Are Tech Companies Responsible for Solving the Global AI Divide?

A Practical Exploration of Libertarian, Rawlsian and Utilitarian Points of View

Laura Rachevsky
Lucy Cavendish, University of Cambridge

The global AI divide, marked by the unequal distribution of AI benefits between developed and developing countries, is a pressing ethical concern. This paper examines the moral responsibility of tech companies in addressing this divide, analysing it through the lenses of libertarianism, Rawlsianism, and utilitarianism. It delves into the nuances of each perspective, particularly highlighting their limitations in a global context, and contrasts the current focus on productivity-enhancing AI applications in developed countries with the potential of life-saving AI applications in developing countries. The paper explores empirical examples of tech companies' investments in developing countries, revealing that libertarian and Rawlsian perspectives, despite initial differences, converge in their practical implications on a global scale. Ultimately, it argues that utilitarianism, although not without its challenges, provides the most actionable framework for addressing the global AI divide due to its emphasis on measurable outcomes and its ability to transcend national boundaries. It further performs a simplistic redistribution calculation as a proof of concept to demonstrate how incorporating life-saving AI applications into the benefits calculation can result in different investment recommendations.

Keywords: Global AI Divide, Philosophy, Corporate Responsibility, Global Justice

Introduction

It is frequently noted that one of the issues in contemporary AI ethics is the "AI divide," or the unequal distribution of benefits produced through use of AI technology between the developed and developing countries. Multiple authors point out the fact that "the economic and social benefits of AI remain geographically concentrated, primarily in the Global North" (World Economic Forum, 2023) and some, such as Yuval Harari even go as far as questioning "will the rest of the world just become algorithmic data colonies for AI-dominating countries?" (LSE, 2023).

This divide also manifests itself empirically. A recent report by PwC, "Sizing the prize," seeks to size the benefits most prominent applications of AI would bring as measured by GDP gains. Use cases they consider mainly include productivity enhancing AI use cases, such as driverless cars and trucks, and scaled financial advice and customisation. Based on this their estimate of $15.7 trillion of GDP gains by 2030 is currently split by 79% of the benefit going to developed countries and 21% to developing countries - a stark contrast with population distribution as illustrated in Table 1 below.

Table 1: *Population distribution vs AI Benefits distribution in 2030*

	Population	GDP gains under current applications ($, bn)	Population distribution	Benefits distribution
Developed countries*	1368m	12,500	17%	57%
Developing countries	6632m	9,300	83%	43%
	8000m	**21,800**	**100%**	**100%**

Note: the GDP gains column of the table is sourced from the "Pwc-Ai-Analysis-Sizing-the-Prize-Report", 2017; UN definition of developing and developed countries used (UNCTDA, 2022)

It is rarely debated who is responsible for alleviating the global AI divide. Rare solutions offered emphasise the need for collaboration among governments, multilateral agencies and technology providers without explicitly attributing responsibility to either party.

Given that the vast majority of investment in AI is currently done by private companies, in my paper I explore whether it is the responsibility of companies developing AI to ensure that the benefits of their technology extend to people in the developing countries. I examine this issue from libertarian, Rawlsian and utilitarian points of view in turn. I first explain each stance's likely approach to the debate and delve into nuances within each that highlight the ambiguity of potential conclusions. For each school of thought I choose an example of tech companies' existing investments in developing countries where the ethos behind the investment seems to have been informed by this particular school. I then evaluate the pros and cons of each theory along two dimensions: (1) their usefulness in addressing the issue in a global context as opposed to that of a nation state as this transcendence of national borders is paramount in addressing issue of global AI divide and (2) the plausibility of real implementation of the practical recommendations offered by each theory as apart from providing the most useful methodological framework the best approach should also be judged on its impact.

