. CTR. (Sept. 13, 2018), https://perma.cc/YAX5-5XRK.

7 *See generally* Anna Jobin et al., Artificial Intelligence: The Global Landscape of Ethics Guidelines (2019) (preprint version), https://perma.cc/9AES-2X8D

8 *See* NICK BOSTROM, SUPERINTELLIGENCE 254 (2014).

9 Stylized examples of possible distribution mechanisms are offered *infra* app. II.

10 Additional legal aspects of the Clause are analyzed *infra* app. I.

11 *See* Nick Bostrom et al., *Public Policy and Superintelligent AI: A Vector Field Approach*, *in* ETHICS OF ARTIFICIAL INTELLIGENCE (S. M. Liao ed., forthcoming 2019) (manuscript version 4.3 at 5–6, 12–14), https://perma.cc/SN54-HKEG.

12 *See generally* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 126–46.

13 *See* Bostrom et al., *supra* note 11, at 10.

14 *See, e.g.*, SMITH & ANDERSON, *supra* note 6, at 5, 44–48; Frey & Osborne, *supra* note 6.

15 *See* BOSTROM, *supra* note 8, at 59–61; *cf.* GREG ALLEN & TANIEL CHAN, ARTIFICIAL INTELLIGENCE & NATIONAL SECURITY 17 (2017), https://perma.cc/HL2J-6Y4K ("Humans do not yet know how to replicate all the technologies and capabilities of nature, but the fact that these capabilities exist in nature proves that they are indeed possible. . . . Humans do not know what the ultimate technological performance limit for autonomous robotics is, but the ultimate limit can be no lower than the very high level of performance that nature has proven possible . . ."); FORD, *supra* note 6, at 83–128 (describing automation of white-collar jobs); Grace et al., *supra* note 5, at 2 ("Taking the mean over each individual, the aggregate forecast gave a 50% chance of [the arrival of machines that can accomplish, without aid, every task better and more cheaply than humans] occurring within 45 years and a 10% chance of it occurring within 9 years.").

16 *Cf.* Daron Acemoglu & Pascual Restrepo, *Modeling Automation* (Massachusetts Inst. Tech. Dep't Econ. Working Paper No. 18-02, 2018), http://ssrn.com/abstract=3123798 ("[I]n a task-based framework automation, conceptualized as the expansion of the set of tasks that can be produced by machines, has very different effects. It always reduces the labor share and it reduces labor demand and the equilibrium wage unless the productivity gains from automation are sufficiently large.").

17 *See* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 178; KAPLAN, *supra* note 6, at 13, 131–58.

18 The relationship between automation and job availability is not straightforward, and one cannot assume that automation (or even elimination of some jobs) will inevitably lead to net job loss. *See infra* note 25 (collecting sources predicting that automation will not result in mass unemployability). In the past automation has not led to widespread unemployability. *See Robots*, IGM FORUM (February 25th, 2014, 1:55 PM), https://perma.cc/N4UD-HE7N?type=image. This could continue to be the case due to, for example, AI-enabled further specialization of white-collar labor, continued demand for specifically human-produced goods and services, or high demand for human inputs into automation research and implementation. *See* KAI-FU LEE, AI SUPERPOWERS 155–57 (2018); *cf.* Daron Acemoglu & Pascual Restrepo, *Artificial Intelligence, Automation, and Work* (Nat'l Bureau Econ. Research, Working Paper No. 24196, 2018). "However, the historical tendency for employment and wages to increase as technological progress occurs is an empirical, historical phenomenon; it is not a law of nature or of economics." Stephen J. DeCanio, *Robots and Humans—Complements or Substitutes*, 49 J. MACROECON. 280, 281 (2016), https://perma.cc/HE2G-98TQ; *accord* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 128. Thus, the net effect of automation on human employability depends crucially on the elasticity of substitution between human and robotic labor. *See* DeCanio, *supra.* Under plausible—but by no means inevitable or uncontroversial—assumptions about this parameter, automation will depress median wages, *see id.* at 289, and possibly lead to persistent mass unemployment, *see, e.g.*, KAREN HARRIS ET AL., LABOR 2030: THE COLLISION OF DEMOGRAPHICS, AUTOMATION AND INEQUALITY 25 (2018), https://perma.cc/LPU8-DRRA ("By 2030, employers will need 20% to 25% fewer workers, equivalent to 30 million to 40 million jobs in the US . . . To put these numbers in context, during the Great Recession, US employment fell rapidly from its peak in January 2008 to its trough in February 2010 by nearly 9 million jobs, or 6.3% of total employment."); *Robots and Artificial Intelligence*, IGM FORUM (June 30, 2017, 9:28 AM), https://perma.cc/2FB9-HQSR?type=image (28% of surveyed economists agreed or strongly agreed that "[h]olding labor market institutions and job training fixed, rising use of robots and artificial intelligence is likely to increase substantially the number of workers in advanced countries who are unemployed for long periods.").

19 This framing taken from CGP Grey, *Humans Need Not Apply*, YOUTUBE (Aug. 13, 2014), https://www.youtube.com/watch?v=7Pq-S557XQU.

20 *See* FORD, *supra* note 6, at xv, 83–128. *Contra* Wike & Stokes, *supra* note 6 ("Most also see a role for individuals. This is particularly true in the U.S, Argentina and Brazil, where more than seven-in-ten say individuals themselves have a lot of responsibility for making sure they are prepared for the future economy. Again, Japan—where just 39% express this opinion—is an outlier.").

21 *See* FORD, *supra* note 6, at 32.

22 Dani Rodrik, *New Technologies, Global Value Chains, and the Developing Economies* 14 (Pathways for Prosperity Comm'n Background Paper Series, Paper No. 1, 2018), https://perma.cc/ZJ7Q-E4YZ; *see also* RIFKIN, *supra* note 6, at 203–07; Brahima Sangafowa Coulibaly, *Africa's Race Against the Machines*, BROOKINGS (June 16, 2017), https://perma.cc/K6NH-WGRA; Andrew Norton, *Automation Will End the Dream of Rapid Economic Growth for Poorer Countries*, GUARDIAN (Sept. 20, 2016), https://perma.cc/XX9V-XHTN; *cf.* RIFKIN, *supra* note 6, at 69–82 (discussing how earlier bouts of automation disproportionately impacted another group of disadvantaged laborers: African-Americans); Lukas Schlogl & Andy Sumner, *The Rise of the Robot Reserve Army: Automation and the Future of Economic Development, Work, and Wages in Developing Countries* 1 (Ctr. for Global Dev., Working Paper No. 487, July 2018), https://perma.cc/69JP-BJXA ("Automation is likely to affect developing countries in different ways to the way automation affects high-income countries. The poorer a country is, the more jobs it has that are in principle automatable because the kinds of jobs common in developing countries—such as routine agricultural work—are substantially more susceptible to automation than the service jobs—which require creative work or face-to-face interaction—that dominate high-income economies. This matters because employment generation is crucial to spreading the benefits of economic growth broadly and to reducing global poverty. We argue that the rise of a global 'robot reserve army' will have profound effects on labor markets and structural transformation in developing countries, but rather than causing mass unemployment, AI and robots are more likely to lead to stagnant wages and deindustrialisation. As agricultural and manufacturing jobs are automated, workers will continue to flood the service sector, driving down wages. This will itself hinder poverty reduction and likely put upward pressure on national inequality, weakening the poverty-reducing power of growth, and potentially placing the existing social contract under strain, or even possibly limiting the emergence of more inclusive social contracts.").

23 *See* FORD, *supra* note 6, at 117–28; *cf.* HARRIS ET AL., *supra* note 18, at 38 ("The primary macroeconomic consequence of higher inequality [due to automation] is to constrain growth by limiting the growth of effective demand vs. the growth in supply. Despite many technological innovations that will increase the capacity for goods and services, the inability of effective demand to keep pace may ultimately reduce growth.").

24 *See, e.g.*, SMITH & ANDERSON, *supra* note 6, at 5 ("Half of the [surveyed] experts (48%) envision a future in which robots and digital agents have displaced significant numbers of both blue- and white-collar workers—with many expressing concern that this will lead to vast increases in income inequality, masses of people who are effectively unemployable, and breakdowns in the social order. The other half of the experts who responded to this survey (52%) expect that technology will not displace more jobs than it creates by 2025."); *cf.* Grace et al., *supra* note 5 ("Researchers believe there is a 50% chance of AI outperforming humans in all tasks in 45 years and of automating all human jobs in 120 years . . .").

25 *See generally supra* note 6 (collecting citations). *But see, e.g.*, OECD, PUTTING FACES TO THE JOBS AT RISK OF AUTOMATION 3 (2018), https://perma.cc/2T5V-SX7U ("[E]ven if technology makes certain jobs redundant, it also creates new jobs and complements existing ones."); PwC, 2018 AI

PREDICTIONS 3 (2018), https://perma.cc/9F9U-WHY6 ("Th[e] unparalleled combination [of human and AI labor] will become the new normal in the workforce of the future."); PWC, WILL ROBOTS REALLY STEAL OUR JOBS? 8 (2018), https://perma.cc/Q63R-CFLE ("[T]he net long term job impact of automation w[ill] . . . likely to be neutral or even slightly positive."); Jason Furman, Chairman, Council of Econ. Advisers, Is This Time Different? The Opportunities and Challenges of Artificial Intelligence, Remarks at AI Now: The Social and Economic Implications of Artificial Intelligence Technologies in the Near Term (July 7, 2016) (expanded transcript available at https://perma.cc/S84L-LZZX); Bryan Caplan, *Technological Unemployment and Work*, LIB. ECON. & LIBERTY (Feb. 20, 2019), https://perma.cc/ME4C-W5WB; Paul Krugman, *Don't Blame Robots for Low Wages*, N.Y. TIMES (Mar. 14, 2019), https://perma.cc/85N7-N8DF; Noah Smith, *Basic Income's Backers Complicate Their Cause with Some Bad Arguments*, SALT LAKE TRIB. (June 22, 2019), https://perma.cc/RZ55-Y2WD; Noah Smith, *Don't Expect Robots to Take Everyone's Job*, BLOOMBERG: OPINION (Nov. 1, 2018, 6:00 AM), https://perma.cc/4FVU-ZC58; Noah Smith, *Robots Might Turn Out to Be Great Co-Workers*, BLOOMBERG: OPINION (June 18, 2019, 10:48 AM), https://perma.cc/6XTP-JM4L; *but cf.* John D. Stoll, *Humans Are Winning the Battle with Robots*, WALL STREET J. (Nov. 1, 2018) (describing difficulties in adopting robotic manufacturing).

26 *See, e.g.*, CELINDA LAKE, PUBLIC OPINION ON DISABILITY ISSUES 21 (2016), https://perma.cc/JJJ7-JRMT (81% of Americans support funding disability insurance); NORTHEASTERN U. & GALLUP, OPTIMISM & ANXIETY: VIEWS ON THE IMPACT OF ARTIFICIAL INTELLIGENCE AND HIGHER EDUCATION'S RESPONSE 21 (2018), https://perma.cc/5MS5-LYYB ("A clear majority of U.S. adults overall (61%) say employers should fund these programs for those who lose their jobs because of new technology. Of the options presented to Americans surveyed, this was the most popular source of funding for retraining, with the federal government second at 50%."); Khiara M. Bridges, *The Deserving Poor, The Undeserving Poor, and Class-Based Affirmative Action*, 66 EMORY L.J. 1049, 1070–79 (2017); Wike & Stokes, *supra* note 6 ("[R]oughly half or more believe employers have a lot of responsibility [for making sure the workforce has the education and skills necessary for success] in Argentina, Brazil, South Africa, Hungary, Italy, the U.S., Canada and Greece.").

27 *See supra* note 25.

28 *See* Bostrom et al., *supra* note 11, at 2, 5–6, 12–15.

29 *See* SMITH & ANDERSON, *supra* note 6, at 53–59; Bostrom et al., *supra* note 11, at 10–11; Betsey Stevenson, *AI, Income, Employment, and Meaning*, *in* THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA, *supra* note 2, ch. 7, at 2, https://perma.cc/ZB45-EG6G; Martin Ford, *Economic Growth Isn't Over, But It Doesn't Create Jobs Like It Used To*, HARV. BUS. REV. (Mar. 14, 2016), https://perma.cc/229L-5793 ("[T]echnology (and specifically artificial intelligence) is going to intertwine with any innovations that occur in the future, making them less labor-intensive. Unless we change our economic rules—perhaps with something like a guaranteed income—broad-based prosperity will remain elusive, even if those robust innovations do eventually show up."); Sachs & Kotlikoff, *supra* note 6; Wike & Stokes, *supra* note 6 ("Many in the nations surveyed also believe that the greater use of robots and computers will worsen inequality between the rich and the poor. More than eight-in-ten in Greece, Argentina, Japan and Brazil express this view, as do more than seven-in-ten in Canada, South Africa, the U.S. and Hungary. Worsening inequality due to technological advances is a particular concern among the more highly educated in countries such as Japan, South Africa and Brazil."); *cf.* WEST, *supra* note 6, at 127–54 (describing political implications of automation, including unequal distribution of its benefits).

30 *See generally* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 148–62.

31 *See, e.g.*, HARRIS ET AL., *supra* note 18, at 37; RIFKIN, *supra* note 6, at 13; Jeffrey D. Sachs, *R&D, Structural Transformation, and the Distribution of Income*, *in* THE ECONOMICS OF ARTIFICIAL INTELLIGENCE: AN AGENDA, *supra* note 2, ch. 13, https://perma.cc/RE8F-TBLA.

32 *See* HARRIS ET AL., *supra* note 18, at 31–41.

33 There are primarily two ways in which an increase in capital share as a result of automation could come about. First, the marginal productivity of automation technology could rise, and that of human work could rise less quickly or fall. In such a scenario, the capital share will increase. *See, e.g.*, HARRIS ET AL., *supra* note 18, at 37.

Second, the owners of AI capital could have market power, allowing them to extract profits from their capital and potentially increase the capital share regardless of the marginal productivity of AI capital. Market concentration across most industries has indeed increased substantially over recent decades, at least in the United States. *See* Jan De Loecker & Jan Eeckhout, The Rise of Market Power and the Macroeconomic Implications (Nat'l Bureau Econ. Research, Working Paper No. 23687, 2017), https://perma.cc/KSL4-MYA3. This trend appears to be due at least in part to technological developments which favor economies of scale and has indeed increased profit shares in the affected industries. *See* Gustavo Grullon et al., *Are US Industries Becoming More Concentrated?*, 23 REV. FIN. 697 (2019), https://perma.cc/N7Y7-2PQ5. If the future AI market obeys this trend, transformative technological developments in AI will result in substantial profits for the leading firm(s), at the expense of prices and wages.

34 *See* LEE, *supra* note 18, at 146 ("The positive-feedback loop generated by increasing amounts of data means that AI-driven industries naturally tend towards monopoly, simultaneously driving down prices and eliminating competition between firms."); Gus Bekdash, *Using Human History, Psychology, and Biology to Make AI Safe for Humans*, *in* ARTIFICIAL INTELLIGENCE SAFETY AND SECURITY 167, 171 (Roman V. Yampolskiy ed., 2018); Allan Dafoe, Co-Director, Future of Humanity Inst. Governance of AI Program, Governance of Artificial Intelligence, Talk at Effective Altruism Global 2017: London (Nov. 5, 2017), https://www.youtube.com/watch?v=RWKHx2bE1H4 ("It's plausibly the case that AI industries and services are natural global oligopolies or monopolies."); Carrick Flynn, *AI Governance Landscape*, EFFECTIVE ALTRUISM (Sept. 25, 2018), https://perma.cc/29ZY-FWXA ("[M]ost AI services have characteristics of a natural monopoly or an oligopoly. There's not much room for competition, and there's not really much room for small businesses."); *cf.* Bostrom et al., *supra* note 11, at 10 (discussing first-mover advantages in AI).

35 To be clear, our concern here is not one of the typical concerns over supracompetitive oligopoly pricing. *Cf.* Haydn Belfield, From Tech Giants to a Tech Colossus: Objections to the Windfall Clause (2019) (unpublished manuscript) (on file with author) (asserting that the Windfall Clause fails to address monopoly power). The Windfall Clause is poorly equipped to combat monopoly pricing. *Accord id.* Rather, the primary concern is simply the extreme concentration of wealth and economic power resulting from a hypothetical AI oligopoly.

At the same time, the Windfall Clause does address *some* harms from monopoly pricing. Classical antitrust theory identifies two distinct harms from monopoly pricing: deadweight loss from supramarginal pricing and transfer of surplus from consumers to producers. The Windfall Clause does not address the first but offers a partial remedy to the second through redistribution. Traditional antitrust remedies are more appropriate for the first problem and anticompetitive worries generally.

36 *See, e.g.*, ALLEN, *supra* note 6, at 14–15, 30–32, 49–51; BRYNJOLFSSON & MCAFEE, *supra* note 6, at 148–62 (discussing "winner-take-all" dynamics in the digital economy); FORD, *supra* note 6, at 75–79.

37 *See* DAFOE, *supra* note 6, at 32.

38 Hal R. Varian, High-Technology Industries and Market Structure (Sept. 4, 2001) (unpublished manuscript), https://perma.cc/DZ2B-E7GT; *see also* COUNCIL OF ECON. ADVISERS, BENEFITS OF COMPETITION AND INDICATORS OF MARKET POWER (2016), https://perma.cc/CF7L-6PS4; *See* Loecker & Eeckhout, *supra* note 33.

