Generative AI as interpreted through a normative framework of knowledge labour well-being

Name student: Alexander van Erven

Name supervisor: Prof. dr. Jack Vromen

Name advisor: Dr. Christoph Brunner

Main study: International Bachelor of Economics and Business Economics

Date of completion: 15th of July, 2023

Number of words: 9853 (excl bibliography)

Table of Contents

List of Abbreviations

AI	-	Artificial Intelligence
GPT	-	Generative Pre-trained Transformer
LMM	-	Large Language Model
API	-	Application Programming Interface

Introduction

How would the well-being of knowledge labour be affected if the cost of cognition and creativity were to follow Moore's law? In a 2021 blog post titled *Moore's Law for Everything*, Sam Altman uses the well known observation by Gordon Moore of the long term exponential productivity growth of semiconductor manufacturing to warn that AI and robotics are on a path to expose virtually all human labour to technological automation (Altman 2021). The recent proliferation of generative AI in the form of commercially available large language models is provoking some to wonder exactly how indispensable humans are to cognitive or creative tasks, and what the associated consequences would be for the cost of knowledge labour in the long term. As with any promising new technology, predictions regarding the future of generative AI diverge significantly. Some hail the promises of this technology by pointing out its capacity to facilitate economic growth (Goldman Sachs 2023), while others are either sceptical of its capabilities (Brynjolfsson 2023), or fear the ramifications of its success (Merchant 2023).

This paper will explore the effect mass-adoption of generative AI will have on well-being by proposing a normative framework designed specifically to evaluate the effect generative AI will have on the well-being of knowledge labour. By virtue of consisting largely of tasks that require complex cognition and long periods of formal training, knowledge labour has historically faced fewer threats of technological disruption relative to other segments of labours. Capabilities particular to generative AI subvert this trend, specifically threatening technological disruption in the knowledge domain of labour (Eloundou et al. 2023). The concept of labour inherently describes the relationship of a human being to the economy, so although the framework aims to capture a holistic account of a knowledge worker's well-being, this analysis represents only a subset of a working human being's overall state of well-being. The framework established in this paper separates the well-being of knowledge labour into two normative dimensions. The first dimension reflects the effect generative AI has on a knowledge worker's capacity to flourish. The second dimension of well-being reflects the technology's impact on the economic relevance of knowledge labour. The implication that follows from the framework, and the thesis of this paper, is that generative AI reduces the economic relevance of knowledge labour while enhancing the capacity of knowledge workers to flourish.

This paper consists of five distinct sections that enable this analysis. Section 1 begins with a description of generative AI followed by a definition of knowledge labour. Section 2 explains the mechanism and construction of the framework. Section 3 analyses the consequences of generative AI on the capacity of knowledge labour to flourish. Section 4 analyses the consequences of generative AI on the economic relevance of knowledge labour. Section 5 derives implications from the framework, reflects on the framework in discussion, and provides a conclusion.

Section 1: Context

How does generative artificial intelligence work?

Generative AI, such as GPT-3, leverages deep learning models called generative neural networks to create new content. These networks are trained on vast amounts of data, learning patterns and relationships. During inference, the model takes an input, such as a prompt or partial text, and generates a coherent output based on its learned knowledge. It employs techniques like attention mechanisms, recurrent neural networks, and transformers to capture contextual information and generate creative responses. By iteratively refining its outputs through training and feedback, generative AI systems can produce increasingly realistic and human-like outputs.

In the spirit of seeing it to believe it, the preceding paragraph was entirely written by OpenAI's ChatGPT in response to the following prompt: "give a high-level abstract explanation of how generative AI technically works, in a 100 words". To ensure the response it generated corresponded with the contextual requirements of this paper, some rudimentary prompt engineering was employed. The descriptive terms "high-level abstract" informs the system that the explanation of how it "technically works" ought not be too technical without oversimplification, while the "100 words" command conveys the preferred response length. If the generated response did not correspond with the desired intention, the user has three options. Regenerate a response based on the same prompt, leveraging the indeterminate nature of the system to produce a different plausible response. Alternatively, the user could leverage the conversational nature of the system to inform it what specifically was undesirable in the original response, resulting in a new response that combines the original response with the user feedback. Finally, the user can simply reformulate their prompt based on their unmet requirements.