I choose the above three schools of thought to analyse the issue as they present seemingly radically different answers to the debate in question from denying responsibility of tech companies (libertarianism) to advocating for it (Rawlsian) with utilitarianism falling somewhere in the middle depending on the calculations involved. However, when delving into nuance, it becomes apparent that libertarian and Rawlsian points of view despite initially offering contrarian points of view actually conflate when it relates to handling the issue from a global point of view. In both, responsibility could be assigned to the global tech companies and the degree to which is inconclusive in either. It then becomes a matter of usefulness in a global context, where I argue utilitarianism prevails based on its better accountability for the global nature of AI production and deployment and its grounding in methods similar to those used by tech corporations themselves which increases its likelihood to drive actionable change. I acknowledge that GDP is a metric that originated in the Global North and has since been challenged for undervaluing non-monetary values and explore the potential positive implications of including alternative measures such as Amartya Sen's capabilities systems and community wellbeing under the ubuntu philosophy in later sections. However, I base my argument mainly on GDP-based measures of benefits based on its current prevalence in global decision making and its importance to developing countries for whom a minimum level of GDP is often a prerequisite for achieving other development goals.

Definitions

I am fully aware of the complexities in grouping countries into developed and developing. This is done for conciseness of expression and largely aligns my definition with the UN definition based on scores on various Sustainable Development Goal dimensions (UNSTATS, 2024). My intention is to distinguish between populations of countries that produce AI and/or have economies advanced enough to benefit from productivity-enhancing use-cases of AI and populations of countries where human development and infrastructure indicators are on the lower end and hence where life-saving applications of AI are most needed.

Therefore, I define "AI divide" as inequality in AI benefits distribution as it relates to individuals' effective access to and benefits from AI on a scale comparable across countries.

This is roughly aligned with Beitz's definition of global inequality: "when I speak of inequalities among societies or states, unless otherwise noted, I shall mean this as shorthand for inequalities among the persons who inhabit them taken as a single group." (Beitz, 2001).

Finally, I define responsibility as a fundamental moral obligation and legal accountability stemming from ownership or control over something, in this case of AI technology. This is to contrast the current situation where corporate ESG efforts aimed towards benefiting the developing countries are seen as a "good thing to do" rather than a moral obligation.

1. Review

1.1. Libertarian

I begin examining the question on responsibility over equal spread of AI benefits from the point

of view of libertarianism, given its dominance in the legislative landscape of the countries where AI is created. Libertarian thinking is often evoked by tech CEOs when advocating for government non-intervention, such as Tim Cook's statement on US government requests to decrypt iPhone OS that "would undermine the very freedoms and liberty our government is meant to protect" (Cook, 2016).

Libertarian stance on responsibilities of tech companies towards citizens of developing countries would emphasise non-intervention on the grounds of (1) violation of intellectual private property rights, (2) free markets' superiority in addressing issues.

Classic Libertarians such as Locke in his "Two Treatises on Civil Government" established natural rights, which included the right to property, protecting which he saw as one of the key functions of governments: "The great and chief end…of men…putting themselves under government is the preservation of their property" (Locke, 1884). It could be argued that strong IP laws of Western democracies and libertarian-like culture of Silicon Value was what contributed to AI breakthroughs in the first place and therefore companies and their shareholders are entitled to the full benefits of their innovation without an obligation to share it with others.

Libertarian economists would argue that free markets would achieve the goal of bridging the "global AI divide" better than any forced redistribution. Milton Friedman famously said: "The great virtue of a free-market system is that it does not care what colour people are; it does not care what their religion is; it is the most effective system we have discovered to enable people who hate one another to deal with one another and help one another." (Friedman, 1993). Among examples of market efficiency given by libertarians is healthcare provision, where government intervention can lead to price inflation, decreased quality of care due to reduced competition. Drawing analogies with AI, removing obstacles to data access and technology deployment in developing countries will be incentive enough for tech companies to provide most efficient entrepreneurial solutions to developing countries' needs.

1.2. Nuance

According to some libertarian thinkers, resource redistribution could be justified based on the following arguments: (1) rectifying past injustices and (2) protecting positive rights. Nozick, for example, mentions just acquisition and rectifying past injustices where it is plausible to do so. His principle of rectification of injustices in holdings requires that parties be returned to the situation they would have been in had the injustice not occurred. (Nozick, 1974). If a corporation has come to possess technology and profits from it in an unjust way, then they should be redistributed to its original holders. In the context of AI, data ownership comes up often and redistribution on the grounds of data collection practices could be made given not only the ownership of data collected from the developing countries but also an outsized role of developing countries in database creation through data annotation.