39 Note, however, that there is substantial uncertainty about the impacts of AI. *See* DAFOE, *supra* note 6, at 18–21. For example, the already-diminishing returns to data, *see* Hal Varian, Chief Economist, Google, Keynote Address at the Lear Conference: Is There a Data Barrier to Entry? (June 25, 2015), https://perma.cc/U2LF-W4RJ, and greater learning efficiency of new algorithms may combine to obviate the need for the massive datasets in the future—perhaps all that will be needed will be a connection to the Internet. Simulations can perform a similar role, effectively substituting computing power for data. *See* Tim Hwang, Computational Power and the Social Impact of Artificial Intelligence 34–44 (Mar. 23, 2018) (unpublished manuscript), https://ssrn.com/abstract=3147971.

40 *Cf.* K. Eric Drexler, *Reframing Superintelligence: Comprehensive AI Services as General Intelligence* 50–62, 73–78 (Future of Humanity Inst., Technical Rept. No. 2019-1), https://perma.cc/4Y7P-EPGB (describing a model in which advanced AI services are performed not by a single agent, but by a coordinated system of multiple AI service providers).

41 *See generally* Stefan Gosepath, *Equality*, STAN. ENCYCLOPEDIA PHIL. (Edward N. Zalta ed., 2011), https://perma.cc/33GP-FULN.

42 *See, e.g.*, GALLUP NEWS SERV., GALLUP SOCIAL POLL SOCIAL SERIES: MOOD OF THE NATION 2 (2018), https://perma.cc/3PZL-MZRR (66% of Americans unsatisfied with "[t]he way income and wealth are distributed in the U.S."); PEW RESEARCH CTR., ECONOMIES OF EMERGING MARKETS BETTER RATED DURING DIFFICULT TIMES 19–24 (2013), https://perma.cc/2HN9-GM6N ("In 31 of the 39 countries surveyed, half or more of the population believe that the gap between the rich and the poor is a *very* big problem in their societies.").

43 *See, e.g.*, Néstor Gandelman & Rubén Hernández-Murillo, Risk Aversion at the Country Level (Fed. Reserve Bank of St. Louis Working Paper No. 2014-005B), https://perma.cc/YB3U-98GM.

44 Of course, such transfers have distortionary effects that must be accounted for. *See generally* Daniel Carroll, *The Economics of Taxation*, FOREFRONT, Spring 2012, at 16, https://perma.cc/YU7D-YGLA.

45 The upwards redistribution of income from workers to shareholders from AI might also suppress aggregate demand, hampering overall economic growth. *See* HARRIS ET AL., *supra* note 18, at 31–41. However, the overall effect of inequality on growth is ambiguous:

> We've surveyed one channel whereby inequality may increase growth [higher investment levels] and three channels by which it may slow growth [less total human capital, higher likelihood of efficiency-reducing taxation, and higher likelihood of unrest]. A natural question is: Which of these effects dominate? Does inequality raise or lower growth? Unfortunately, the available statistical data are unable to answer this question. Although some economists claim to find evidence that inequality is on average bad for growth, others claim the data point in the opposite direction. One of the obstacles to getting a clear answer is that inequality itself is difficult to measure. Thus, we cannot say that inequality has no effect on growth, only that the data are not yet sufficient to tell us what the possible effect is.

DAVID WEIL, ECONOMIC GROWTH: INTERNATIONAL EDITION 409 (3d e., Kindle ed. 2016).

46 *See* Luca Agnello et al., *Income Inequality, Fiscal Stimuli and Political (In)stability*, 24 INT'L TAX & PUB. FIN. 484 (2017); Alberto Alesina & Roberto Perrotti, *Income Distribution, Political Instability, and Investment*, 40 EUR. ECON. REV. 1203 (1996); *Inequality, Populism, and Redistribution*, IGM FORUM (Oct. 9, 2019, 11:08 AM), https://perma.cc/S333-ECM7?type=image. *But see* Richard A. Posner, *Equality, Wealth, and Political Stability*, 13 J.L., ECON., & ORG. 344 (1997); *but cf.* Luisa Blanco & Robin Grier, *Long Live Democracy: The Determinants of Political Instability in Latin America*, 45 J. DEV. STUD. 76 (2009) (finding a nonlinear relationship between income inequality and political instability). For more on the relationship between the Windfall Clause and stability, *see infra* § I(B)(1)(iii).

47 *See* Derek A. Epp & Enrico Borghetto, Economic Inequality and Legislative Agendas in Europe (Feb. 5, 2018) (unpublished working paper), https://

enricoborghetto.netlify.com/working_paper/EuroInequality.pdf; Martín A. Rossi, Wealth and Political Power: Evidence from the Foundation of Buenos Aires (Aug. 2011) (unpublished manuscript), https://perma.cc/FT6Q-TV9J; *cf.* Benjamin I. Page et al., *Democracy and the Policy Preferences of Wealthy Americans*, 11 PERSP. ON POL. 52 (2013) (finding that the policy preferences of very wealthy Americans are more likely to be represented in law, despite those preferences differing from most Americans).

48 *Cf.* Bekdash, *supra* note 34, at 171 (discussing the risk of "perpetual incumbency"). *But cf.* Madeleine I.G. Daepp et al., *The Mortality of Companies*, 12 J. ROYAL SOC'Y INTERFACE 20150120 (2015) ("[T]he typical half-life of a publicly traded company is about a decade, regardless of business sector.").

49 For discussion of how the Windfall Clause relates to other policy proposals that can achieve this end, see *infra* § V(D).

50 Bostrom et al., *supra* note 11, at 8.

51 This concern already looms large among many workers globally. *See* Wike & Stokes, *supra* note 6 ("In all 10 advanced and emerging economies polled, large majorities say that in the next 50 years robots and computers will probably or definitely do much of the work currently done by humans. In three countries—Greece, South Africa and Argentina—four-in-ten or more believe this will definitely happen. And most believe that increasing automation will have negative consequences for jobs. Large majorities in each nation surveyed think ordinary people will have a hard time finding jobs as a result of automation. Relatively few predict new, better-paying jobs will be created by technological advances.").

52 *Cf.* HOROWITZ ET AL., *supra* note 6, at 15–19. Most individuals believe that governments are responsible for ensuring workers' success in the face of job automation. *See* Wike & Stokes, *supra* note 6 ("Government looms large in the minds of many. Nearly eight-in-ten Argentines say government has a lot of responsibility for ensuring that the nation's workforce has the right skills and education to succeed in the future, and more than seven-in-ten hold this view in South Africa, Brazil, Greece and Italy. Only in the U.S. do fewer than half [35%] believe the government has a lot of responsibility for preparing the nation's workforce.").

53 *Cf.* RIFKIN, *supra* note 6, at 208–17; Seymour Martin Lipset, *Some Social Requisites of Democracy: Economic Development and Political Legitimacy*, 53 AM. POL. SCI. REV. 69 (1959); Daniel Treisman, *Income, Democracy, and Leader Turnover*, 59. AM. J. POL. SCI. 927 (2015); Therese F. Azeng & Thierry U. Yogo, *Youth Unemployment and Political Instability in Selected Developing Countries* (Afr. Dev. Bank Grp., Working Paper No. 171, 2013), https://perma.cc/4SG3-6HVM; Aniruddha Bagchi & Jomon A. Paul, *Youth Unemployment and Terrorism in the MENAP (Middle East, North Africa, Afghanistan, and Pakistan) Region*, 64 SOCIO-ECON. PLANNING SCI. 9 (2018) ("Youth unemployment increases domestic terrorism in MENAP countries."); Efraim Benmelech et al., *Economic Conditions and the Quality of Suicide Terrorism* (Nat'l Bureau Econ. Research, Working Paper No. 16320, 2010), https://perma.cc/RRN3-ASLP ("High levels of unemployment enable terror organizations to recruit more educated, mature and experienced suicide terrorists who in turn attack more important Israeli targets."). *But see* Benmelech et al., *supra*, at 2 ("A growing body of empirical literature shows that poverty and economic conditions are not directly correlated with the occurrence of terrorism. The lack of correlation between the frequency of terror attacks and economic conditions, combined with the privileged background of suicide terrorist found in earlier studies, suggests that economic conditions do not have a direct and straightforward effect on terrorism." (citations omitted))

54 *Cf.* Bruno Biais & Enrico Perotti, *Machiavellian Privatization*, 92 AM. ECON. REV. 240 (2002) ("When median-class voters a priori favor redistributive policies, a strategic privatization program allocating them enough shares can induce a voting shift away from left-wing parties whose policy would reduce the value of shareholdings."); Denise DiPasquale & Edward L. Glaeser, *Incentives and Social Capital: Are Homeowners Better Citizens?*, 45 J. URB. ECON. 354 (1999) ("Homeownership may encourage investment in local amenities and social capital, because homeownership gives individuals an incentive to improve their community and because homeownership creates barriers to mobility."); Saumitra Jha, *Financial Asset Holdings and Political Attitudes: Evidence from Revolutionary England*, 130 Q.J. ECON. 1485 (2015) (suggesting that incentives created by financial asset-holding enabled the end of dictatorial rule and democratic reforms in England); Markku Kaustia et al., *Stock Ownership and Political Behavior: Evidence from Demutualizations*, 62 MGMT. SCI. 945 (2016) (finding that exogenous increased stock ownership generated greater support for right-of-center policies).

55 *See, e.g.*, JEFFREY DING, DECIPHERING CHINA'S AI DREAM 30–31 (2018), https://perma.cc/SUV6-3K4S; IEEE, ETHICALLY ALIGNED DESIGN (version 2, n.d.), https://perma.cc/X7PN-FTFA; *Asilomar AI Principles*, *supra* note 6; Brin, *supra* note 6; *DeepMind Ethics & Society*, DEEPMIND, https://perma.cc/H8MY-QPG9 (archived Oct. 22, 2019); Alex Hern, *Chinese Search Firm Baidu Joins Global AI Ethics Body*, GUARDIAN (Oct. 17, 2018), https://perma.cc/Z4C3-FRXD; Elsa Kania & Rogier Creemers, *Xi Jinping Calls for 'Healthy Development' of AI (Translation)*, NEW AMERICA: DIGICHINA (Nov. 5, 2018), https://perma.cc/9JD2-M5WT; *OpenAI Charter*, OPENAI (Apr. 9, 2018), https://perma.cc/J29Y-2PU8; *Our Approach: Microsoft AI Principles*, MICROSOFT, https://perma.cc/52D9-E7PV (archived Oct. 31, 2018); Sundar Pichai, *AI at Google: Our Principles*, GOOGLE: BLOG (June 7, 2018), https://perma.cc/UV53-7SV3; *Tenets*, PARTNERSHIP ON AI, https://perma.cc/SJQ9-VH4Q (archived Oct. 31, 2018).

56 *See, e.g.*, *AI for Good Global Summit*, INTL. TELECOMM. UNION, https://aiforgood.itu.int/ (last visited Nov. 6, 2019).

57 *See, e.g.*, Thomas Metzinger, *EU Guidelines: Ethics washing made in Europe*, DER TAGESSPIEGEL (Apr. 8, 2019, 3:48 PM), https://perma.cc/S5TH-A444.

58 *See, e.g.*, Eric Johnson, *How will AI change your life? AI Now Institute founders Kate Crawford and Meredith Whittaker explain.*, RECODE DECODE (Apr 8, 2019, 6:20 AM), https://perma.cc/2BM2-J9VC.

59 *See, e.g.*, Daniel Susser, *Ethics Alone Can't Fix Big Tech*, SLATE (Apr. 17, 2019, 11:45 AM), https://perma.cc/8HQU-UFKH.

60 *Cf. infra* § III.

61 *See, e.g.*, PWC, HANDBOOK ON CORPORATE SOCIAL RESPONSIBILITY IN INDIA 8 (2013), https://perma.cc/5J48-R3CG; Veronica Besmer, *The Legal Character of Private Codes of Conduct: More Than Just A Pseudo-Formal Gloss on Corporate Social Responsibility*, 2 HASTINGS BUS. L.J. 279, 291–92 (2006); Miriam A. Cherry, *The Law and Economics of Corporate Social Responsibility and Greenwashing*, 14 U.C. DAVIS BUS. L.J. 281, 283 (2014); Paul R. Portney, *The (Not So) New Corporate Social Responsibility: An Empirical Perspective*, 2 REV. ENVTL. ECON. & POL'Y 261, 265 (2008); James J. Waters, *Achieving World Trade Organization Compliance for Export Processing Zones While Maintaining Economic Competitiveness for Developing Countries*, 63 DUKE L.J. 481, 518 (2013).

62 *See* Marc Orlitzky & John D. Benjamin, *Corporate Social Performance and Firm Risk: A Meta-Analytic Review*, 40 BUS. & SOC'Y 369, 391 (2001) ("[I]n the data set covering the past two decades, risk is negatively correlated with CSP [i.e., corporate social performance]. In fact, among all risk measures, high CSP appears to be most highly negatively correlated with total market risk. Furthermore, the better a firm's CSP reputation, the lower is its risk. Thus, a firm that is socially responsible and responsive may be able to increase interpersonal trust between and among internal and external stakeholders, build social capital, lower transaction costs, and, therefore, ultimately reduce uncertainty about its financial performance."); *see also* Judd F. Sneirson, *The Sustainable Corporation and Shareholder Profits*, 46 WAKE FOREST L. REV. 541, 543 n.10 (2011) (collecting sources). *But see* David J. Vogel, *Is there a Market for Virtue?: The Business Case for Corporate Social Responsibility*, 47 CAL. MGMT. REV. 19 (2005).

63 *See* Marc Orlitzky et al., *Corporate Social and Financial Performance: A Meta-analysis*, 24 ORG. STUD. 403, 417–19 (2003).

64 *See id.* at 423–24.

65 DAVID J. VOGEL, THE MARKET FOR VIRTUE 73 (2006) ("For a subset of firms, CSR does appear to make business sense. These firms fall into two broad categories . . . For a few firms, CSR is a part of their corporate strategy and business identity. . . . A second category of firms for whom CSR makes business sense are those that have been targeted by activists, who are concerned that they could be targeted, largely because of the visibility of their brands."); Michael E. Porter & Mark R. Kramer, *Strategy And Society: The Link Between Competitive Advantage and Corporate Social Responsibility*, HARV. BUS. REV., Dec. 2006, at 37.

66 *See* Employees of Microsoft, *An Open Letter to Microsoft: Don't Bid on the US Military's Project JEDI*, MEDIUM (Oct. 12, 2018), https://perma.cc/PGU4-HHZR; Sheera Frenkel, *Microsoft Employees Question C.E.O. Over Company's Contract With ICE*, N.Y. TIMES (July 26, 2018), https://perma.cc/Z2CZ-9WHZ; *Google Should Not Be In Business Of War, Say Employees*, BBC NEWS (Apr. 5, 2018), https://perma.cc/J63U-TVDG; Andrea Shalal, *Researchers To Boycott South Korean University Over AI Weapons Work*, REUTERS (Apr. 4, 2018), https://perma.cc/A2H6-NGN9; Laura Sydell, *Tech Workers Demand CEOs Stop Doing Business With ICE, Other U.S. Agencies*, NPR (July 14, 2018), https://perma.cc/494V-VHEL; *cf.* Ryan Gallagher, *Senior Google Scientist Resigns over "Forfeiture of Our Values" in China*, THE INTERCEPT (Sept. 13, 2018 12:15 PM), https://perma.cc/S8RQ-YDE2; Naomi Nix, *Google Drops Out of Pentagon's $10 Billion Cloud Competition*, BLOOMBERG (Oct. 8, 2018), https://www.bloomberg.com/news/articles/2018-10-08/google-drops-out-of-pentagon-s-10-billion-cloud-competition ("'We are not bidding on the . . . contract because . . . we couldn't be assured that it would align with our AI Principles,' a Google spokesman said in a statement. . . . The Tech Workers Coalition, which advocates for giving employees a say in technology company decisions, said in a statement that Google's decision to withdraw from the cloud competition stemmed from 'sustained' pressure from tech workers who 'have significant power, and are increasingly willing to use it.'"); Brad Smith, *Facial Recognition Technology: The Need For Public Regulation And Corporate Responsibility*, MICROSOFT: MICROSOFT ON THE ISSUES (July 13, 2018), https://perma.cc/LQ28-Z2BL.

67 *Cf.* Grace et al., *supra* note 5, at 2 ("Forty-eight percent of respondents [i.e., AI experts] think that research on minimizing the risks of AI should be prioritized by society more than the status quo (with only 12% wishing for less)."); Peter Holley, *Tech Leaders: Killer Robots Would Be 'Dangerously Destabilizing' Force In The World*, WASH. POST (July 19, 2018), https://perma.cc/R83K-GNVL ("In total, more than 160 organizations and 2,460 individuals from 90 countries promised this week not to participate in or support the development and use of lethal autonomous weapons.").

68 *Cf.* Daniel B. Turban & Daniel W. Greening, *Corporate Social Performance and Organizational Attractiveness to Prospective Employees*, 40 ACAD. MGMT. J. 658 (1997), https://perma.cc/HA4U-FKKW.

69 *See* ROBERT H. FRANK, WHAT PRICE THE MORAL HIGH GROUND? 71–91 (2004); Karine Nyborg & Tao Zhang, *Is Corporate Social Responsibility Associated with Lower Wages?*, 55 ENVTL. & RESOURCE ECON. 107 (2013).

70 *See* Nyborg & Zhang, *supra* note 69, at 113. A more recent study similarly found a 30% wage premium for executives in the "sin" industries of alcohol, gambling, and tobacco. *See* Jiri Novak & Pawel Bilinski, *Social Stigma and Executive Compensation*, 96 J. BANKING & FIN. 169 (2015).