Within the vast domain of artificial intelligence technologies, generative AI is a subset of machine learning, where machine learning refers to the technology's capacity to learn and discover patterns from a large set of training data. If the data is labelled and curated, this form of learning is categorised as supervised machine learning. This supervision only extends to the contents of the dataset, since all the parameter setting is performed by the algorithm independent from human input. Training machine learning algorithms broadly speaking consist of two elements, models attempting to optimally represent the dataset with some parameters and a loss function to evaluate a model's performance. These performances are subsequently ranked, with only the best models surviving each training iteration. The parameters are then carried over to the next training, new models are subjected to reinforcement learning where the parameters are independently tweaked once again in the hope that the changes make the model represent the dataset even better. Generative AI combines this capability with the capacity to generate new outputs that correspond with the patterns it learned from its training data. For the purposes of this paper, the main generative AI models to consider are Large Language Models and Stable Diffusion models. LLMs are trained with enormous amounts of text, allowing them to complete tasks that require natural language. In essence, it does so by functioning as a supercharged auto-complete algorithm, predicting the likelihood of each subsequent letter or word based on what it or the prompt had already written. Through this training, commercially available technology like GPT-3 has become so familiar with natural language that it appears to mimic an ability to reason (OpenAI 2023). Stable Diffusion is a technology typically trained with visual data, and has consequently learned patterns that allow it to recognize visual objects and reproduce artistic styles (OpenAI 2021).

What is the potential of generative AI?

As with any new technology, the relevant question is as much 'what can it do?' as it is 'what will it be able to do?'. The prevailing notion surrounding generative AI at the time of this paper's writing, is that it has the promise to become a general purpose technology (Eloundou et al. 2023). This prediction invokes several astonishing technological comparisons, such as electricity, internet and the steam engine. The category is typically defined by technologies meeting the following criteria: widespread and diverse adoption, continuous improvement, and enabling complementary innovations (Jovanovic and Rousseau 2005). At the time of writing, generative AI is not demonstrating obvious deficiencies in any of these dimensions, which elicits a lot of uncertainty surrounding the future capabilities of this technology. Particularly significant is the ease with which this technology can be incorporated into new or existing software, since the large AI companies are making generative AI available through standard APIs. These allow any authenticated third party software developer to leverage the capabilities of generative AI within the code of their own software application. Microsoft recently incorporated OpenAI's GPT-4 into its Bing search engine, and has promised to do the same with its productivity software such as Word, Powerpoint and Excel (Spataro 2023). The technology here is leveraged as an assistant that makes it simpler to use advanced features within their software as well as aiding the user in reviewing, adapting, continuing or summarising text.

Aside from the explicitly intended functions of the technology, advanced generative AI systems are also demonstrating emergent properties. Upon publicly unveiling GPT-4, OpenAI showed that the system scored in the 89th percentile for SAT Math and in the 90th percentile on the Uniform Bar Exam, which is a test every aspiring lawyer in the United States needs to pass (OpenAI 2023). When OpenAI released Github Copilot in 2021, it proved that generative AI is also capable of writing and assessing computer code (Kalliamvakou 2022). Since then, ChatGPT has inherited this capability, and people have shown that this cocktail of capabilities allows the technology to automatically prompt itself until it accomplishes its assigned task (Ortiz and Windsor 2023). To some, this is evidence of a unique potential feature of generative AI, namely its capability to write code that would improve itself (Schmidt and Wang 2023). If its ability to improve itself exceeds the programming abilities of its creators, and the technology is granted the freedom to do so at will, the future capabilities of this technology will be entirely unpredictable. Nonetheless, it will remain entirely predictable that with each additional capability, the economic impact of generative AI will grow. While not completely relying on it, this paper will continue under the assumption that it is realistic for generative AI to realise some of the aforementioned potential.

What are inherent flaws in the technology?