The idea of positive rights also allows for some redistribution in cases where severe inequality is preventing citizens from exercising control over their lives and therefore limiting their freedom. For example Vallentyne argues on egalitarian grounds that profits based on natural resource exploitation should be redistributed among global citizens through a "global fund" in an egalitarian manner (2000). It could be argued then that there are some minimum entitlements that each individual has and in order to protect those, some redistribution from corporates to individuals is warranted. Companies in other industries such as construction and healthcare are often mandated by governments to protect such positive rights of their citizens through compulsory licensing and social housing projects.

1.3. Accounting for Global Context

Libertarian thinking often tends to focus on legislation within nation states and tends to downplay the role of global natural rights. Strict interpretations of classical Libertarianism only mandate individual states to protect the natural right to life within its borders. So, Locke viewed protecting rights to life, health and liberty as within the state's mandate (Locke, 1884). There seems to be a dissonance, however, between the idea of the right to life that is "natural" and the

fact that in a lot of the developing world this right is routinely violated through deaths from preventable causes. While during the time of Locke's writing the focus on a nation state might have been justified given the context of the emerging US independence movement, in the current globalised world, it is hard to confine the idea of natural rights to a single state. Resource redistribution arguments within libertarianism, such as that of Nozick as referenced above also tend to focus within the boundaries of nation states and while they could be extrapolated to international contexts, they do not explicitly address global injustices. It could be argued that current Western digital technological dominance is based on the history of colonialism and resource extraction that enabled select elites in the developed world to enjoy living and educational standards that allowed them to reach the levels of current technological innovation. "Europe is literally the creation of the Third World. The wealth which smothers her is that which was stolen from the underdeveloped peoples." (Fanon, 2001). This context, however, is largely absent from Nozick's thinking as he focuses on the justification of property rights within a society.

Vallentyne's idea of a "global fund" is a rare example of left-libertarians addressing the issue of global injustice. While the focus on natural resources which are viewed as a commonly owned good is not directly transferable to the issue in question, extrapolating this line of thought could serve as an argument for distributing the benefits of AI broader. As AI is trained on user data and hence could be seen as a public good, the distribution of its benefits should be more equal globally.

1.4. Example

Libertarian thinking is behind some of the current ESG efforts of tech companies. Open-source models are often lauded as prioritising social good over profits. Connectivity-focused projects, such as Meta's Free Basics and Express Wi-Fi aimed at providing affordable Wi-Fi to emerging markets through hotspots are underpinned by the idea that providing opportunities and removing barriers would allow the free market to alleviate economic disparity (Meta, 2024). Mark Zuckerberg in his argument for connectivity mentions "The richest 500 million have way more money than the next 6 billion combined. You solve that by getting everyone online, and into the knowledge economy." (Wired, 2013). One issue with this is that providing connectivity and source models alone rarely leads to progress in developing countries. Developed world software developers benefited most from open sources systems, while applications for developing countries are much harder to come across online. There is a long temporal lag between providing connectivity and benefits of technology being felt economically within communities. Interestingly, the META Express Wi-Fi project has now been scaled down and no tangible results were reported, which is perhaps telling of the efficiency of the approach based solely on providing access without further distributive assistance (TechRadar, 2024).

1.5. Evaluation

Overall, libertarianism is not very instrumental in evaluating whether or not the responsibility over AI benefits redistribution lies with tech companies. The school's coverage of global interdependencies is limited and its strong focus on a single nation state makes it difficult to apply to the globalised nature of AI production and consumption. Additionally, when informing ESG efforts in the real world, the school's recommendations fall short of delivering meaningful results to developing markets.

2. Rawlsian

A classical Rawlsian stance would posit that it is indeed the moral responsibility of companies that create transformational technologies to ensure a more equal global distribution of the benefits of such technologies. Based on the idea of the "veil of ignorance" if a neutral objective person would be deciding which use cases to deploy AI towards, she or he would direct it towards solving the most pressing global issues such as climate change, food security, illiteracy etc (Rawls, 1999). A lot of these use cases are relevant to developing countries, unlike the current productivity focused use cases.

A general criticism of Rawlsian thinking is its impracticality in a world where existing resource distribution is far from the original state. It is implausible that tech companies

would agree to develop AI systems without commercial interest at stake and it is likely that in such a case the pace of AI development would be impeded.