71 *See* Hannes Koppel & Tobias Regner, *Corporate Social Responsibility in the Work Place*, 17 EXPERIMENTAL ECON. 347 (2014); Lea Cassar, *Job Mission as a Substitute for Monetary Incentives: Experimental Evidence* (Univ. of Zurich, Dep't of Econ., Working Paper No. 177, 2014), https://perma.cc/7NXW-R3EG.

72 *See, e.g.*, Hal Varian, *Artificial Intelligence, Economics, and Industrial Organization* (Nat'l Bureau Econ. Research, Working Paper No. 24839, July 2018) ("[T]he main bottleneck [to industrial implementation of AI is] now . . . analysts who can implement these machine learning systems."); Caleb

Watney, *Reducing Entry Barriers in the Development and Application of AI* 2 (R Street Pol'y Study No. 153, Oct. 2018), https://perma.cc/52C7-C53G ("Perhaps the single biggest bottleneck in AI development and application today is the supply of skilled data scientists and machine-learning engineers. Typical AI specialists can expect to earn between $300,000 and $500,000 at top tech firms; numbers that are significantly higher than their peers in other computer-science-related subfields. In addition to these ballooning salaries, industry experts . . . have pointed to the scarcity of adequate AI talent as the largest factor behind slow application in the economy." (footnote omitted)); Jeremy Kahn, *Sky-High Salaries Are the Weapons in the AI Talent War*, BLOOMBERG BUSINESSWEEK (Feb. 13, 2018), https://www.bloomberg.com/news/articles/2018-02-13/in-the-war-for-ai-talent-sky-high-salaries-are-the-weapons; Cade Metz, *A.I. Researchers Are Making More Than $1 Million, Even at a Nonprofit*, N.Y. TIMES (Apr. 19, 2018), https://perma.cc/LQ9H-6GSD; Cade Metz, *Tech Giants Are Paying Huge Salaries for Scarce A.I. Talent*, N.Y. TIMES (Oct. 22, 2017), https://perma.cc/5LAH-QJFE. *But see* Peter Yuen, *Million Dollar Salaries for AI Researchers? Well, We Quants Have Seen this Movie Before*, MEDIUM: GOOD AUDIENCE (Apr. 22. 2018), https://perma.cc/Q455-BFWA.

73 *See* Daniel W. Greening & Daniel B. Turban, *Corporate Social Performance as a Competitive Advantage in Attracting a Quality Workforce*, 39 BUS. & SOC'Y 254 (2000); M. Todd Henderson & Anup Malani, *Corporate Philanthropy and the Market for Altruism*, 109 COLUM. L. REV. 571, 578 (2009); Jeanne M. Logsdon et al., *Corporate Philanthropy: Strategic Responses to the Firm's Stakeholders*, 19 NONPROFIT & VOLUNTARY SECTOR Q. 93, 104 (1990); Lance Moir & Richard J. Taffler, *Does Corporate Philanthropy Exist?: Business Giving to the Arts in the U.K.*, 54 J. BUS. ETHICS 149, 154–57 (2004).

74 Although the Windfall Clause is a form of self-regulation, "[i]ndustry prefers voluntary compliance to avoid the costs of formal compliance." Bekdash, *supra* note 34, at 192.

75 *See, e.g.*, Richard Waters, *The Great Silicon Valley Land Grab*, FIN. TIMES (Aug. 24, 2017), https://www.ft.com/content/82bc282e-8790-11e7-bf50-e1c239b45787.

76 *See supra* § I(B)(1)(ii).

77 *See, e.g.*, Lina M. Khan, Note, *Amazon's Antitrust Paradox*, 126 YALE L.J. 710, 746–83 (2017), Matt Stoller, *The Evidence is Piling Up—Silicon Valley is Being Destroyed*, U.K. BUS. INSIDER (Apr. 19, 2017), https://perma.cc/3GY7-6ESV; *Technology Firms Are Both the Friend and the Foe of Competition*, ECONOMIST (Nov. 15, 2018), https://perma.cc/Q4BE-R6TK; Kenneth P. Vogel, *New America, a Google-Funded Think Tank, Faces Backlash for Firing a Google Critic*, N.Y. TIMES (Sept. 1, 2017), https://perma.cc/HY2G-U3WU; David Weigel, *Breaking From Tech Giants, Democrats Consider Becoming An Antimonopoly Party*, WASH. POST, (Sept. 4, 2017), https://perma.cc/2NL3-6TA4; *cf. History's Biggest Firms*, ECONOMIST, (July 5, 2018), https://perma.cc/RQD2-THKS.

78 *See, e.g.*, Jim Brusden, *EU Efforts to Impose Heavier Taxes on Tech Giants Gather Momentum*, FIN. TIMES (Sept. 16, 2017), https://www.ft.com/content/10e81e8c-9ace-11e7-8cd4-932067fbf946 (detailing EU efforts to levy higher taxes on tech firms).

79 *See, e.g.*, David Remnick, *Cambridge Analytica and a Moral Reckoning in Silicon Valley*, NEW YORKER (Apr. 2, 2018), https://perma.cc/N7GJ-W9G2.

80 *Cf.* Eve Smith, *The Techlash against Amazon, Facebook, and Google—And What They Can Do*, ECONOMIST (Jan. 20, 2018), https://perma.cc/ZQS8-6MDU ("And there's another lesson from the robber barons [on how to avoid regulation]—one that some of you and your peers have already embraced. Philanthropy can change people's opinions and shape your legacies into the far future. In part because you do not employ as many people as corporate giants of previous eras, it is critical to think about local initiatives that can woo public opinion around the world. Mark [Zuckerberg] has made the biggest strides in setting up a foundation. The rest of you could form personal or corporate foundations, too.").

81 *See* VED P. NANDA ET AL., *The Rise of Stakeholder Related Risks to Large Development and Extractive Projects*, 1 THE LAW OF TRANSNATIONAL BUSINESS TRANSACTIONS § 1:9 ("In this way, CSR embodies more than simply corporate philanthropy. Rather, in some industries (such as the extractive industry), it has become deeply embedded in business risk mitigation strategies and has been accompanied by a widespread belief that CSR efforts (more specifically community investment programs) are critical to obtaining a social license to operate."); Robert C. Bird et. al., *Corporate Voluntarism and Liability for Human Rights in A Post-*Kiobel *World*, 102 KY. L.J. 601, 630 (2014); Jeremy S. Goldstein, *CSR Best Practice for Abolishing Child Labor in the Travel and Tourism Industry*, 44 DENV. J. INT'L L. & POL'Y 475, 515 (2016); Alexandra Guáqueta, *The New Social License to Operate and the Role of Legal Advisors*, 105 AM. SOC'Y INT'L L. PROC. 303, 304 (2011) ("Some companies responded to these multi-faceted [sic] challenges [of operating in developing nations] by working with NGOs and states to come up with practical solutions both to prevent negative impacts on the ground and to better manage their reputational problems."); Henderson & Malani, *supra* note 73, at 578; Ashton B. Inniss, *Rethinking Political Risk Insurance: Incentives for Investor Risk Mitigation*, 16 SW. J. INT'L L. 477, 493 (2010) ("The profit motive drives corporate aid programs."); Lisa J. Laplante & Suzanne A. Spears, *Out of the Conflict Zone: The Case for Community Consent Processes in the Extractive Sector*, 11 YALE HUM. RTS. & DEV. L.J. 69, 84 (2008) ("Even if companies embraced sustainable development out of a sense of charitable philanthropy at the outset, they will be required to make progress on the issue in order to maintain not only their social license to operate, but also their competitive advantage."); Thomas P. Lyon & John W. Maxwell, *Corporate Social Responsibility and the Environment: A Theoretical Perspective*, 2 REV. ENVTL. ECON. & POL'Y 240, 245 (2008); Ibironke T. Odumosu-Ayanu, *Foreign Direct Investment Catalysts in West Africa: Interactions with Local Content Laws and Industry-Community Agreements*, 35 N.C. CENT. L. REV. 65, 82 (2012) ("[Community development agreements] are strategic community investment initiatives fostering social licenses to operate in host communities."); *Political Risk Insurance: Is it Necessary?* § 5, PRACTICAL LAW PRACTICE NOTE 5-503-9151 (2017); Rachael E. Salcido, *Enduring Optimism: Examining the Rig-to-Reef Bargain*, 32 ECOLOGY L.Q. 863 (2005); Gare A. Smith, *An Introduction to Corporate Social Responsibility in the Extractive Industries*, 11 YALE HUM. RTS. & DEV. L.J. 1, 4 (2008) ("[E]ffective and meaningful CSR policies are increasingly regarded as a necessary prerequisite to securing a social license to operate . . ."); *cf.* David P. Baron, *Private Politics, Corporate Social Responsibility, and Integrated Strategy*, 10 J. ECON. & MGMT. STRATEGY 7, 44 (2001) ("Firms that are sensitive to how such [activist] groups might evaluate their practices and whether those groups are likely to act on that evaluation may proactively change their practices to avoid even the threat of becoming the target of an activist."); Neil Gunningham et al., *Social License and Environmental Protection: Why Businesses Go Beyond Compliance*, 29 LAW & SOC. INQUIRY 307 (2004) (describing the "social license to operate").

82 *Cf.* NANDA ET AL., *supra* note 81, § 1:9; Odumosu-Ayanu, *supra* note 81, at 82.

83 Reasonable definiteness of essential terms is generally necessary for a binding contract. *See generally* RESTATEMENT (SECOND) OF CONTRACTS § 33 (AM. LAW INST. 1981).

84 *See supra* §§ I(B)(1)(i)–(ii).

85 *See supra* § I(B)(1).

86 *See, e.g.*, R.G. Grp., Inc. v. Horn & Hardart Co., 751 F.2d 69, 75–76 (2d Cir. 1984) (applying New York law).

87 *See infra* § II(B).

88 A number of studies have suggested that increasing the corporate tax rate in a country tends to decrease the rate of investment by corporations, *see* Jonathan Lewellen & Katharina Lewellen, *Investment and Cash Flow: New Evidence*, 51 J. FIN. & QUANT. ANALYSIS 1135 (2016), https://perma.cc/XXS4-F92H; Eric Ohrn, *The Effect of Corporate Taxation on Investment and Financial Policy: Evidence from the DPAD*, 10 AM. ECON. J.: ECON. POL'Y 272 (2018), https://perma.cc/SA22-Q98T; Simeon Djankov et al., The Effect Of Corporate Taxes On Investment And Entrepreneurship (Nat'l Bureau Econ. Research, Working Paper No. 13756, 2008), https://perma.cc/R7N8-6WGL, although the magnitude of the effect remains controversial. A sufficiently large decrease in the rate of investment could be expected to adversely affect the rate of economic growth. We expect the adverse effects of an increase in a signatory's giving obligations to be similar to the adverse effects of an increase in the tax rate experienced by the signatory.

89 Saudi Aramco was the most profitable company in 2018, with approximately $111 billion in profits. *See* Fred Imbert, *Saudi Aramco made $111 billion in 2018, topping Apple as the world's most profitable company*, CNBC (Apr. 1, 2019, 3:55 PM). 2018 GWP was $85.791 trillion. *See GDP (current US$)*, WORLD BANK (2019), https://perma.cc/YQ7L-DJU6. This puts Saudi Aramco's profits at 0.129% of GWP, placing it in the second bracket.

90 Apple Inc. was the second most profitable company in 2018, with approximately $59,531,000,000 in profits. *See* Imbert, *supra* note 89.

91 A typical company in the Fortune 500 gives approximately 1% of its pre-tax earnings. *See* O'Neil et al., *Corporate Cash Giving Rises 5%, Exclusive Chronicle Survey Shows*, CHRON. PHILANTHROPY (Sept. 5, 2018), https://www.philanthropy.com/article/Corporate-Cash-Giving-Rises/244391.

92 *Cf.* Daniel Bunn, *Corporate Tax Rates Around the World, 2018*, TAX FOUNDATION (Nov. 27, 2018), https://perma.cc/MAD4-E8H4.

93 However, designation of a specific, unbiased, expert body to determine GWP, *see supra* note 89, should ease administration of this determination. *See generally* Cullen O'Keefe, Stable Agreements in Turbulent Times: A Legal Toolkit for Constrained Temporal Decision Transmission 25–27 (Future of Humanity Inst. Technical Report, 2019), https://perma.cc/W894-ADVV.

94 *See id.*

95 *See infra* § V(A)(2).

96 2061 is the median year by which AI experts predicted that AI would outperform humans at all economically relevant tasks. *See* Grace et al., *supra* note 5, at 2.

97 *See Real GDP Long-Term Forecast*, OECD, https://perma.cc/4DZF-LBZZ (archived July 27, 2019).

98 This is more than Americans' combined annual charitable giving. *See* Giving USA, Giving USA 2019*: Americans Gave $427.71 Billion To Charity In 2018 Amid Complex Year For Charitable Giving*, GIVING USA (June 18, 2019, 3:07 PM), https://perma.cc/4JNC-QNJ4. It is about 10% of the 2019 US federal budget. *See Budget*, CONG. BUDGET OFFICE, https://perma.cc/3MCH-Y4SJ (archived July 27, 2019).

99 *See supra* note 91.

100 The current US federal corporate income tax rate is 21%. *See* I.R.S., I.R.S. PUBL'N 542, at 16 (2019), https://perma.cc/8UP5-XXB2.

101 *See infra* § V(B)(1).

102 *See Contingent Convertible (CoCo)*, *in* A DICTIONARY OF FINANCE AND BANKING (Jonathan Law ed., 5th ed. 2014) (ebook). For tax purposes, CoCos are part of a class of assets known as "contingent payment debt instruments (CDPIs)". *See* Robert Willens, *Cuckoo for Coco Puffs?*, CFO (May 22, 2002), https://perma.cc/B6B8-K49K.

103 A CoCo Windfall Clause that would give a controlling stake to bondholders in case of a windfall would probably be subject to enhanced scrutiny under Revlon, Inc. v. MacAndrews & Forbes Holdings, Inc., 506 A.2d 173 (Del. 1986). *Revlon* requires that, in "every case in which a fundamental change of corporate control occurs or is contemplated," the board must attempt to maximize sale price. *See* Barkan v. Amsted Indus., Inc., 567 A.2d 1279, 1286–88 (Del. 1989). This leaves the firm with two options. If they wish to transfer control to a single bondholder in case of a Windfall,

they must engage in a process that satisfies the *Revlon* duties. *Cf. id.* Alternatively, they can give bondholders a minority stake and thereby avoid invoking *Revlon*. *Cf id.*

If the corporation already has a controlling shareholder, however, then the "controlling stockholder is free to consider its interests alone in weighing the decision to sell its shares or, having made such a decision, evaluating the adequacy of a given price." *In re* Delphi Fin. Grp. S'holder Litig., No. CIV.A. 7144-VCG, 2012 WL 729232, at *15 (Del. Ch. Mar. 6, 2012); *see also* Abraham v. Emerson Radio Corp., 901 A.2d 751, 753 (Del. Ch. 2006) ("Under Delaware law, a controller remains free to sell its stock for a premium not shared with the other stockholders except in very narrow circumstances."); *In Re:* Sea-Land Corp. Shareholders Litig., No. CIV. A. 8453, 13 Del. J. Corp. L. 795, 804 (Del. Ch. May 22, 1987). These cases are not exactly analogous since the contemplated transaction is not truly a sale of control but rather a contingent shift of control. Nevertheless, the underlying principles should be the same: in either case, the controlling shareholder causes control to shift from herself to another entity, so minority shareholders do not lose the opportunity to extract a control premium. I do not know of any cases on point, however.

There is, however, an important exception to a controlling shareholder's freedom to sell: a controller cannot sell to someone she suspects "[i]s either a looter or [i]s dishonest and ha[s] improper plans for [the corporation]." *See Abraham*, 901 A.2d at 759. A charity-bondholder would likely be considered such a "looter" if it plans to take assets out of the corporation. In any case, once bondholders became controllers, they would owe fiduciary duties to minority shareholders. Controlling bondholders could, however, force the corporation to pay out dividends so long as it shares those dividends *pro rata* with minority shareholders. *See* Sinclair Oil Corp. v. Levien, 280 A.2d 717, 721–22 (Del. 1971).

A final complication with a CoCo Windfall Clause that gives bondholders a majority stake is that a minority shareholder could rightly argue that, because of a control premium, the value of the majority stake is greater than the pro rata value of the minority stakes distributed in the pre-Clause market test.

Of course, a Windfall Clause could easily avoid these messy implications by paying contingent/index-linked coupons to bondholders or contingently converting bonds into nonvoting equity.

104 *See generally* Greg Brockman et al., *OpenAI LP*, OpenAI: Blog (Mar. 11, 2019), https://perma.cc/9LDZ-CPX8.

For transparency, we note that OpenAI employs lead author Cullen O'Keefe.

105 *See supra* § I(B).

106 We borrow this goal from principles of the Effective Altruism community. *See generally* Centre for Effective Altruism, Effective Altruism Handbook (2d ed. n.d.), https://perma.cc/Z8JQ-4S26; William MacAskill, Doing Good Better (2015); Peter Singer, The Life You Can Save (2009). One organization that applies these principles is the Open Philanthropy Project, *see generally* Open Philanthropy Project, https://www.openphilanthropy.org/ (last visited Jan. 6, 2019). Their mission is "to give as effectively as [they] can . . ." *Mission & Values*, Open Philanthropy Project, https://perma.cc/Y8UE-QYHV (archived Jan. 6, 2019). We hope that research of this and other organizations involved in Effective Altruism will aid in the design of a windfall distribution plan.