Recognising the impracticality of "veil of ignorance", Rawls also argued under his difference principle that inequality could be justified as long as it makes the worst off in society better off: "While the distribution of wealth and income need not be equal, it must be to everyone's advantage, and at the same time, positions of authority and offices of command must be accessible to all." (Rawls, 2005). Every policy and investment decision then needs to consider its impact on the worst-off in society, something which current tech companies' investment principles do not and something that is quite different from libertarian thinking where concern for the worst-off is not a given.

2.1. Nuance

While the difference principle is powerful in putting a condition on inequality-producing decisions, it is difficult to measure what "to everyone's advantage" means. Rawls does not offer a single definition to subjective notions of "worst-off" and "improvement". This could hence be interpreted in a number of ways as it relates to the global AI divide (Rawls, 2005). One interpretation could be that as long as AI development for commercial purposes also funds some socially positive applications no further redistribution is needed. So, if people in the developing countries are slightly better off than what they would have been without any AI development, this is sufficient. In this instance, Rawlsian thinking could potentially paradoxically recommend a similar or lower redistribution of resources than that warranted by Nozick's redistribution principle discussed earlier. While this does not seem to be the intention of the theory, there is the danger that this vagueness could be used to "green wash" corporate ESG efforts.

2.2. Accounting for Global Context

Similarly to libertarianism, classical Rawlsian theory focuses on justice within the domestic nation state. While in his *Law of the Peoples* Rawls does state that "Peoples have a duty to assist other people's living under unfavourable conditions that prevent their having a just or decent political and social regime," he mainly places assistance responsibility with the developed states rather than individuals or corporates (Rawls, 1999). Under this constraint, tech corporations would be responsible for addressing inequalities within their own countries and communities as a priority to those of other nations, which given the fact that developing countries lack AI development capabilities will likely exacerbate rather than alleviate the global AI divide.

Ideas of Charles Beitz extended Rawlsian ideas in a domestic society to our duties as global citizens. He argues that the differences between the domestic and global realms have been overestimated and a lot of the arguments in favor of equality as domestic justice could be applied in a similar way to equality as global justice. "There is a dispute about whether we should understand global justice, so to speak, as an enlarged image of justice in one society – and correspondingly demanding – or rather as a distinct construction, suited to a world that cannot be described as a single society, and therefore as demanding less than its domestic analogy." (Beitz, 2001). He argues that for reasons of shared humanity and interdependence our duties to citizens of other nations are the same as to those in our own countries. Beitz's idea of the "Global Resource" dividend is surprisingly similar to that of left-libertarian Vallentyne and once again could be extended to be a "Global AI" dividend through the notion that AI is a common global good trained on global data and therefore that its proceeds could be distributed among developing nations.

2.3. Example

Some organisations such as OpenAI's original mission was to "ensure that artificial general intelligence benefits all of humanity."(OpenAI, 2024). which appears close to the Rawlsian ethos. While conceding the need for a commercial arrangement to reach scale, OpenAI argues that it "continued to advance our mission by building widely-available beneficial tools" in its recent blog. (OpenAI, 2024). The example given by OpenAI as it relates to developing countries is the Digital Green collaboration in Kenya aimed at improving agricultural knowledge in the current climate change

affected environment. From reading the customer success story on OpenAI's website – it is not clear what role OpenAI itself played beyond providing the model. It would seem that the bulk of the effort fell with the Digital Green organisation itself – an NGO with diverse funding sources. Furthermore, this is the only developing country case study listed on OpenAI's website with other examples covering developed country applications. Additionally, and similarly to other tech firms, OpenAI was criticised for its working standards used for human labellers in Kenya in terms of wage levels (between around $1.32 and $2 per hour) and working conditions (Time, 2023).

OpenAI's example is illustrative of the idealism of the Rawlsian school of thought that appeared to have been "reality checked" in this case.

2.4. Evaluation

Overall, although later Rawlsian thinkers do explore the context of justice in the globalised world, the solutions they offer are surprisingly similar to those offered by libertarians and rather impractical and hence unlikely to have impact in the real world.

3. Act Utilitarianism

Act utilitarianism would approach the question of how a company should invest its resources, based on what would produce the best ultimate outcome for the majority of people (Bentham, 2012). I will proceed with illustrating a hypothetical approach to such evaluation below, with the purpose of illustrating the utilitarian approach rather than reaching a conclusion on the recommended benefit reallocation amount.

Going back to the PWC report referenced earlier – AI applications included in this report are those related to increased efficiency and accuracy in applications most relevant to developed countries. They measure gains in productivity and extrapolate this to resultant economic benefits. Based on this, applications in developing countries are relatively limited considering the smaller sizes of their economies.

AI use cases not considered in the report are those related to death prevention and improvements in basic quality of life, such as alleviating malnutrition, increasing literacy and preventing death and displacement through natural disasters. It could be argued that such applications will improve outcomes in emerging markets by the product of the lives saved and the current GDP per capita (or a significant proportion of it) as lives saved will create additional economic benefits proportionate to their number.

I attempt to make a calculation below where including such AI applications into the equation will suggest an optimum redirection of investment from the development of current commercial AI applications into life-saving AI applications more relevant to developing countries. Before doing so, I would like to reiterate that this is purely to illustrate the benefits and pitfalls of a utilitarian approach to this issue. Through performing this simplistic calculation I illustrate the possibility of doing so with an alternative objective in mind – calculating the economic benefits of life-saving AI applications. To my knowledge no comprehensive attempt to do so has been done globally and therefore there is no existing body of expertise. Through my simplistic demonstration I call on this viewpoint to be taken into account in similar future evaluations.

Following this approach, redistributing 12% of the benefits from current commercial uses to those designed to save lives in the developing world (aimed at preventing hunger, natural disaster, and treatable diseases) is the equilibrium point. This is illustrated in Table 2 below with corresponding assumptions and caveats.

Table 2: Illustration of hypothetical utilitarian approach to sizing benefits of life-saving AI applications in developing countries

	Population	GDP gains: current applications ($, bn)	Lives saved: life-saving applications	Lives with significantly improved quality	GDP per Capita ($)	Benefits Redistribution ($)	Benefits Redistribution (%)
Developed countries	1368m	12,500	18m	409m	6,770	-1,508	-12%
Developing	6632m	9,300				1,508	16%
	8000m	**21,800**					
	Population	GDP gains: current applications ($, bn)	Lives saved: life-saving applications*	Lives with significantly improved quality**	GDP per Capita ($)	Benefits Redistribution ($)	Benefits Redistribution (%)
Developed countries	1368m	12,500	18m	409m	6,770	-1,508	-12%
Developing	6632m	9,300				1,508	16%
	8000m	**21,800**					

The methodology of my simplistic exercise could have been greatly improved given time and access to experts in relevant domains. However, I want to acknowledge that even the most comprehensive methodology would pose a number of challenges.

1. **Uncertainty**: one of the main criticisms of utilitarianism is that the outcomes of actions are extremely hard to predict especially when complex concepts or new technologies are involved. Predicting the impact of AI on a broad range of applications such as healthcare and agriculture accurately is extremely difficult and might lead to misleading conclusions.

2. **One unit of measurement**: Quantifying the value of a human life, education, and food safety along in the same monetary units (GDP gains) as improvements in financial planning and business productivity is not only methodologically difficult but ethically hugely problematic. As the above exercise shows, the result of factoring in global suffering only results in a modest recommendation for redistribution which is the by-product of the assumption that productivity in the workplace could be compared to the value of a human life.

3.1. Nuance

Using another quantitative measure of investment such as the OECD Better Life index, or Amartya Sen's capabilities framework, would have likely resulted in an even more favourable recommendation distribution in favour of life saving applications of AI. While admittedly facing the same methodological issue of scales, such a measurable, visual approach could act as a meaningful call to action, such was the case with Peter Singer's famous work "The Life You Can Save". Impact dimensions on "The Life You Can Save" websites such as "Health", "Education" and "Living Standards" are well defined with corresponding indicators, definitions and measurement systems. This acts as a motivator for investors based on utilitarian grounds. This approach works within frameworks understood and accepted by major philanthropic investors and corporations and in the case of individual donations has significantly elevated the profile of global philanthropy among the general public. There are some reality checks however that need to be kept in mind despite the powerful call to action of this methodology. Singer has famously advocated for directing up to ⅓ of one's income to philanthropy; however most individuals are not close to this suggested amount (WSG, 2015). Similarly, in the case of tech companies' investment distribution, we are unlikely to reach this "north star", however we have a chance of making some progress towards this goal.