For transparency, we note that the Open Philanthropy Project has funded the publishers of this Report: the Future of Humanity Institute and the Berkeley Existential Risk Initiative. The publishers of this Report are also members of the wider Effective Altruism community.

107 *See The Wrong Donation Can Accomplish Nothing*, GiveWell, https://perma.cc/E7LG-FWNQ (archived June 4, 2019); *Your Dollar Goes Further When You Fund the Right Program*, GiveWell, https://perma.cc/XUG7-5V6U (archived June 4, 2019).

108 *See* Toby Ord, *The Moral Imperative toward Cost-Effectiveness in Global Health*, Ctr. for Global Dev. (Mar. 2013), https://perma.cc/63W6-ZZRJ ("In many cases, ignoring cost-effectiveness . . . means losing almost all the value that we could create. Thus there is a moral imperative to fund the most cost-effective interventions.").

109 *See supra* § I(B).

110 *Cf. supra* § I(B)(1)(ii).

111 *See also* Jon Elster, Securities Against Misrule (2013).

112 *Cf. Reasoning Transparency*, Open Philanthropy Project (Dec. 2017), https://perma.cc/HLQ5-AE37.

113 *See What is the Four-eyes Principle?*, United Nations Indus. Dev. Org., https://perma.cc/BW3U-AY3G (archived Jan. 14, 2019).

114 Several US agencies have such whistleblower rewards programs in place. *See generally* Jason Zuckerman & Matthew Stock, *What is a Whistleblower Reward?*, Zuckerman Law (June 12, 2019), https://perma.cc/KL2J-2938.

115 *See, e.g.*, Robert H. Sitkoff & Jesse Dukeminier, Wills, Trusts, and Estates 587–693 (10th ed. 2017).

116 *See, e.g.*, Global Impact Investing Network, Impact-Based Incentive Structures (2011), https://perma.cc/VY9X-HXMC.

117 *See, e.g.*, Geoffrey A. Manne, *Agency Costs and the Oversight of Charitable Organizations*, 1999 Wis. L. Rev. 227.

118 For example, there are existing proposals to use AI to advance the United Nations' Sustainable Development Goals. *See, e.g.*, U.N. Secretary-General, *Strategy on New Technologies* (2018), https://perma.cc/VXY5-PEJU; *AI for Good Global Summit*, *supra* note 56; Declan Butler, *AI Summit Aims to Help World's Poorest*, Nature: News (June 6, 2017), https://perma.cc/5QTN-4YT8; Michael Chui et al., *Notes from the AI Frontier: Applying AI for Social Good*, McKinsey Global Inst. (2018), https://perma.cc/3D9R-L8MG.

119 For the importance of China in global AI development, *see generally* Ding, *supra* note 55.

120 Cayman Islands company law might also be important, since several major Chinese AI firms are incorporated there. *See* Alibaba Grp. Holding Ltd., Annual Report (Form 20-F) (July 27, 2018), https://perma.cc/Q6GQ-3BU8; Baidu, Inc., Annual Report (Form 20-F) (Mar. 15, 2019), https://perma.cc/M8Y9-K63Z; Tencent Holdings Ltd., 2018 Interim Report (2018), https://perma.cc/V8SM-K4KP. Although we are not aware of any decisions directly on-point, relevant Caymanian company law appears broadly similar to Delaware's. *Cf.* Douglas Freeman, *New Frontiers for Asia M&A: Strategies in the Cayman Islands for Asia-Based Companies*, 2014 WL 4160103, at *3 ("[U]nder Cayman Islands law, directors must act *bona fide* and in what they consider to be the best interests of the company. Directors are considered in breach of their fiduciary duties if the court concludes that no reasonable director would have found the director's decision to be in the best interest of the company. In making its assessment, the court will not be concerned with the merits of the contested director decision from a commercial standpoint. . . . While this regime bears resemblance to the business judgment rule in Delaware, the standard of care, in practice, may be easier for a director to satisfy."). So, the Windfall Clause should be permissible for Caymanian companies too.

121 Jill E. Fisch, Questioning Philanthropy from A Corporate Governance Perspective, 41 N.Y.L. Sch. L. Rev. 1091, 1096 (1997).

122 *Id.* at 1094 (citing Kahn v. Sullivan, 594 A.2d 48, 61 (Del. 1991); Theodora Holding Corp. v. Henderson, 257 A.2d 398, 405 (Del. Ch. 1969); A.P. Smith Mfg. Co. v. Barlow, 98 A.2d 581, 590 (N.J. 1953)).

123 "A director or officer who makes a business judgment in good faith fulfills the duty [of care] if the director or officer: (1) is not interested . . . in the subject of the business judgment; (2) is informed with respect to the subject of the business judgment to the extent the director or officer reasonably believes to be appropriate under the circumstances; and (3) rationally believes that the business judgment is in the best interests of the corporation." Principles of Corp. Governance § 4.01(a) (Am. Law Inst. 1994).

124 *See Kahn*, 594 A.2d at 61 (quoting *Theodora*, 257 A.2d at 405).

125 *See id.* (quoting *Theodora*, 257 A.2d at 405).

126 I.R.C. § 170(a)(1).

127 *See id.* § 170(b)(2)(A).

128 No. CIV. A. 10823, 16 Del. J. Corp. L. 1621 (Del. Ch. Aug. 7, 1990), *aff'd sub nom. Kahn*, 594 A.2d 48.

129 *See* Appellant's Statement in Lieu of Reply Brief, *Kahn*, 594 A.2d 48 (No. 312, 1990), 1990 WL 10090250, at *45.

130 *See* Opening Brief of Objector Below-Appellant Alan R. Kahn, *Kahn*, 594 A.2d 48 (No. 301, 1990), 1990 WL 10090236, at *32.

131 *See Hammer*, 16 Del. J. Corp. L. at 1633.

132 *See Kahn*, 594 A.2d at 61.

133 *See* Dodge v. Ford Motor Co., 70 N.W. 668 (Mich. 1919). *See generally* John A. Pearce II, *The Rights of Shareholders in Authorizing Corporate Philanthropy*, 60 Vill. L. Rev. 251, 253 (2015); Hayden W. Smith, *If Not Corporate Philanthropy, Then What?*, 41 N.Y.L. Sch. L. Rev. 757, 760–62 (1997).

134 *Kahn*, 594 A.2d 48; Ella M. Kelly & Wyndham, Inc. v. Bell, 266 A.2d 878 (Del. 1970) (upholding voluntary payments to county government); *Theodora*, 257 A.2d 398; Shlensky v. Wrigley, 237 N.E.2d 776 (Ill. App. Ct. 1968) (upholding president's choice not to host lucrative nighttime baseball games, allegedly out of concern for stadium neighborhood's welfare); *Barlow*, 98 A.2d 581; Union Pac. R.R. v. Trustees, Inc., 329 P.2d 398 (Utah 1958). Note that of these, *Barlow* and *Union Pacific* were friendly suits. *See* Smith, *supra* note 133, at 761 n.21.

135 *See In re* KKR Fin. Holdings LLC S'holder Litig., 101 A.3d 980, 995–96 (Del. Ch. 2014); Off v. Ross, No. CIV.A. 3468-VCP, 2008 WL 5053448, at *3, *11 (Del. Ch. Nov. 26, 2008); *In re* Oracle Corp. Derivative Litig., 824 A.2d 917, 920, 923–24, 930–35, 940, 942 (Del. Ch. 2003).

136 *See* Se. Pennsylvania Transportation Auth. v. Volgenau, No. CV 6354-VCN, 2013 WL 4009193, at *15–16 (Del. Ch. Aug. 5, 2013).

137 *See Barlow*, 98 A.2d at 590 (upholding donation where "[t]here [wa]s no suggestion that it was made . . . to a pet charity of the corporate directors in furtherance of personal rather than corporate ends."); *see also Theodora*, 257 A.2d at 404 (quoting *Barlow*, 98 A.2d at 590). *See generally* Jayne W. Barnard, *Corporate Philanthropy, Executives' Pet Charities and the Agency Problem*, 41 N.Y.L. Sch. L. Rev. 1147 (1997).

138 *See Khan*, 594 A.2d at 61 (approving as protected by BJR corporate donation to build a museum to house CEO's personal art collection); *see also* Joseph K. Leahy, Are Corporate Super PAC Contributions Waste or Self-Dealing? A Closer Look, 79 Mo. L. Rev. 283, 366 (2014) ("Although [*Theodora*] and cases from outside Delaware have suggested that a plaintiff might successfully challenge donations to a director's pet charity, *Kahn* may foreclose this possibility in Delaware. The overarching concern of these commentators seems to be that donations to pet charities ought to be deemed self-dealing. If a charity created with the sole purpose of building an entire museum dedicated to housing the CEO's personal art collection is not a pet charity—then what is?").

139 *See* Frank H. Easterbrook & Daniel R. Fischel, The Economic Structure of Corporate Law 93, 98–100 (1991); Todd M. Aman, *Cost-Benefit Analysis of the Business Judgment Rule: A Critique in Light of the Financial Meltdown*, 74 Alb. L. Rev. 1, 13–14 (2011); Richard A. Booth, *A Minimalist Approach to Corporation Law*, 34 Ga. L. Rev. 431, 442 (2000); Melvin Aron Eisenberg, *The Director's Duty of Care in Negotiated Dispositions*, 51 U. Miami L. Rev. 579, 589 (1997); Melvin Aron Eisenberg, *The Divergence of Standards of Conduct and Standards of Review in Corporate Law*, 62 Fordham L. Rev. 437, 444–45 (1993); Daniel R. Fischel, *The Business Judgment Rule and the Trans Union Case*, 40 Bus. Law. 1437, 1442 (1985);

Daniel R. Fischel & Michael Bradley, *The Role of Liability Rules and the Derivative Suit in Corporate Law: A Theoretical and Empirical Analysis*, 71 Cornell L. Rev. 261, 270 (1986); Wayne O. Hanewicz, *Director Primacy, Omnicare, and the Function of Corporate Law*, 71 Tenn. L. Rev. 511, 547–49 (2004); Peter V. Letsou, *Implications of Shareholder Diversification on Corporate Law and Organization: The Case of the Business Judgment Rule*, 77 Chi.-Kent L. Rev. 179 (2001); Steven L. Schwarcz & Gregory M. Sergi, *Bond Defaults and the Dilemma of the Indenture Trustee*, 59 Ala. L. Rev. 1037, 1056, 1061–62 (2008); E. Norman Veasey & Christine T. Di Guglielmo, *What Happened in Delaware Corporate Law and Governance from 1992–2004? A Retrospective on Some Key Developments*, 153 U. Pa. L. Rev. 1399, 1423–24 (2005).

140 *See* Veasey & Di Guglielmo, *supra* note 139, at 1424.

141 *Cf.* Duncan v. Theratx, Inc., 775 A.2d 1019, 1022 (Del. 2001) ("[T]he standard remedy for breach of contract is based upon the reasonable expectations of the parties *ex ante*."); *In re* Tyson Foods, Inc., 919 A.2d 563, 586 (Del. Ch. 2007) ("Even were plaintiffs to maintain that this payment [of spousal benefits] constituted a pure waste of corporate assets, the relevant value for consideration would not be the *ex post* cost of benefits paid to [a consultant's wife] after her husband's death, but the cost of the *risk* placed on the company at the time of the contract."); Andaloro v. PFPC Worldwide, Inc., No. CIV.A. 20289, 2005 WL 2045640, at *9 (Del. Ch. Aug. 19, 2005) ("The [discounted cash flow] model of valuation is a standard one that gives life to the finance principle that firms should be valued based on the expected value of their future cash flows, discounted to present value in a manner that accounts for risk."); *In re* Trans World Airlines, Inc. Shareholders Litig., No. CIV.A. 9844, 1988 WL 111271, at *11 (Del. Ch. Oct. 21, 1988) ("Surely the expected value of the debt security is a most material consideration to a shareholder."), *abrogated on other grounds by* Kahn v. Lynch Commc'n Sys., Inc., 638 A.2d 1110 (Del. 1994); Easterbrook & Fischel, *supra* note 139, at 98 ("A decision is good to the extent it has a high expected value, although it may also have a high variance.").

142 *See infra* § IV(A)(1) (giving an example of the present expected cost of a Windfall Clause).

143 *See supra* § I(B)(2).

144 *See* Smith v. Van Gorkom, 488 A.2d 858, 873 (Del. 1985), *overruled on other grounds by* Gantler v. Stephens, 965 A.2d 695 (Del. 2009); Aronson v. Lewis, 473 A.2d 805, 812 (Del. 1984), *overruled on other grounds by* Brehm v. Eisner, 746 A.2d 244 (Del. 2000).

145 For related scholarship on the BJR as a justification for engaging in CSR practices, *see, e.g.,* Janet E. Kerr, *Sustainability Meets Profitability: The Convenient Truth of How the Business Judgment Rule Protects a Board's Decision to Engage in Social Entrepreneurship*, 29 Cardozo L. Rev. 623 (2007); Sneirson, *supra* note 62, at 548–54.

146 *See* Facebook Inc., 2005 Officers' Stock Plan (Nov. 8, 2005), https://perma.cc/XK6J-CUFH.

147 *See* Facebook Inc., Statement of Changes in Beneficial Ownership (Form 4) (May 21, 2012), https://perma.cc/4Q8N-SF2N.

148 *See* Facebook Inc., Statement of Changes in Beneficial Ownership (Form 4) (Dec. 20, 2013), https://perma.cc/9YLC-EGB9.

149 All options had an exercise price of $0.06, and Zuckerberg acquired 60 million shares in each transaction. *See* Facebook Inc., *supra* note 147; Facebook Inc., *supra* note 148. Facebook's stock was valued at $34.03/share for the 2012 transaction. *See FB Historical Prices*, Yahoo! Finance, https://perma.cc/GE8A-CBKZ (archived Feb. 7, 2019). Thus, in the 2012 transaction, Zuckerberg paid $3.6 million for $2.0418 billion in stock, for a net pre-tax gain of $2.0382 billion. Facebook's stock was valued at $55.12/share for the 2013 transaction. *See FB Historical Prices*, Yahoo! Finance, https://perma.cc/JJJ5-VM46 (archived Feb. 7, 2019). Thus, in the 2012 transaction, Zuckerberg paid $3.6 million for $3.3072 billion in stock, for a net pre-tax gain of $3.3036 billion.

150 *See* Dealbook, *Tracking Facebook's Valuation*, N.Y. Times: Dealbook (Feb. 1, 2012), https://perma.cc/8AHE-FHM8.

151 In fact, the stock options technically had a fair market value of zero at the time they were issued, since the exercise price was equivalent to the fair market value of the shares. *See* Facebook Inc., *supra* note 146.

152 Zhōnghuá Rénmín Gònghéguó Gōngsī Fǎ (中华人民共和国公司法) [Company Law of the People's Republic of China] (promulgated by Standing Comm. Nat'l People's Cong., Dec. 28, 2013, effective Mar. 1, 2014) [hereinafter Company Law], art. 5.

153 Virginia Harper Ho, *Beyond Regulation: A Comparative Look at State-Centric Corporate Social Responsibility and the Law in China*, 46 Vand. J. Transnat'l L. 375, 399 (2013) (footnotes omitted); *see also* Afra Afsharipour & Shruti Rana, *The Emergence of New Corporate Social Responsibility Regimes in China and India*, 14 U.C. Davis Bus. L.J. 175, 199–200 (2014).

154 *See* Ho, *supra* note 153, at 400.

155 Company Law art. 5.

156 *See, e.g.,* Jessica Marie Conrad, The Business of Business: Comparing Corporate Social Responsibility Initiatives in China and the United States, 41 Ga. J. Int'l & Comp. L. 747, 753 (2013); *cf.* Shen Wei, Corporate Law in China 91 (2015) ("The current regulatory approach [for CSR in China] is still to rely on corporate conscience.").

157 Wei, *supra* note 156, at 89; *see also id.* at 90–91.

158 Ho, *supra* note 153, at 409.

159 *See id.* at 418.

160 *See id.*; *cf.* OECD, China-OECD Project on Government Approaches to Encouraging Responsible Business Conduct ¶ 78 (2008), https://perma.cc/RP2M-X5KU ("Corporate philanthropy is not a substitute for [CSR]. In some cases, it is used as a mere public relations tool, giving 'face' to enterprises and their leaders. When this is done to divert attention from irresponsible conduct, negative consequences may counteract any social benefit. Internationally-recognised standards of corporate conduct such as the OECD Guidelines for Multinational Enterprises and the Ten Principles of the UN Global Compact do not include sections explicitly endorsing charitable giving as a form of responsible business conduct.").

161 Zhōnghuá Rénmín Gònghéguó Císhàn Fǎ (中华人民共和国慈善法) [Charity Law of the People's Republic of China] (promulgated by Nat'l People's Cong. Mar. 16, 2016, effective Sept. 1, 2016).

162 *See id.* art. 80.
In response to a general complaint that businesses that make large charitable donations do not receive enough tax concessions, the government waived the tax on the donations made from company profits, for amounts that constitute up to 12% of the profits. Under the new Law, if the donation exceeds 12% annually, the balance can be deducted from the taxable income over the following few years.
Wendy Zeldin, *China: Charity Law Adopted*, Libr. Cong.: Global Legal Monitor (May 17, 2016) (citation omitted), https://perma.cc/W3GR-KHK7.
Note that this is more generous than the American tax deduction, which is capped at 10% annually. *See supra* § III(A)(1).

163 *See* Ho, *supra* note 153, at 413.

164 *Id.* at 416.

165 *See supra* § II(A)(3).

166 *See* James R. Hitchner, Financial Valuation: Applications and Models 138-40, 190-92 (4th ed. 2016); Aswath Damodaran, *Cost of Capital by Sector (US)* (Jan. 2019), Damodaran Online, https://perma.cc/EGN9-LN5W.

167 *See* Caroline Preston, *The 20 Most Generous Companies of the Fortune 500*, Fortune (June 22, 2016), https://perma.cc/L2SM-8W2K.

168 *See CPI Inflation Calculator*, Bureau of Labor Statistics, https://perma.cc/4QRE-JAQA (archived Aug. 3, 2019).

169 *See* Preston, *supra* note 167.

170 *See, e.g.*, Christopher May, *Who's in Charge? Corporations as Institutions of Global Governance*, 1 Palgrave Comm. 15042 (2015).

171 *See generally* Responsible Bus. Alliance, http://www.responsiblebusiness.org/ (last visited Mar. 12, 2019).

172 *See* Jonathan E. Helmreich, Gathering Rare Ores 74 (1986).

173 *See id.* at 74–77.

174 The original sum of 100 million Belgian francs, *see id.* at 75, is worth between 22 million and 109 million 2015 US dollars (depending on use of goods or labor equivalence). *See* Rodney Edvinsson, *Historical Currency Converter (Test Version 1.0)*, Historical Statistics (Jan. 10, 2016), https://perma.cc/TV7N-RQNZ (archived Mar. 12, 2019).

175 *See* Helmreich, *supra* note 172, at 75.

176 *See id.* This worry would be diminished today due to widespread accessibility of public companies' financial statements.

177 *See* Vogel, *supra* note 65, at 18.

178 *See* Michael Bloom, *Doing Good Well*, Fed. Trade Comm'n: Blog (Oct. 20, 2015, 3:38 PM), https://perma.cc/9NUZ-C8WF.

179 *See* Aseem Prakash & Matthew Potoski, The Voluntary Environmentalists 46–53 (2006); David Vogel, *The Private Regulation of Global Corporate Conduct*, *in* The Politics of Global Regulation 151, 164–71 (Walter Mattli & Ngaire Woods eds., 2009).

180 FSC, https://us.fsc.org/en-us (last visited Nov. 6, 2019).

181 Rainforest Alliance, https://www.rainforest-alliance.org/ (last visited Nov. 6, 2019).

182 Fair Trade Certified, https://www.fairtradecertified.org/ (last visited Nov. 6, 2019).

183 (RED), https://www.red.org/ (last visited Nov. 6, 2019).

184 *See generally Total Giving*, Newman's Own Foundation, http://newmansownfoundation.org/about-us/total-giving/ (last visited Nov. 6, 2019).

185 Prominent examples include Ikea, Carlsberg, Novo Nordisk, Maersk, and Indian conglomerate Tata. *See* Henry Hansmann, The Governance of Foundation-Owned Firms 4–7 (2018) (unpublished manuscript), https://perma.cc/X2E4-P4QR.

186 For example, Carlsberg uses funds to support basic research and Danish national causes. *See* Anders Garrigues, *The Carlsberg Foundation Revenue and Distribution of Funds*, Carlsberg Foundation, https://perma.cc/3MVV-6SDB (last visited Nov. 6, 2019).

187 Google Inc., Registration Statement vi (Form S-1) (Apr. 29, 2004), https://perma.cc/XAR3-CJ5N.

188 *See* Associated Press, *Google Commits $1 Billion to Charity*, Bos. Globe (Oct. 12, 2005), https://perma.cc/VCS5-BCNN.

189 More accurately, OpenAI is now two legal entities: a nonprofit ("OpenAI Nonprofit") and a capped-profit limited partnership ("OpenAI LP"). *See generally* Brockman et al., *supra* note 104. LP profits above a predetermined cap on investor returns will be donated to the OpenAI Nonprofit—hence "capped profit." *See id.*

190 *See id.* Mechanisms meant to ensure fidelity to that humanitarian mission include: binding the LP to first promote OpenAI's broad humanitarian mission, even if it requires foregoing profits; having the Nonprofit's board manage the LP; requiring that board to have a majority of disinterested members; and excluding financially interested board members from decisions where the humanitarian and profit motives might conflict, such as decisions on payouts to investors and employees. *See id.*

191 *See Top 86 Largest Sovereign Wealth Fund Rankings by Total Assets*, SWFI, https://perma.cc/5FP4-YAZ7 (archived Nov. 6, 2019).
192 *See, e.g.*, John M. Hartwick, *Intergenerational Equity and the Investing of Rents from Exhaustible Resources*, 67 AM. ECON. REV. 972 (1977); Anthony Venables & Samuel Wills, *Resource Funds: Stabilization, Parking and Intergenerational Transfer*, BROOKINGS (Sept. 18, 2015), https://perma.cc/LEH5-E97C.
193 *About the Fund*, NORGES BANK INV. MGMT., https://perma.cc/GZL7-PMHD (archived Nov. 6, 2019).
194 We consider a sort of sovereign wealth fund for AI as an alternative to the Windfall Clause *infra* § V(D)(3).
195 *See* Martin J. Greenberg & Michael R. Gavin, *The Windfall Clause*, LAW OFFICES OF MARTIN J. GREENBERG (Nov. 1, 2016), https://perma.cc/7G2W-B8B8.
196 *See, e.g.*, Dhammika Dharmapala & Vikramaditya Khanna, *The Impact of Mandated Corporate Social Responsibility: Evidence from India's Companies Act of 2013*, 56 INTL. REV. L. & ECON. 92 (2018).
197 *See generally* SANFORD V. BERG & JOHN TSCHIRHART, NATURAL MONOPOLY REGULATION (1988).
198 *See, e.g.*, John Thornhill, *Why Facebook Should Pay Us A Basic Income*, FINANCIAL TIMES: OPINION (Aug. 7, 2017) ("[Mark Zuckerberg] should now live up to his rhetoric [about wealth redistribution] and launch a Facebook Permanent Fund to cover a broader universal basic income experiment. He should encourage other data businesses, such as Google, to contribute too."); Felix Salmon, *Takeaway from the New Billionaires Ranking: Zuckerberg and Bezos Don't Give Away Enough Money*, SLATE (July 8, 2018).
199 THE GIVING PLEDGE, https://givingpledge.org/ (last visited Nov. 6, 2019).
200 FOUNDERS PLEDGE, https://founderspledge.com/ (last visited Mar. 13, 2019).
201 *But see infra* § V(C)(2) (restating and responding to arguments that analogizing very profitable companies to wealthy individuals is inappropriate).
202 *See* Christian Kellner et al., *Ex-ante Commitments To "Give If You Win" Exceed Donations After A Win*, 169 J. PUB. ECON. 109 (2019).
203 *See* JOAN LUFT & MICHAEL D. SHIELDS, PSYCHOLOGY MODELS OF MANAGEMENT ACCOUNTING (2010); Herbert A. Simon, *Rational Decision Making in Business Organization*, 69 AM. ECON. REV. 493 (1979); John D. Sterman, *Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment*, 35 MGMT. SCI. 321 (1989).
204 Proposed alternatives not considered at-length here include, e.g., a value-added tax on technology services, *see* RIFKIN, *supra* note 6, at 270, a negative income tax, *see infra* note 288, and paying individuals more for providing data for AI, *see* ERIC A. POSNER & E. GLEN WEYL, RADICAL MARKETS 205–49 (2018). This is surely not an exhaustive list.
205 *See* Belfield, *supra* note 35, at 3–4.
206 *See id.*
207 *See generally* Michelle Clark Neely & Larry D. Sherrer, *A Look at Corporate Inversions, Inside and Out*, REGIONAL ECONOMIST, first quarter 2017, at 1.
208 *Cf.* Practical Law Bankruptcy & Restructuring and Practical Law Finance, *Fraudulent Conveyances: Issues and Strategies for Lenders and Private Equity Sponsors*, PRACTICAL LAW PRACTICE NOTE 8-382-2478 § 1 (2019) (discussing "downstream guarant[ies]"—where "the parent guarantees the debt of its subsidiary."). Some contracts contain similar "upstream" guaranties (i.e., a subsidiary guaranteeing the debts of its parent) and "cross-stream" guaranties (i.e., a subsidiary guaranteeing the debts of its sibling corporation). *See id.* However, such guarantees are more difficult to enforce given the guarantor's lack of control over the debtor. *See id.* Beyond these specific guarantees, other contractual protections—based on common corporate creditor protections—are possible. *See generally id.*; Practical Law Finance, *Loan Agreement: Affirmative Covenants*, PRACTICAL LAW PRACTICE NOTE 6-382-8184 (2019); Practical Law Finance, *Loan Agreement: Negative Covenants*, PRACTICAL LAW PRACTICE NOTE 5-383-3077 (2019).
209 *See* FIN. ACCOUNTING STANDARDS BD., STATEMENT OF FINANCIAL ACCOUNTING STANDARDS NO. 94 (1987), https://www.fasb.org/jsp/FASB/Document_C/DocumentPage?cid=1218220128231.
210 *See generally* UNIF. VOIDABLE TRANSACTIONS ACT § 4(a) (UNIF. LAW COMM'N 2014) ("A transfer made or obligation incurred by a debtor is voidable as to a creditor, whether the creditor's claim arose before or after the transfer was made or the obligation was incurred, if the debtor made the transfer or incurred the obligation: (1) with actual intent to hinder, delay, or defraud any creditor of the debtor; or (2) without receiving a reasonably equivalent value in exchange for the transfer or obligation, and the debtor: (i) was engaged or was about to engage in a business or a transaction for which the remaining assets of the debtor were unreasonably small in relation to the business or transaction; or (ii) intended to incur, or believed or reasonably should have believed that the debtor would incur, debts beyond the debtor's ability to pay as they became due."). Note that, under the Uniform Voidable Transactions Act, transfer to an affiliated corporation (such as a parent, subsidiary, or sibling) may be suggestive of actual intent to defraud the creditor. *See id.* § 4(b).
211 *See generally* RESTATEMENT (SECOND) OF CONTRACTS § 205 (AM. LAW INST. 1981) ("Every contract imposes upon each party a duty of good faith and fair dealing in its performance and its enforcement.").
212 *See* DGCL §§ 170(a)(1), 173.
213 "[DGCL § 170] contains a narrow exception under which a corporation that does not have surplus may pay a dividend out of 'net profits' for the fiscal year in which the dividend is declared and the prior year. These so-called 'nimble dividends' are rarely paid." James D. Honaker & Eric S. Klinger-Wilensky, *Dividends, Redemptions, and Stock Purchases*, PRACTICAL LAW PRACTICE NOTE 1-519-2507 § 2 (2019). Nevertheless, this might provide a means to circumvent the Clause if unaddressed. The Clause should explicitly guard against circumvention by nimble dividends, perhaps by contractually specifying that a signatory must satisfy its Clause obligations before paying nimble dividends. Furthermore, use of nimble dividends to circumvent Clause obligations could trigger fraudulent conveyance liability. *See id.*
214 "'Surplus' is defined in [DGCL § 154] as, '[t]he excess, if any, at any given time, of the net assets of the corporation over the amount . . . determined to be capital.' 'Net assets' is defined by [DGCL § 154] to mean 'the amount by which total assets exceed total liabilities.' This definition of net assets is essentially a balance sheet test." *Id.*; *see also* 11 FLETCHER CYC. CORP. § 5340 (2018) ("The duty of a corporation to pay its debts is superior to the duty to pay dividends. Courts have generally held that indebtedness must be deducted as a liability in determining compliance with statutory equity insolvency and balance sheet limitations for the payment of dividends." (footnotes omitted)). *See generally* RICHARD A. BOOTH, FINANCING THE CORPORATION § 5:12 (2018).
215 *See supra* § I(B)(2)(iii).
216 *See supra* § I(B)(1)(ii).
217 *See supra* § I(B)(1)(i).
218 *See, e.g.*, Brin, *supra* note 6.
219 *Asilomar AI Principles*, *supra* note 6.
220 *See id.*
221 *See supra* § I(B).
222 *See supra* § IV(A)(1).
223 *See supra* § I(B)(1)(iii).
224 *See supra* § I(B)(1)(iv).
225 White-collar workers are more likely to report that they feel "'positive' or 'mostly positive' about the impact of [AI] on people's lives and work over the next 10 years" than blue-collar workers. NORTHEASTERN U. & GALLUP, *supra* note 26, at 24. They also feel more prepared to find a new job if they lose theirs to AI. *See id.* at 25. A similar gap in both optimism and perceived preparedness exists between college-educated and non-college educated Americans. *See id.* at 26–27.
Zhang and Dafoe found similar results: groups that show greater support development of AI capable of displacing almost all human labor (high-income Americans, men, and Americans with tech experience) tend to be better positioned than their counterparts. *See* BAOBAO ZHANG & ALLAN DAFOE, ARTIFICIAL INTELLIGENCE: AMERICAN ATTITUDES AND TRENDS 36–37 (2019), https://perma.cc/X3G6-3NN4.
226 *See supra* § IV(A)(1).
227 *See* Hilary Greaves, *Discounting Future Health*, *in* GLOBAL HEALTH PRIORITY-SETTING: COST-EFFECTIVENESS AND BEYOND (Ole Frithjof Norheim et al. eds, forthcoming), https://perma.cc/F9YG-4J99.
228 *See supra* § II(A).
229 *But see* Robert E. Hall & Susan E. Woodward, *The Incentives to Start New Companies: Evidence from Venture Capital* 19 tbl.4 (Nat'l Bureau Econ. Research, Working Paper No. 13056, Apr. 2007), https://perma.cc/KUK8-Y5RQ; Seth Levine, *Venture Outcomes are Even More Skewed Than You Think*, VCADVENTURE (Aug. 12, 2014), https://perma.cc/Q4CQ-972G; Jerry Neumann, *Power Laws in Venture*, REACTION WHEEL (June 25, 2015), https://perma.cc/P662-74EJ.
230 *See supra* § II(A)(2).
231 Furthermore, there is some evidence that we could increase the personal income tax rate at the highest brackets without seriously adversely affecting economic growth. *See generally* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 227.
232 *See supra* §§ I(B)(2), III(A).
233 *See supra* § I(B)(2)(i).
234 *See supra* § I(B)(2)(ii).
235 *See supra* § I(B)(2)(iii).
236 *See, e.g.*, BOSTROM, *supra* note 8; MILES BRUNDAGE ET AL., THE MALICIOUS USE OF ARTIFICIAL INTELLIGENCE (2018), https://perma.cc/JZ4Y-PA7L; *Potential Risks from Advanced Artificial Intelligence*, OPEN PHILANTHROPY PROJECT (Aug. 2015), https://perma.cc/NV54-CDWK.
237 *See* Bostrom et al., *supra* note 11, at 9; *cf.* Sebastian Farquhar et al., *Pricing Externalities to Balance Public Risks and Benefits of Research*, 15 HEALTH SEC. 401 (2017), https://perma.cc/WUK3-554K (proposing internalization of risks from biotechnology research).
238 For background on racing dynamics in AI development, see Stephen Cave & Seán Ó hÉigeartaigh, *An AI Race for Strategic Advantage: Rhetoric and Risks*, 2018 PROC. OF THE AAAI/ACM CONF. ON AI, ETHICS, & SOC'Y 36, https://doi.org/10.1145/3278721.3278780; Edward Moore Geist, *It's Already Too Late to Stop the AI Arms Race—We Must Manage It Instead*, 72 BULLETIN OF THE ATOMIC SCIENTISTS 318 (2016); The Anh Han et al.,

Modelling and Influencing the AI Bidding War: A Research Agenda, 2019 Proc. of the AAAI/ACM Conf. on AI, Ethics, & Soc'y 5; Peter Apps, *Commentary: Are China, Russia Winning the AI Arms Race?*, Reuters (Jan. 15, 2019, 6:06 AM), https://perma.cc/RN9M-XFJQ; *Autonomous Weapons: An Open Letter From AI & Robotics Researchers*, Future of Life Inst., https://perma.cc/7Z6E-89NR (archived Aug. 28, 2019); Julian E. Barnes & Josh Chin, *The New Arms Race in AI*, Wall St. J. (Mar. 2, 2018, 11:47 AM); Paul Scharre, *Killer Apps*, Foreign Affairs, May/June 2019; Stuart Armstrong et al., *Racing to the Precipice: A Model of Artificial Intelligence Development* (Future of Humanity Inst., Technical Rept. No. 2013-1), https://perma.cc/S56B-Z5UZ; Amanda Askell et al., The Role of Cooperation in Responsible AI Development (July 10, 2019) (unpublished manuscript), https://perma.cc/WYR4-LM96; *OpenAI Charter*, OpenAI (Apr. 9, 2018), https://perma.cc/J29Y-2PU8.

For criticism of the AI arms race narrative, see, e.g., Heather M. Roff, *The Frame Problem: The AI "Arms Race" Isn't One*, 75 Bulletin of the Atomic Scientists 95 (2019); Justin Sherman, *Stop Calling Artificial Intelligence Research an 'Arms Race'*, Wash. Post (Mar. 6, 2019), https://perma.cc/X8S2-M7HR?type=image; Remco Zwetsloot et al., *Beyond the AI Arms Race*, Foreign Affairs (Nov. 16, 2018).

239 Belfield, *supra* note 35, at 2.

240 *See id.* at 5.

241 *See id.*

242 *See id.* at 6.

243 *See generally supra* § III(A).