The highly numerical and material nature of utilitarianism also opens it to criticism from non-western schools of ethical thought such as ubuntu for example, which emphasises the wellbeing of the entire community and values the wellbeing of all individuals in its own right. My proposed methodology implicitly justifies actions that benefit the majority in a zero-sum game with implicit need for sacrificing the well-being of one group to increase the well-being of another. Ubuntu, on the other hand, rejects the notion that the well-being of some can be sacrificed for the benefit of others. Exploring whether both goals can be achieved simultaneously is a valid direction of enquiry

that would enrich the argument as part of further exploration.

3.2 Accounting for Global Context

Utilitarianism does provide a useful basis for addressing the issue through allowing for transcendence of national boundaries. Peter Singer in "The Most Good You Can Do" provides a compelling argument why philanthropists should invest in alleviating global poverty as opposed to poverty in the developed world (Singer, 2015) . While not denying the negative effects of poverty in developed countries, he points out the various security mechanisms available to citizens of the developed countries through taxation for example. He concludes that there is a wide gulf between poverty in the developed and developing world and the value of donations (in life outcomes) is far greater in the developing world: "their dollars go much further when used to aid those outside the affluent nations" (Singer, 2015). Singer's writing has been influential in alleviating global poverty and its explicit address of the global nature of inequality makes it relevant to evaluating the issue of the global AI divide.

Utilitarianism also has a visual and measurable quality to it, something that makes it compatible with driving change within corporations. The presentation of productivity and lifesaving use cases side by side, while problematic, does allow for visual and measurable accountability in a format familiar to the corporate world. Like Peter Singer's case for individual donations, it helps substantiate the claim in measurable terms. The investment disparity, if presented internally within corporations, could ignite employee activism towards influencing corporate investment decisions or towards institutionalising individual voluntary time donations in an arrangement similar to the legal profession's pro bono practice.

3.3. Example

As opposed to Meta's connectivity projects, Google's Build for Africa investment announced in 2021 does attempt to combine connectivity provision (through a subsea cable) with initiatives aimed at local talent development (through an AI research facility in Ghana) and product usability improvements (increased language inclusion and Maps coverage) along with start-up empowerment programmes. Google's announcement invokes utilitarian values through its explicit mention of benefiting the lives of most people: "benefits of the digital economy for more people by providing useful products, programmes and investments" (Gajria, 2022). Benefits are also quantified in GDP terms similarly to the PWC report and anchoring to $1Bn as the investment sum, shows how numerical grounding could be instrumental in motivating corporate action.

While the focus on talent development and product localisation is a move in the right direction – it is worth noting that the $1Bn investment (spread over 5 years) is equivalent to only 0.4% of Alphabet's 2023 R&D budget (Alphabet Annual Report, 2023). This is illustrative of the pitfall of utilitarian thinking in not recommending enough redistribution despite being a useful framework for motivating corporate action in general. The results of this programme are yet to be seen.

3.4. Evaluation

Overall, however, despite its methodological and ethical challenges, I would rate utilitarianism as the more useful theory among the three considered in evaluating the issue of the AI global divide. This is based on its rich recent body of thought on global justice and its measurable and visual basis that is likely to lead to actionable change both at the corporate level and at the level of individual employees.

Conclusion

In conclusion, the three schools of thought come to three different conclusions on the moral debate of whether or not it is the responsibility of tech companies to alleviate the global AI divide. Libertarianism would argue that it is not the responsibility of the tech companies, Utilitarianism that it is, to a certain quantifiable degree, while Rawlsian that it definitely is as part of a moral imperative. Despite these broad-based conclusions, when factoring in nuance, there is room for an increased benefits redistribution within each school of thought, based either on rectifying past injustices, reframing benefits in terms of increases in "capabilities" rather than GDP or fully extending the veil of ignorance principle to the global context.