244 *See, e.g.*, Jobin et al., *supra* note 7; *Asilomar AI Principles*, *supra* note 6; Brin, *supra* note 6; Catherine Clifford, *Billionaire Bill Gates on the Impact of A.I.: 'Certainly' We Can Look Forward to Longer Vacations*, CNBC (Jan. 29, 2018, 1:25 PM), https://perma.cc/AQ5L-A9M3; *Open Letter: Research Priorities for Robust and Beneficial Artificial Intelligence*, Future of Life Inst., https://perma.cc/7SRX-NY3Y (archived Sept. 30, 2019).

245 For example, politicians of both main American political parties frequently talk about wanting to support the development of large companies. *See, e.g.,* Lisa Mascaro, *JOBS Bill Clears Congress Despite Warnings*, L.A. Times (Mar. 27, 2012), https://perma.cc/MEZ7-NWYT ("Rep. Spencer Bachus (R-Ala.), the chairman of the Financial Services Committee, . . . said the legislation could help foster the next Google or Amazon.com."); *Jobs & The Economy*, Scott Peters for Congress (2018), https://perma.cc/WCR6-EQEP ("We want the next Google or Qualcomm or the next medical breakthrough to happen in [our district] . . . ."). *But cf.* Astead W. Herndon, *Elizabeth Warren Proposes Breaking Up Tech Giants Like Amazon and Facebook*, N.Y. Times (Mar. 8, 2019), https://perma.cc/8C42-8LGD. This suggests that many politicians want to be supportive of the development of such large companies, perhaps conditional on those companies acting responsibly. The Clause accomplishes exactly this.

246 *See generally* Uzma Khan & Ravi Dhar, *Licensing Effect in Consumer Choice*, 43 J. Marketing Res. 259 (2006); Anna C. Merritt et al., *Moral Self-Licensing: When Being Good Frees Us to Be Bad*, 4 Soc. & Personality Psychol. Compass 344 (2010); Benoît Monin & Dale T. Miller, *Moral Credentials and the Expression of Prejudice*, 81 J. Personality & Soc. Psychol. 33 (2001); Sonya Sachdeva et al., *Sinning Saints and Saintly Sinners: The Paradox of Moral Self-Regulation*, 20 Psychol. Sci. 523 (2009).

247 This resemblance is also why the Clause will not harm consumers through higher prices—corporate income taxes are ultimately paid by shareholders and (to a lesser degree) labor, not consumers. *See, e.g.*, Richard Goode, *The Corporate Income Tax and the Price Level*, 35 Am. Econ. Rev. 40 (1945); Juan Carlos Suárez Serrato & Owen Zidar, *Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms*, 106 Am. Econ. Rev. 2582 (2016); Bruce Bartlett, *Who Pays the Corporate Income Tax*, N.Y. Times: Economix (Feb. 19, 2013), https://perma.cc/6YYC-MLRB.

248 *See, e.g.*, Megan McArdle, *Here's a Good Corporate Tax Rate: Zero*, Wash. Post: Post Partisan (July 31, 2018), https://perma.cc/23PS-THAH; *see also* Robert E. Lucas, Jr., *Supply-Side Economics: An Analytical Review*, 42 Oxford Econ. Papers 293 (1990).

249 *See* McArdle, *supra* note 248.

250 *See id.*

251 *See supra* note 247.

252 *See* Belfield, *supra* note 35, at 4; *cf.* David Callahan, The Givers: Wealth, Power, and Philanthropy in a New Gilded Age (2017); Anand Giridharadas, Winners Take All: The Elite Charade of Changing the World (2018); Sachs & Kotlikoff, *supra* note 6; Ryan Pevnick, *Philanthropy and Democratic Ideals*, in Philanthropy in Democratic Societies 226 (Rob Reich et al. ed., 2016); Ed Burmila, *Jeff Bezos, Amazon and Why 'Charity' Is the Wrong Solution*, Rolling Stone (Feb. 14, 2018), https://perma.cc/KS5V-LTA7; Annie Lowrey, *Jeff Bezos's $150 Billion Fortune Is a Policy Failure*, Atlantic (Aug. 1, 2018), https://perma.cc/GSD5-MLWW; Rob Reich, *What Are Foundations For?*, Bos. Rev.: Forum (Mar. 1, 2013), https://perma.cc/9HB8-26V9.

253 *See generally supra* note 252 (collecting sources).

254 *See supra* § I(B)(2).

255 Indeed, large companies sometimes lobby heavily against taxation efforts. *See, e.g.,* Maya Kosoff, *Amazon Crushes a Small Tax That Would Have Helped The Homeless*, Vanity Fair (June 12, 2018), https://perma.cc/T8KP-NDC8.

256 *See Your Dollar Goes Further Overseas*, GiveWell, https://perma.cc/GNP3-4RHJ (archived Apr. 16, 2019).

257 *See* Curt Tarnoff & Marian L. Lawson, Cong. Res. Serv., R40213, Foreign Aid: An Introduction To U.S. Programs And Policy 18 (2018), https://perma.cc/245Z-GHJ4 (foreign aid spending accounted for 1.2% of the US's 2016 budget).

258 *Cf.* MacAskill, *supra* note 106, at 62 ("In the United States, public health experts regard any program that provides one [quality-adjusted life year (QALY)] for less than $50,000 as a good value, and health programs will often be funded even if the cost per QALY is much higher than $50,000. In contrast, providing the same benefit in poor countries [through philanthropy] can cost as little as one hundred dollars.").

259 *Cf.* Max Roser, *Global Income Inequality*, Our World in Data § I(4) (Oct. 2016), https://perma.cc/BX5Y-74BZ ("[T]he inequality of incomes between different countries is much higher than the inequality within countries.").

260 Despite the fact that it is much cheaper to save or improve lives in the developing world, *see id.*; *see also Your Dollar Goes Further Overseas*, GiveWell: Giving 101, https://perma.cc/GNP3-4RHJ (archived Apr. 16, 2019), the US spends very little of its budget on humanitarian foreign aid, *see* Tarnoff & Lawson, *supra* note 257, at 18.

261 *See* Belfield, *supra* note 35, at 4–5.

262 *See id.* at 4.

263 *Accord id.* at 4–5.

264 Verizon Commc'ns Inc. v. Law Offices of Curtis V. Trinko, LLP, 540 U.S. 398, 407 (2004); *see also* United States v. Grinnell Corp., 384 U.S. 563, 570–71 (1966) ("The offense of monopol[ization] . . . has two elements: (1) the possession of monopoly power in the relevant market and (2) the willful acquisition or maintenance of that power as distinguished from growth or development as a consequence of a superior product, business acumen, or historic accident.").

265 This current state of law is being contested. *See, e.g.*, Khan, *supra* note 77; Herndon, *supra* note 245.

266 *See supra* § I(B)(1)(ii).

267 *See* Belfield, *supra* note 35, at 4–5.

268 *See supra* note 35.

269 *See supra* note 35 (explaining how an AI monopoly would cause both deadweight loss and transfer of consumer surplus to producers, and how the Windfall Clause would address the latter but not the former).

270 E-mail from Anthony Aguirre, Associate Professor of Physics, Univ. of Cal., Santa Cruz, to Allan Dafoe, Director, Ctr. for the Governance of AI, Future of Humanity Inst., Univ. of Oxford (Jan. 8, 2019, 15:47 EST) (on file with author). Similar proposals can be found in Brynjolfsson & McAfee, *supra* note 6, at 246 (reporting, without attribution, a proposal to "[c]reate a national mutual fund distributing the ownership of [all] capital widely and perhaps inalienably, providing a dividend stream to all citizens and assuring that capital returns do not become highly concentrated"); Ford, *supra* note 6, at 273–75; Kaplan, *supra* note 6, at 178–82; Daniel C. Dennett, *Will AI Achieve Consciousness? Wrong Question*, Wired (Feb. 19., 2019, 7:00 AM), https://perma.cc/QM8Y-3X8W ("[W]e should enact legislation that puts [tech entrepreneurs'] deep pockets in escrow for the public good."); Rana Foroohar, *California's New Model for Fixing Inequality*, Financial Times (Oct. 13, 2019), https://app.ft.com/content/04830250-eb45-11e9-a240-3b065ef5fc55; Noah Smith, *The End of Labor: How to Protect Workers From the Rise of Robots*, Atlantic (Jan. 14, 2013), https://perma.cc/SU7B-QGNQ ("Everyone is born with an endowment of labor; why not also an endowment of capital? What if, when each citizen turns 18, the government bought him or her a diversified portfolio of equity? . . . This portfolio of capital ownership would act as an insurance policy for each human worker; if technological improvements reduced the value of that person's labor, he or she would reap compensating benefits through increased dividends and capital gains.").

271 *See supra* § IV(B).

272 *See* E-mail from Anthony Aguirre, *supra* note 270.

273 *See id.* The agreement establishing the SAF would also include some mechanism for aggregating constituents' votes. *See id.*

274 *See id.*

275 *See* E-mail from Anthony Aguirre, Associate Professor of Physics, Univ. of Cal., Santa Cruz, to Cullen O'Keefe, Research Affiliate, Ctr. for the Governance of AI, Future of Humanity Inst., Univ. of Oxford (Jan. 9, 2019, 11:54 EST) (on file with author).

276 *See id.*

277 *See id.*

278 *Id.*

279 *See id.*

280 *See generally* Practical Law Corporate & Securities & Scott Towers, *Poison Pills: Defending Against Takeovers/Stockholder Activism and Protecting NOLs*, Practical Law Practice Note 3-386-0340 (2019).

281 *See id.* § 2.

282 *Cf.* 15 U.S.C. § 18 ("No person engaged in commerce or in any activity affecting commerce shall acquire, directly or indirectly, the whole or any part of the stock or other share capital and no person subject to the jurisdiction of the Federal Trade Commission shall acquire the whole or any part of the assets of another person engaged also in commerce or in any activity affecting commerce, where in any line of commerce or in any activity

affecting commerce in any section of the country, the effect of such acquisition may be substantially to lessen competition, or to tend to create a monopoly."). *See generally* Einer Elhauge, *Horizontal Shareholding*, 129 HARV. L. REV. 1267 (2016).

283 *See* Mathew Ingram, *At Alphabet, There Are Only Two Shareholders Who Matter*, FORTUNE (June 7, 2017), https://perma.cc/QB5S-7ATG.

284 *See* E-mail from Anthony Aguirre, *supra* note 270.

285 The British Labour Party is proposing forcibly transferring shares of thousands of large companies to workers. *See* Jim Pickard, *UK's Labour Would Seize £300bn of Company Shares*, FINANCIAL TIMES (Sept. 1, 2019), https://www.ft.com/content/dc17d7ee-ccab-11e9-b018-ca4456540ea6.

286 *See id.*

287 *Cf.* U.S. CONST. amend. V ("[N]or shall private property be taken for public use, without just compensation.").

288 A number of sources have suggested UBI as a solution to AI-induced unemployment. *E.g.,* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 232–33; *The Freedom Dividend*, YANG 2020, https://perma.cc/3M76-7ARZ (archived Feb. 2, 2019). Other sources have endorsed a similar program—the negative income tax—which incentivizes work more than UBI. *E.g.,* BRYNJOLFSSON & MCAFEE, *supra* note 6, at 234–41; RIFKIN, *supra* note 6, at 261. For criticisms of UBI, *see, e.g.,* Furman, *supra* note 25, at 8–9; Daron Acemoglu, *Why Universal Basic Income Is a Bad Idea*, PROJECT SYNDICATE (June 7, 2019), https://perma.cc/S89F-HXBP.

289 For transparency, note that GiveWell was formerly affiliated with the Open Philanthropy Project, see Catherine Hollander, *Separating GiveWell and the Open Philanthropy Project*, GIVEWELL: THE GIVEWELL BLOG (July 6, 2017), https://perma.cc/6HSE-V86Q, which provided some funding for this work, *see supra* note 106.

290 *See Cash Transfers*, GIVEWELL (Nov. 2018), https://perma.cc/K8M3-AJ3F.

291 *See* Catherine Hollander, *Our Updated Top Charities for Giving Season 2018*, GIVEWELL: THE GIVEWELL BLOG (Dec. 6, 2018), https://perma.cc/9SNE-QEBN.

292 *Cf.* Vitalik Buterin et al., Liberal Radicalism: Formal Rules for a Society Neutral among Communities 5–6 (Sept. 17, 2018) (unpublished manuscript), https://perma.cc/H65U-FZAV.

293 *See id.*

294 *See* Dylan Matthews, *Basic Income Can't Do Enough to Help Workers Displaced by Technology*, VOX (Oct. 18, 2019, 8:20 AM), https://perma.cc/R8AH-87BA.

295 *See infra* app. II.

296 That the Clause would arguably involve exchange of unequal values does not render the Clause unsupported by consideration as a matter of contract law. *See* RESTATEMENT (SECOND) OF CONTRACTS § 79(c) (AM. LAW INST. 1981).

297 *Id.* § 90. *See generally Lack Of Consideration As Barring Enforcement Of Promise To Make Charitable Contribution Or Subscription—Modern Cases*, 86 A.L.R. 4th 241 (originally published in 1991); *Comment Note—Promissory Estoppel*, 48 A.L.R. 2d 1069 § 7 (originally published in 1956); *Consideration For Subscription Agreements*, 151 A.L.R. 1238 (originally published in 1944); *Consideration For Subscription Agreements*, 115 A.L.R. 589 (originally published in 1938); *Consideration For Subscription Agreements*, 95 A.L.R. 1305 (originally published in 1935); *Consideration For Subscription Agreements*, 38 A.L.R. 868 (originally published in 1925).

298 RESTATEMENT (SECOND) OF CONTRACTS, *supra* note 296, § 90 cmt. f.

299 *Id.* illus. 16.

300 *Id.* illus. 17.

301 It is worth noting, however, that "relief [in such cases] may *sometimes* be limited to restitution or to damages or specific relief measured by the extent of the promisee's reliance rather than by the terms of the promise." RESTATEMENT (SECOND) OF CONTRACTS, *supra* note 296, § 90 cmt. d. (emphasis added). Still, such limitations seem primarily aimed at avoiding unjust enrichment of the promisee, *see id.*, which would not be applicable here. Furthermore, cases applying this doctrine often enforce the full terms of the contract, not merely award reliance damages. *See, e.g.*, Los Angeles Traction Co. v. Wilshire, 67 P. 1086 (Cal. 1902) (promise for $2000 was fully enforced despite reliance of only $1505); Lasar v. Johnson, 58 P. 161, 163 (1899) ("[Reliance] was a good consideration for the promise to pay the full sum sued for, and it is therefore not necessary to consider whether the defendants' liability should be reduced to its proportion of the expense incurred in entertaining the grand parlor measured by the entire amount subscribed; the amount subscribed exceeding the expenditures by the sum of $241.52, but the amount of subscriptions collected being $593.38 less than the amount expended.").

302 *See, e.g.*, *Wilshire*, 67 P. at 1088 ("The contract at the date of its making was unilateral, a mere offer that if subsequently accepted and acted upon by the other party to it would ripen into a binding enforceable obligation. When the respondent purchased and paid upwards of $1,500 for a franchise it had acted upon the contract, and it would be manifestly unjust thereafter to permit the offer that had been made to be withdrawn."); *Lasar*, 58 P. at 162 ("The request of the defendants to the committee to regard their subscription as cash, and to go on and make all their arrangements, and the action of the committee in incurring obligations on the faith of this request, was a sufficient consideration to fix the liability of the defendants upon their subscription . . . ."); Grand Lodge of the Indep. Order of Good Templars of the State of Cal. v. Farnham, 11 P. 592, 592–93 (Cal. 1886); Christian Coll. v. Hendley, 49 Cal. 347, 349 (1874) ("If it clearly appears that a number of subscribers promise to contribute money, on the faith of the common engagement, for the accomplishment of an object of interest to all, and which cannot be accomplished, save by their common performance, then it would seem that the mutual promises constitute reciprocal obligations."); 1 B. E. WITKIN, SUMMARY 11TH CONTRACTS § 254 (2018).

303 *See, e.g.*, I. & I. Holding Corp. v. Gainsburg, 12 N.E.2d 532, 534 (N.Y. 1938) ("Our courts have definitely ruled that such subscriptions are enforceable on the ground that they constitute an offer of a unilateral contract which, when accepted by the charity by incurring liability in reliance thereon, becomes a binding obligation."); *In re* Taylor's Estate, 167 N.E. 434, 436 (N.Y. 1929) ("[I]f a thing is to be done by the promisee, be it ever so small, this is a sufficient consideration to sustain the promise. Was there a promise on the part of the churches, or either of them, to do something in return for which the notes were given?"); Allegheny Coll. v. Nat'l Chautauqua County Bank of Jamestown, 159 N.E. 173, 175 (N.Y. 1927) (finding consideration in implied return promise to maintain memorial fund endowed by donor); Keuka Coll. v. Ray, 60 N.E. 325 (N.Y. 1901) (finding consideration in founding a college by donor's request); Twenty-Third St. Baptist Church v. Cornwell, 23 N.E. 177, 177 (N.Y. 1890) (finding no consideration where subscriber died before any reliance); Barnes v. Perine, 12 N.Y. 18, 25 (1854) ("The evidence, however, discloses a good consideration, in the . . . obligations incurred by the promisee . . . at [the defendant's] request."); JOY T. CARMICHAEL ET AL., 22 N.Y. JUR. 2D CONTRACTS §§ 76, 83 (2018).

304 An injunction is a court order to do or refrain from doing some specific act. *See generally* William M. Lafferty & John P. DiTomo, *Provisional Remedies: Delaware*, PRACTICAL LAW STATE Q&A w-000-4035 (2017).