Libertarian focus on the nation state makes it hard to make judgements in a globalised setting and due to its original non-intervention stance, those recommendations aimed at redistribution on the basis of past injustices or positive rights, such as a Natural resource tax are quite unrealistic. Rawlsian stance although going beyond the nation state eventually is fairly vague in defining what making those globally worse off better is. Actionability of its recommendations such as the Global Resource dividend is also fairly improbable given its highly conceptual nature. Traces of libertarian and Rawlsian thought can be seen in various empirical examples of existing tech companies' investments in developing countries such as ESG and connectivity directed efforts and partnerships with local NGOs where both can result in quite in "greenwashing" one in line with its intention and the other due to its impractical nature.

Utilitarianism scores highly as it relates to evaluating this moral debate for two reasons: its well-established body of work on global inequality and actionability of its recommendations. The visual nature of the results it demonstrates is likely to draw more realistic action such as employee activism and increased collaboration with unilaterals as a result. Therefore I argue that this is the most useful framework in this case albeit not problem-free especially as it relates to its methodological complexity and ethical challenges. While utilitarianism does set a good basis for empirical actionability, in order for resource redistribution to increase in a meaningful way we need to employ cognitive behavioural tools, such as for example drawing from the more successful examples of cooperation within climate change and nuclear regulation domains. This could be a good next step to consider to further enrich this argument further.

Incorporating the viewpoints of other schools of thought, such as duty ethics for example, while out of scope of this paper would enrich the analysis and I recommend these as next steps of this line of inquiry. So, exploring how duty to the company shareholders might negate the recommendations made under utilitarianism. This would also have practical implications, for example due to the need to change the governance structures of corporations to mitigate this obstacle. On the contrary, duties of individual employees to their communities of origin could further strengthen the plausibility of some courses of action such as employee activism.

References

Beitz, Charles R. (2001). Does global inequality matter? *Metaphilosophy*, 32 (1-2), 95–112.

Bentham, J. (2012). *An introduction to the principles of morals and legislation*. Dover Publications.

Fanon, Frantz. 2001. *The wretched of the earth*. Penguin Modern Classics.

"Friedman-Government-Problem-1993.pdf." n.d. Retrieved from https://www.hoover.org/sites/default/files/uploads/documents/friedman-government-problem-1993.pdf.

Gajria, Nitin. (2022). "Delivering on Our $1 Billion Commitment in Africa." Retrieved from__https://blog.google/around-the-globe/google-africa/delivering-on-our-1b-commitmen t-in-africa/

Levy, Steven. (2013). Zuckerberg explains Facebook's plan to get entire planet online. *Wired*. Retrieved from https://www.wired.com/2013/08/mark-zuckerberg-internet-org/.

Locke, J. (1884). *Two treatises on civil government*. George Routledge and Sons.

Nozick, R. (1974). *Anarchy, state, and utopia*. Basic Books.
Pwc-Ai-Analysis-Sizing-the-Prize-Report.pdf.. (2017). Retrieved from https://www.pwc.com/gx/en/issues/analytics/assets/pwc-ai-analysis-sizing-the-prize-report.pdf.

Rawls, John. (1999a). *The law of peoples*. Harvard University Press.
Rawls, John. (2005b). *A theory of justice*. Belknap Press.

Singer, P. (2015). *The most good you can do: How effective altruism is changing ideas about living ethically*. Yale University Press.

Slater-Robins, Max. (2022). "*Meta Is Shutting down Its Express Wi-Fi Service.*" TechRadar Pro. February 1, 2022. https://www.techradar.com/news/meta-is-shutting-down-its-express-wi-fi-service.

United Nations Statistics Division. n.d. Methodology. Retrieved from https://unstats.un.org/unsd/methodology/m49/.

Vallor, Shannon, ed. (2000). *Left-libertarianism and its critics: The contemporary debate*. Palgrave Macmillan.

Yu, Danni, Hannah Rosenfeld, and Abhishek Gupta. (2023). The 'AI Divide' between the global north and global south. Retrieved from https://www.weforum.org/agenda/2023/01/davos23-ai-divide-global-north-global-sout h/

Wolfe, Alexandra. (2015). Peter Singer on the ethics of philanthropy. WSJ Online. April 3, 2015. Retrieved from https://www.wsj.com/articles/peter-singer-on-the-ethics-of-philanthropy-1428083293