305 N. River Ins. Co. v. Mine Safety Appliances Co., 105 A.3d 369, 380 n.47 (Del. 2014), *as revised* (Nov. 10, 2014).

306 *See, e.g.*, Simon v. Pyrites Co., 128 A. 370, 371 (Del. Super. Ct. 1925).

307 *See, e.g.*, Pomilio v. Caserta, 215 A.2d 924, 925 (Del. 1965) ("There is no right to an injunction as a matter of course. . . . The granting or refusal of final injunctive relief rests within the sound judicial discretion of the trial court.").

308 *See supra* § II.

309 *See* Hollinger Int'l, Inc. v. Black, 844 A.2d 1022, 1090 (Del. Ch. 2004); Kohls v. Duthie, 765 A.2d 1274, 1289 (Del. Ch. 2000); State v. Delaware State Educ. Ass'n, 326 A.2d 868, 875 (Del. Ch. 1974).

310 *See, e.g.*, *In re* Cogent, Inc. S'holder Litig., 7 A.3d 487, 513 (Del. Ch. 2010); *Hollinger*, 844 A.2d at 1090; *In re* Staples, Inc. S'holders Litig., 792 A.2d 934, 960 (Del. Ch. 2001); Cantor Fitzgerald, L.P. v. Cantor, 724 A.2d 571, 586 (Del. Ch. 1998).

311 *See, e.g.*, Revlon, Inc. v. MacAndrews & Forbes Holdings, Inc., 506 A.2d 173, 184–85 (Del. 1986); *Hollinger*, 844 A.2d at 1090; ACE Ltd. v. Capital Re Corp., 747 A.2d 95, 110 (Del. Ch. 1999).

312 *See* Bayard v. Martin, 101 A.2d 329, 334 (Del. 1953); Cook v. Oberly, 459 A.2d 535, 540 (Del. Ch. 1983).

313 *See* Weldin Farms, Inc. v. Glassman, 414 A.2d 500, 505 (Del. 1980); *Oberly*, 459 A.2d at 540; Pierce Family, Inc. v. Magness Const. Co., 235 A.2d 268, 270 (Del. Ch. 1967); McQuail v. Shell Oil Co., 183 A.2d 581, 585 (Del. Ch. 1962); Gray Co. v. Alemite Corp., 174 A. 136, 142 (Del. Ch. 1934).

314 *See* Organovo Holdings, Inc. v. Dimitrov, 162 A.3d 102, 115 (Del. Ch. 2017); Troy Corp. v. Schoon, 959 A.2d 1130, 1138 (Del. Ch. 2008); State *ex rel.* Brady v. Pettinaro Enterprises, 870 A.2d 513, 536 (Del. Ch. 2005); Aquila, Inc. v. Quanta Servs., Inc., 805 A.2d 196, 208 (Del. Ch. 2002); Cantor, 724 A.2d at 586; McMahon v. New Castle Assocs., 532 A.2d 601, 606 (Del. Ch. 1987); Delaware State Educ. Ass'n, 326 A.2d at 876; *see also* New Castle Cty. Vocational Tech. Educ. Ass'n v. Bd. of Educ. of New Castle Cty. Vocational Tech. Sch. Dist., 451 A.2d 1156, 1160 (Del. Ch. 1982) ("Nevertheless, the threat of an illegal action . . . must be imminent to justify injunctive relief.").

315 E. I. duPont de Nemours & Co. v. Am. Potash & Chem. Corp., 200 A.2d 428, 432 (Del. Ch. 1964).

316 *duPont* is the only Delaware case mentioning this "well-recognized principle," and it cites no other cases for the proposition. *See id.*

317 *Cf.* Pell v. Kill, 135 A.3d 764, 793 (Del. Ch. 2016) (finding irreparable harm when board action disenfranchised shareholders in derogation of fiduciary obligations, solely on the basis of such disenfranchisement); *duPont*, 200 A.2d at 436 ("[W]here there is a finding that a threat of use or disclosure exists which the court concludes will, if effectuated, constitute a breach of confidence, the court will grant appropriate relief."). *But cf.* Sealy Mattress Co. of New Jersey v. Sealy, Inc., 532 A.2d 1324 (Del. Ch. 1987) (finding breach of fiduciary duty but analyzing irreparable harm separately).

318 *See, e.g.*, *In re* BankAtlantic Bancorp, Inc. Litig., 39 A.3d 824, 845 (Del. Ch. 2012); *In re* Del Monte Foods Co. S'holders Litig., 25 A.3d 813, 839 (Del. Ch. 2011); Abrons v. Maree, 911 A.2d 805, 810 (Del. Ch. 2006); Hollingsworth v. Szczesiak, 84 A.2d 816, 822 (Del. Ch. 1951).

319 *See* Theravectys SA v. Immune Design Corp., No. CV 9950-VCN, 2015 WL 1308273, at *12 (Del. Ch. Mar. 9, 2015); Mitchell Lane Publishers, Inc. v. Rasemas, No. CIV.A 9144-VCN, 2014 WL 4925150, at *10 (Del. Ch. Sept. 30, 2014); *Del Monte*, 25 A.3d at 839; Cantor, 724 A.2d at 587; Lennane v. Ask Computer Sys., Inc., No. CIV. A. 11744, 16 Del. J. Corp. L. 1521, 1534 (Del. Ch. Oct. 11, 1990); Newell Co. v. Wm. E. Wright Co., 500 A.2d 974, 975 (Del. Ch. 1985).

320 *See* Elster v. Am. Airlines, 106 A.2d 202, 204 (Del. Ch. 1954); Allied Chem. & Dye Corp. v. Steel & Tube Co. of Am., 122 A. 142, 157 (Del. Ch. 1923).

321 This does *not* imply that any plaintiff who succeeds in proving a breach of fiduciary duty will also prove irreparable harm. Injunctions only issue for future conduct, *see, e.g.*, Young v. Red Clay Consol. Sch. Dist., 159 A.3d 713, 781 (Del. Ch. 2017), so there must be some threat of future breach to warrant an injunction.

322 *See supra* § II.

323 *Tyson*, 919 A.2d at 584; *see also In re* Am. Int'l Grp., Inc., 965 A.2d 763, 812 (Del. Ch. 2009); Dover Assocs. Joint Venture v. Ingram, 768 A.2d 971, 973 (Del. Ch. 2000).
324 Delaware Tr. Co. v. Partial, 517 A.2d 259, 262 (Del. Ch. 1986).
325 *Ingram*, 768 A.2d at 974 (appointing receiver pursuant to plaintiff's contracted-for right thereto); *cf.* C & J Energy Servs., Inc. v. City of Miami Gen. Employees', 107 A.3d 1049, 1054 (Del. 2014) ("[A]n injunction cannot strip an innocent third party of its contractual rights while simultaneously binding that party to consummate the transaction."). Numerous cases demonstrate the respect equity affords to contractual rights. *See, e.g.*, Diebold Computer Leasing, Inc. v. Commercial Credit Corp., 267 A.2d 586, 590 (Del. 1970) ("The jurisdiction of the Chancery Court to enjoin a threatened breach of contract, for which breach damages would not be adequate, is unquestioned."); Tull v. Turek, 147 A.2d 658, 664 (Del. 1958) ("An historic field of equity jurisdiction is that of suits for an injunction to restrain the breach of the covenant of a contract. Such jurisdiction is by analogy as broad as the jurisdiction of equity to compel the specific performance of a contract."); Randolph v. Wilmington Hous. Auth., 139 A.2d 476, 489 (Del. 1958) ("The jurisdiction of equity to enjoin a breach of contract if the legal remedy is inadequate is too well-established to require discussion."); Richard Paul, Inc. v. Union Imp. Co., 91 A.2d 49, 55 (Del. 1952) ("[A] clear violation of plaintiff's rights under the lease has been established, without the establishment of any counter-balancing equities in the defendant's favor. Under such circumstances, the course of a Court of Equity is clear—it must follow the law and protect the plaintiff's legal rights.").
326 *See, e.g.*, Freedman v. Cont'l Serv. Corp., 622 P.2d 487, 492 (Ariz. Ct. App. 1980) ("As a sound general proposition, equity follows the law and a court will enforce a valid contract according to its terms, even though enforcement may be harsh or result in a forfeiture."); Preferred Prof'l Ins. Co. v. The Doctors Co., 419 P.3d 1020, 1027 (Colo. App. 2018) ("Absent a showing that a contractual provision violates public policy, equity should not be employed to defeat a party's bargained-for contractual rights."); Honolulu Brewing & Malting Co. v. Bartlett, 23 Haw. 192, 196 (1916) ("Equity follows the law, does not oppose the law; it assists parties in the protection of their legal rights when the law is unable to do so; it protects the vigilant, but does not invade private fundamental rights to impose under a contract obligations not assumed in such contract, either by express terms or by necessary implication."); Norcomo Corp. v. Franchi Const. Co., 587 S.W.2d 311, 317 (Mo. Ct. App. 1979) ("The rights of the parties are to be determined from the contracts into which they entered and the consequences of those contracts and not from some generalized concepts of equity."); McCall v. Carlson, 172 P.2d 171, 187–88 (Nev. 1946) ("Our equitable powers do not extend so far as to permit us to disregard fundamental principles of the law of contracts, or arbitrarily to force upon parties contractual obligations, terms or conditions which they have not voluntarily assumed. In this regard, equity respects and upholds the fundamental right of the individual to complete freedom to contract or decline to do so, as he conceives to be for his best interests, so long as his contract is not illegal or against public policy. In this respect, and many others, equity follows the law. Much as we would like to relieve the appellant from his unfortunate situation, we cannot rightfully do so, as we must maintain the necessary certainty, stability and integrity of contractual rights and obligations."); Dunkin' Donuts of Am., Inc. v. Middletown Donut Corp., 495 A.2d 66, 74 (N.J. 1985) ("[T]he equitable maxim 'equity follows the law,' . . . instructs that as a rule a court of equity will follow the legislative and common-law regulations of rights, and also obligations of contract."); ARC LifeMed, Inc. v. AMC-Tennessee, Inc., 183 S.W.3d 1, 26 (Tenn. Ct. App. 2005) ("What appears to be fair as a matter of equity cannot prevail in the face of . . . contractual provisions governing the rights of the parties."); Texas Health Ins. Risk Pool v. Sigmundik, 315 S.W.3d 12, 14 (Tex. 2010) ("[E]quitable doctrines conform to contractual and statutory mandates, not vice versa."); Fortis Benefits v. Cantu, 234 S.W.3d 642, 648–49 (Tex. 2007) ("We generally adhere to the maxim that 'equity follows the law,' which requires equitable doctrines to conform to contractual and statutory mandates, not the other way around. Where a valid contract prescribes particular remedies or imposes particular obligations, equity generally must yield unless the contract violates positive law or offends public policy."); Pac. Fin. Corp. v. Snohomish Cty., 295 P. 110, 112 (Wash. 1931) ("Equity . . . follows the law, and the courts have no authority on any equitable principle to rewrite the contract for the parties.").
327 41 A.3d 432 (Del. Ch. 2012).
328 *See id.* at 435.
329 *Id.* at 436–37 (citations omitted).
330 *See id.* at 435.
331 *See id.* at 444.
332 *See id.* at 449.
333 *See id.* at 450.
334 *Id.* at 449.
335 *Id.* at 450.
336 501 A.2d 1239.
337 *See id.* at 1245–46, 1251.
338 *See id.* at 1252.
339 *See id.* at 1251.
340 *Id.*
341 *See id.* at 1252.
342 Rescission of a contract is technically distinct from an injunction. It is only available in circumstances more extreme those warranting an injunction: [C]ommon-law rescission of a sale or a contract of sale has been allowed on broad principles of justice. Rescission is an equitable remedy available in contract cases when the resulting injury is irreparable, when damages at law are inadequate or when damages would be difficult or impossible to determine. Generally, and subject to statutory provisions, a contract may be rescinded by one party without the consent of the other on grounds such as fraud or mistake [of fact], lack of mutuality, or, in some cases, breach of contract or condition. On the other hand, neither party may alone, and without the consent of the other, rescind the contract arbitrarily, without adequate cause or for his or her own default. One may not unilaterally rescind a contract of sale because it is uneconomic or imprudent.
WILLIAM H. DANNE, JR. ET AL., 77A C.J.S. SALES § 187 (June 2018) (citations omitted). Given the extremity of the circumstances warranting rescission, and the unlikelihood of them arising in connection with a Windfall Clause, we do not consider rescission further. However, planners should be aware that rescission is theoretically possible.
343 No. 12883, 1995 WL 376919 (Del. Ch. June 15, 1995) [hereinafter Arnold III].
344 Arnold v. Soc'y for Sav. Bancorp, Inc., 650 A.2d 1270, 1273, 1277 (Del. 1994) [hereinafter Arnold II].
345 *See id.* at 1287–89.
346 Arnold III, 1995 WL 376919, at *4 (citations omitted).
347 Arnold v. Soc'y for Sav. Bancorp, Inc., No. C.A. 12883, 1993 WL 183698 (Del. Ch. May 29, 1993) [hereinafter Arnold I].
348 509 A.2d 584 (Del. Ch. 1986).
349 *See id.* at 587.
350 *See id.* at 586–87.
351 *See id.* at 590.
352 *Id.* at 591.
353 *See id.* at 592.
354 *Id.* at 600 (emphasis added). Interestingly, this suggests that the court thinks that Bally's contract rights had already vested, despite the fact that the merger had not yet closed.
355 *MGM*, 509 A.2d at 600.
356 *Cf. id.* at 597.
357 *Id.* at 600 n.12 (citations omitted).
358 *See, e.g.*, Reid v. Spazio, 970 A.2d 176, 183 (Del. 2009); Adams v. Jankouskas, 452 A.2d 148, 157 (Del. 1982).
359 Meso Scale Diagnostics, LLC v. Roche Diagnostics GmbH, 62 A.3d 62, 76 (Del. Ch. 2013) (quoting Scureman v. Judge, 626 A.2d 5, 13 (Del. Ch. 1992)).
360 *Id.* at 76–77 (quoting *Jankouskas*, 452 A.2d at 157).
361 *Jankouskas*, 452 A.2d at 157; *see also Spazio*, 452 A.2d at 183; Carsanaro v. Bloodhound Techs., Inc., 65 A.3d 618, 645 (Del. Ch. 2013). *See generally* DEL. CODE tit. 10, § 8106.
362 *Bloodhound*, 65 A.3d at 645; *see also* Wal-Mart Stores, Inc. v. AIG Life Ins. Co., 860 A.2d 312, 319 (Del. 2004).
363 Halpern v. Barran, 313 A.2d 139, 141 (Del. Ch. 1973).
364 *See Wal-Mart*, 860 A.2d at 319.
365 625 A.2d 269 (Del. Ch. 1993).
366 *See id.* at 270.
367 *See id.*
368 *See id.*
369 *See id.* at 270–71.
370 *Id.* at 271 (emphasis added) (footnote omitted).
371 *See supra* § I.
372 This structure is inspired by the Centre For Effective Altruism's EA Funds. *See generally Frequently Asked Questions*, EA FUNDS, https://perma.cc/X4MP-SCSQ (archived Apr. 16, 2019).
For transparency, we note that one such EA Fund—the Long-Term Future Fund—has funded the Berkeley Existential Risk Initiative, which is a publisher of this Report. *See* Nick Beckstead, *Payout Report: Long-Term Future Fund*, EA FUNDS (Mar. 20, 2017), https://perma.cc/Y2LX-STXT. We further note that the Centre for Effective Altruism and the other publisher of this report, the Future of Humanity Institute, share office space, some staff, and work together intimately on a number of projects.
373 One could further add mechanisms to incentivize managers to do their jobs effectively. Managers could earn bonuses based on how many

quality-adjusted life-years (QALYs) per dollar their funded projects achieved in the past two years, based off estimates from an independent auditor (which would also be funded by the Clause). For example, managers could earn a bonus for hitting ambitious QALY/dollar goals. Or, to incentivize maximizing QALYs/dollar, they could earn bonuses based on their performance relative to other managers. To avoid the problem of managers relying on a very small number of unusually effective projects to obtain a high average effectiveness (but undesirably withholding money otherwise), bonuses could be available only after managers had released a certain amount of funding, or comparisons could be made for each tranche of funding (e.g., Which manager spent her most-effective million dollars most effectively? Which manager spent her second-most-effective million dollars most effectively?) then aggregated.

Note that the incentive structure would need to account for the possibility that future funding opportunities might be more effective than present ones. *See generally* Julia Wise, *Giving Now vs. Giving Later: A Summary*, Effective Altruism Forum (July 23, 2013, 12:00 AM EDT), https://perma.cc/34XA-VUX5. So, for example, bonuses should go up over time if managers' grants improve over time (e.g., due to learning, increased expertise, or improved technology).

For more on existing social impact measurement and incentive structures, *see* Global Impact Investing Network, *supra* note 116; Kerr, *supra* note 145, at 642–54.

374 *Cf. supra* § II(B)(2) (promoting the "four-eyes principle").

375 *Cf.* Elster, *supra* note 111, at 109–11, 119–20 (discussing the value of anonymity in jury deliberations and mechanics).

376 Manne, *supra* note 117.

377 *See id.* at 240–52.

378 *See id.* at 252–64.

379 *See id.*

380 For a similar proposal by Anthony Aguirre, *see supra* § V(D)(3).

381 *See generally* Sitkoff & Dukeminier, *supra* note 115, at 624–54.

382 *See generally id.* at 587–693.

383 *See* AI for Good Global Summit, *supra* note 56.

384 U.N. Secretary-General, *supra* note 118.

385 For a list of the SDGs, see *Sustainable Development Goals*, Sustainable Dev. Goals Knowledge Platform, https://perma.cc/LR7L-KFJ9 (archived Jan. 29, 2019).

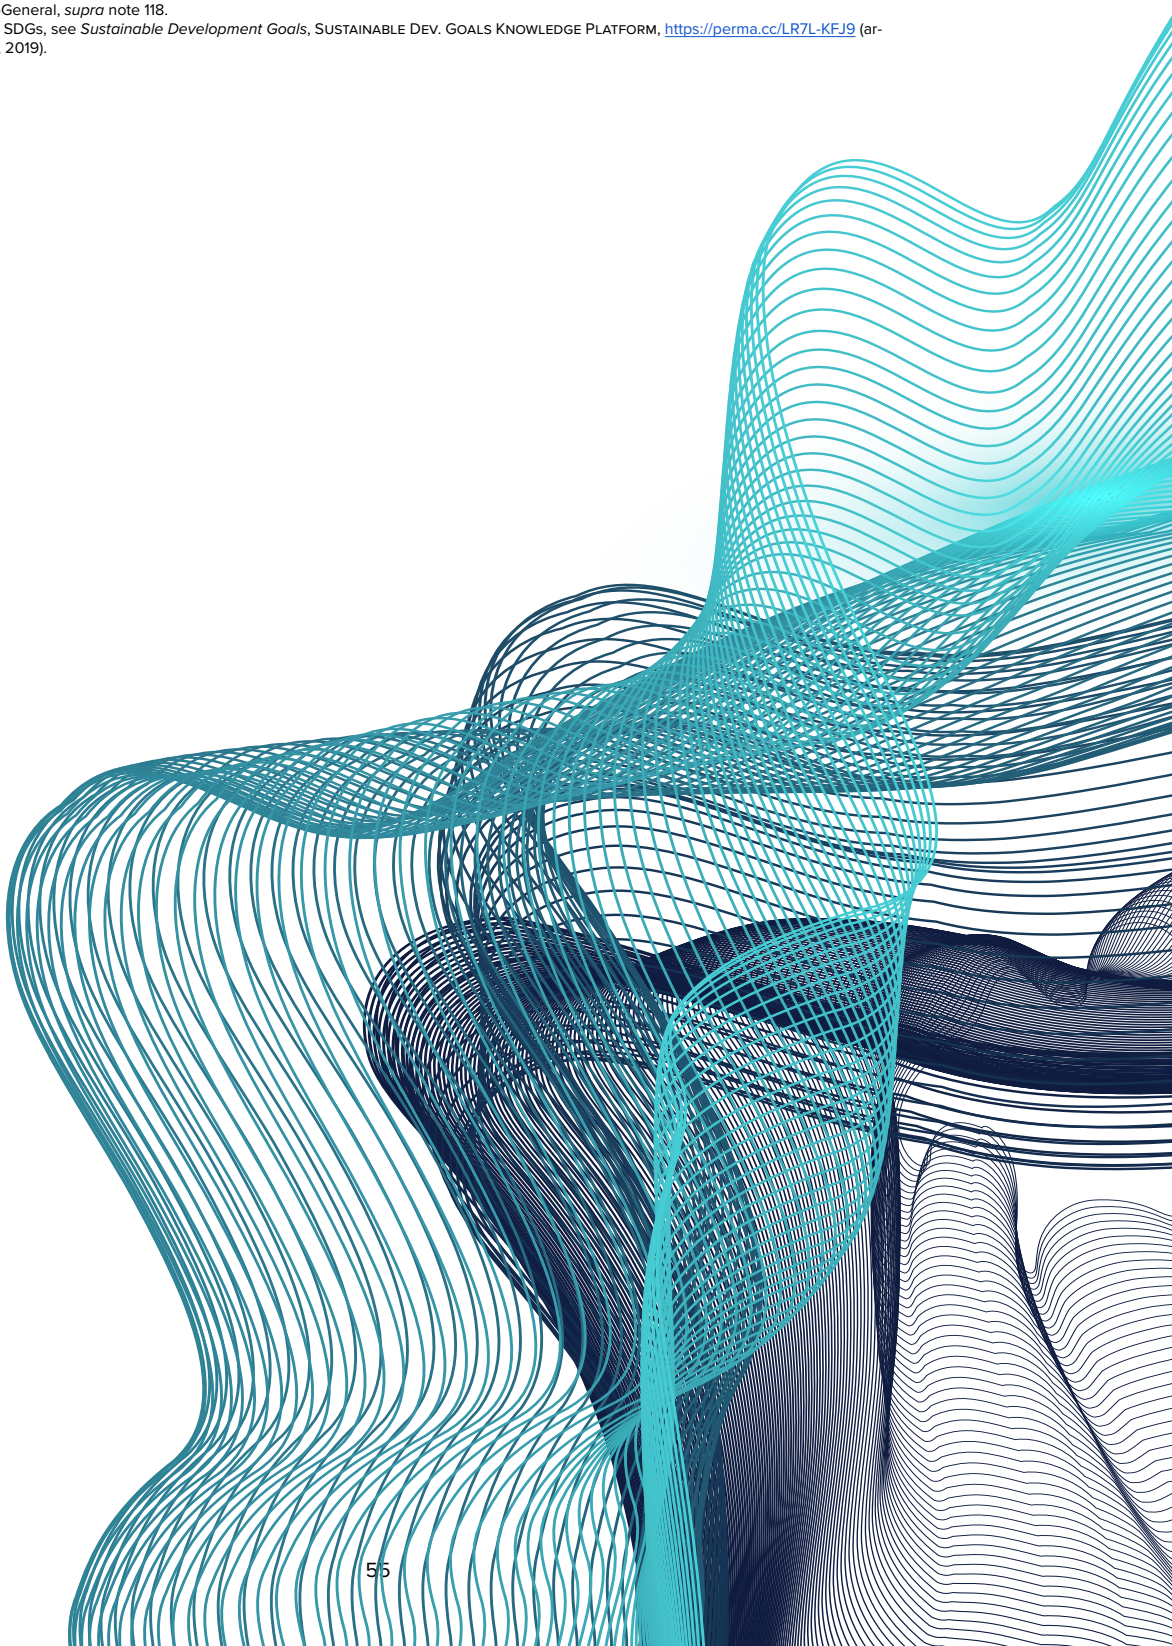

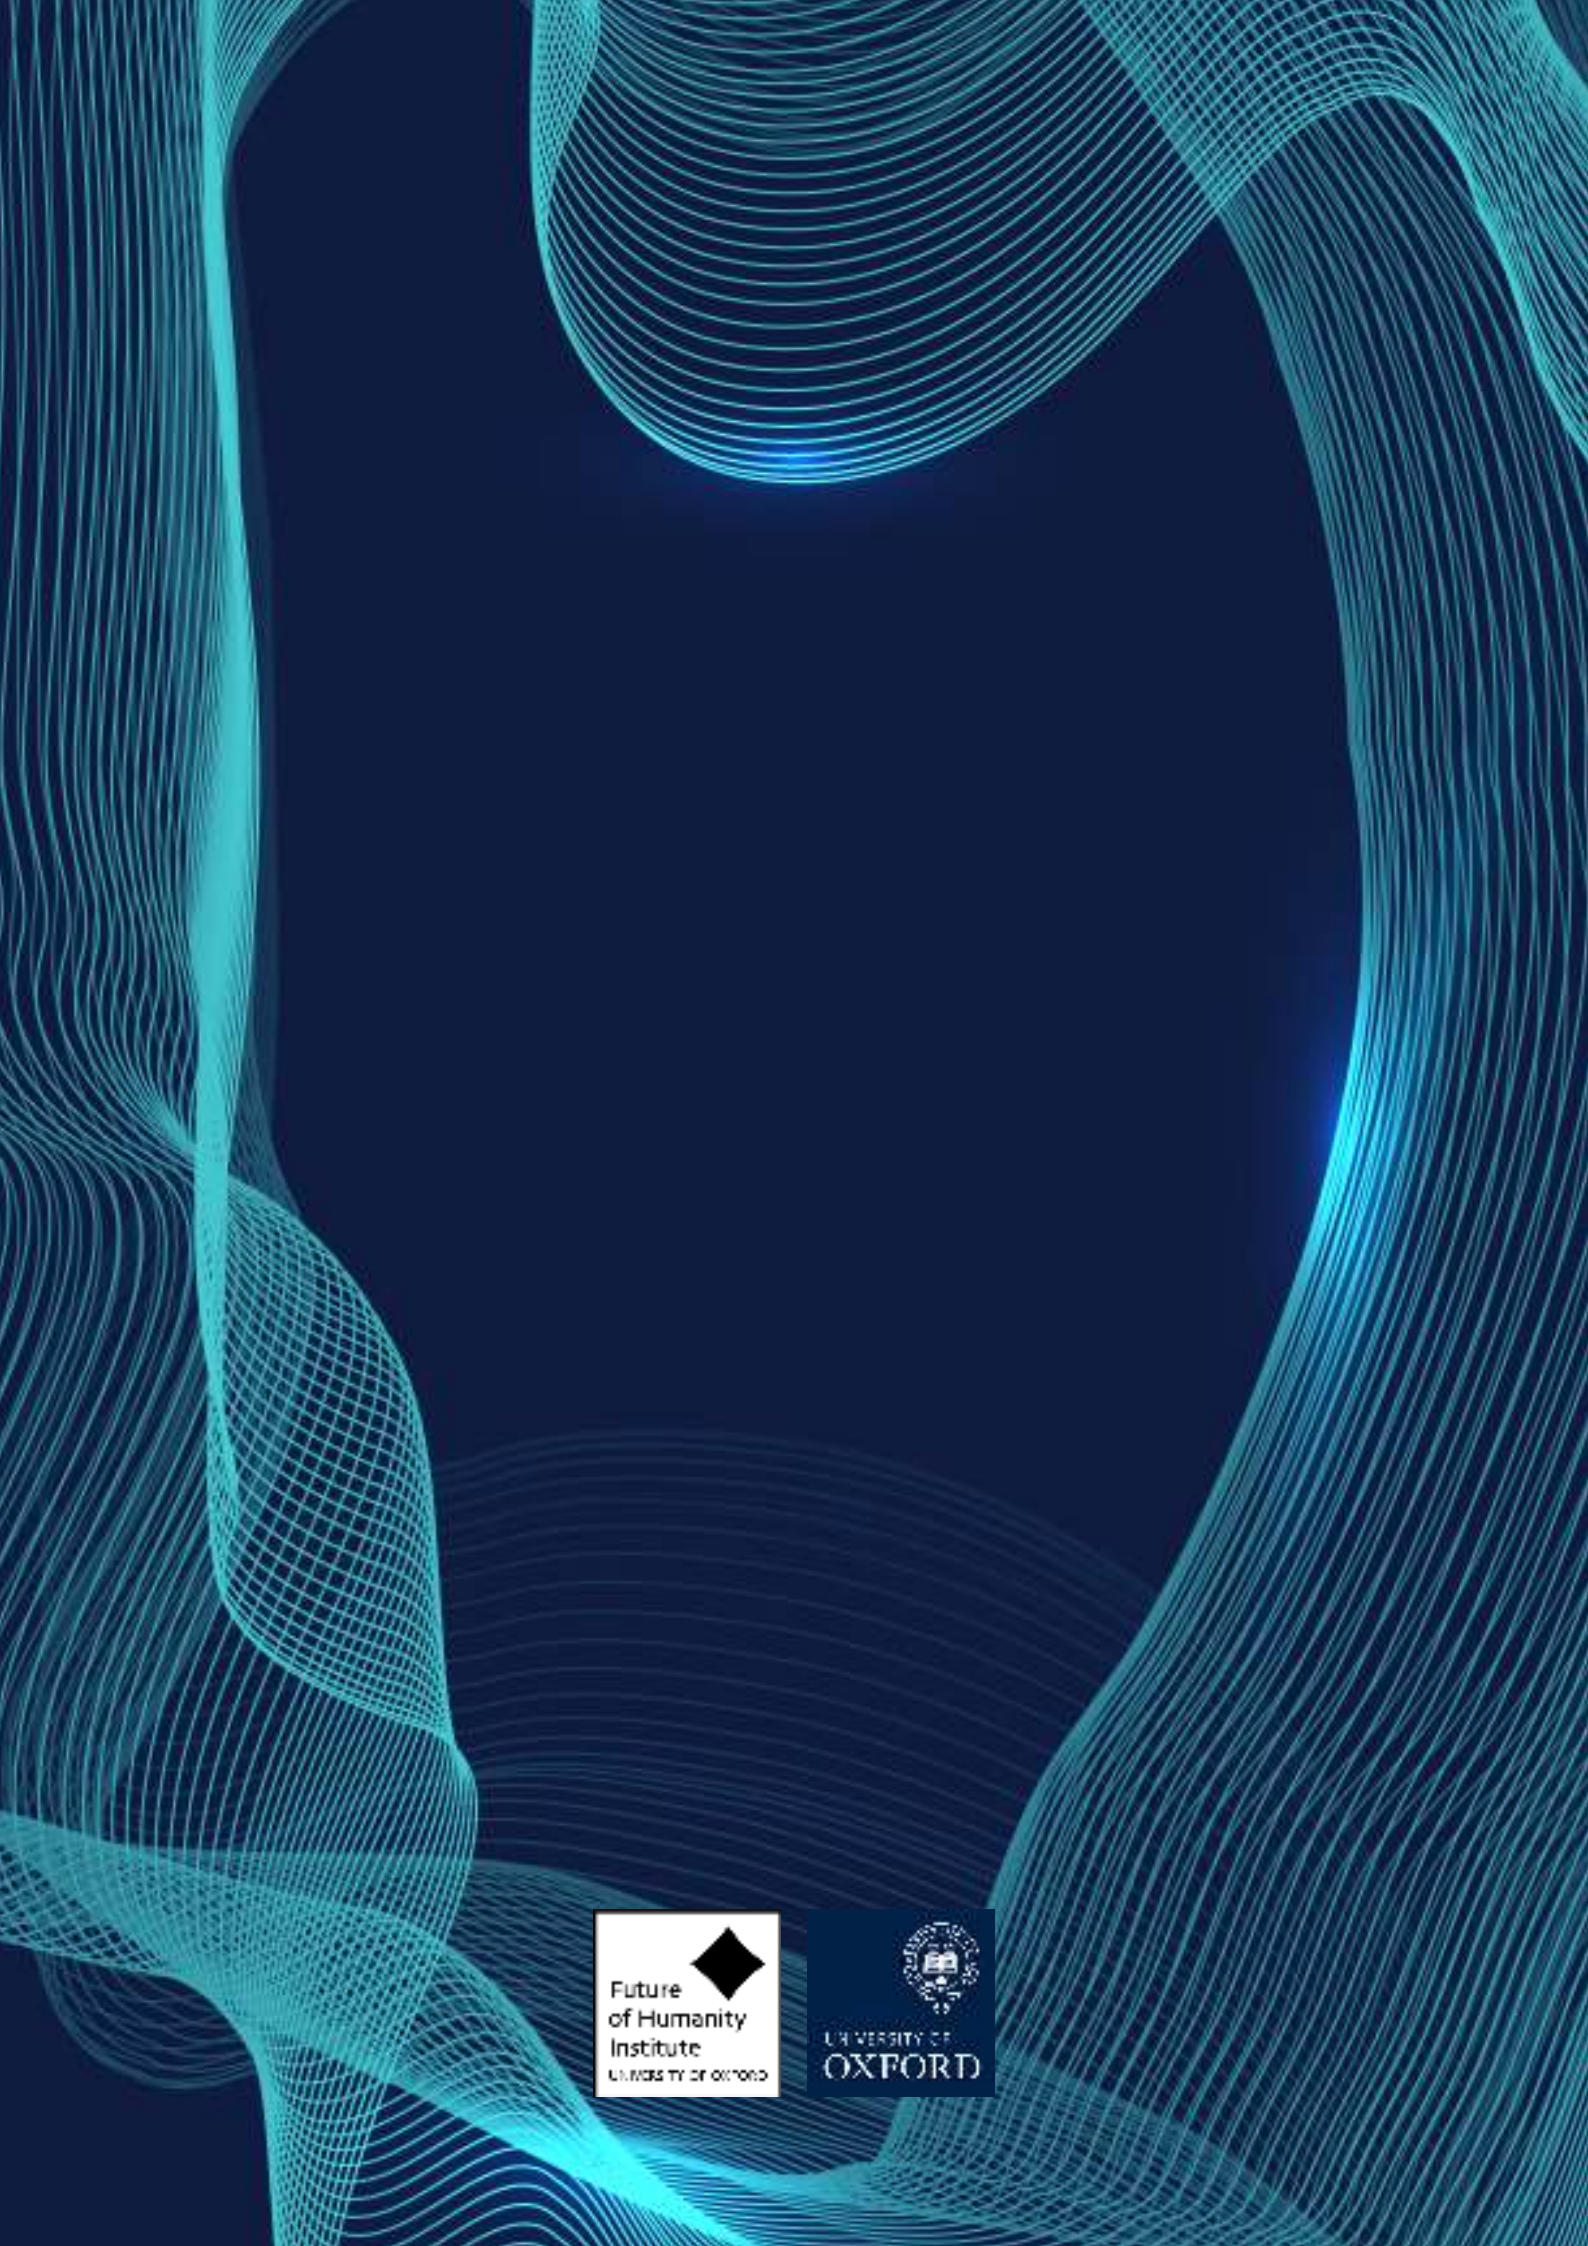
Future
of Humanity
Institute
UNIVERSITY OF OXFORD
UNIVERSITY OF
OXFORD