The main impediment for the mass adoption of generative AI in economically relevant tasks is the inherent epistemic unreliability of its generated responses. This stems from its machine learning architecture that enables it to independently recognise patterns from large datasets. This independence allows it to derive relationships that are far too complicated to code directly, but also means that the relationships it derives are very difficult to manually verify or adjust. This stands in contrast to rules based, deductive AI systems that rely on explicit coding by humans. By relying exclusively on its training dataset to recognize patterns, generative AI systems are susceptible to the same limitations to external validity that arise from inductive reasoning (Kumichev 2022). Regardless of our temptation to anthropomorphize the technology, it has no comprehension of the real world, and it can only 'learn' about this world inductively. AI researchers are referring to this as the alignment problem, and is the subject of intense debate at the time of writing (Christian 2020). The most pertinent consequence of the alignment problem to this paper is the proclivity of generative AI systems to produce convincing hallucinations. These are responses that sound completely plausible, but are either entirely fabricated or based on fictitious facts about the real world (OpenAI 2023). This property actually risks becoming more harmful as the system becomes more reliable, as its inherently probabilistic response will increasingly be accepted at face value. The partial solution to this is to simply give the system access to the data that it requires to respond correctly to the request; "the model does not know what time it is, but you can give it a watch" (OpenAI 2023, p4). This nevertheless leaves the problem of independently acquired biases inherent to machine learning systems unaddressed (Manyika, Silberg, and Presten 2019). Since the system independently adjusts its parameters to recognize patterns in its training dataset, the patterns it establishes will incorporate the intentional or unintentional bias contained in the data of the training dataset. Some of this bias can arise from intuitive sources, such as failing to ensure adequate diversity of data. What is more concerning, is that the pattern generating disposition of the system lends itself to propagate complex cultural biases, or instead hallucinate entirely nonsensical correlations that can have social consequences. If left in charge of human resource management, the system could be inclined to propagate recruitment and promotion biases based on data that lacks counterfactuals, or inexplicably always deny promotion to individuals with more than three m's in their name. While some of these flaws are difficult philosophical pills to swallow, economic participants will adopt the technology purely on pragmatic terms. The inherent epistemic uncertainty associated with the outputs of generative AI will expose society to a new category of risks, however homo economicus is only concerned with rationally maximising its self-interest. As such these systems merely need to achieve a critical level of reliability where the cost of oversight and rework is less than the best forgone alternative.

Defining knowledge labour

Labour is not only a factor of production, but also a dynamic subset of human beings defined by their specific relationship to the economy. This paper defines knowledge labour as a subset of labour possessing specialised, formally trained knowledge or expertise, for which they are employed. Knowledge work is typically associated with a long period of job preparation, and predominantly consists of cognitive tasks (Reinhardt et al. 2011). This includes people who produce knowledge or technology such as researchers, engineers and programmers, but also those who rely on acquired knowledge to perform their economic function, for example: analysts, teachers, lawyers, medical professionals, journalists and consultants. This paper assumes a broad interpretation of cognitive tasks to include aspects of creativity, thereby also referring to writers, artists and designers as knowledge labour. Knowledge labour is primarily contrasted by physical labour, which is defined by tasks that either require the physical presence or physical work of a person. This distinction is not merely an artefact of abstraction, society has cultivated distinct connotations for each ‘form’ of labour. This is in part because knowledge labour earns significantly more than physical labour on average (Glassdoor 2022). William Baumol’s cost disease theory provides a partial explanation for this through describing why the cost of labour intensive services increases over time (Baumol 2012). As wages increase for skilled workers in industries of high productivity growth, skilled workers in labour-intensive services demand an equivalent wage increase without corresponding productivity growth in their industry. Historically, most productivity growth has come from the introduction of new technologies, evoking the concept of skill-biassed technological change, which suggests that technological progress raises the demand for skilled labour over unskilled labour (Katz and Murphy 1992). Defined by their specialised expertise, knowledge labour has largely benefited from this technologically enabled demand, without itself being exposed to significant technological disruptions. Generative AI has the potential to be precisely the kind of technology that would disrupt knowledge labour (Eloundou et al. 2023). This paper considers jobs and occupations to be composed of a collection of tasks, in line with the task-based model of automation (Autor, Levy, and Murnane 2003). In this view, automation leads to task-displacement, only amounting to job-displacement when a critical number of tasks in the job become automated.

Section 2: Normative framework

Mechanics of the framework

In order for the consequences of generative AI to be evaluated, a suitable normative framework needs to be applied, which in turn can be assessed according to the intended scope of the framework and its stated normative claims. Due to the specific technology that is the focus of this paper, the scope of this framework is exclusively confined to knowledge labour. While this self-imposed restriction necessarily leads to blind spots, it allows the normative framework and subsequent analysis to be tailored to the idiosyncrasies of the demographic that is most affected by this technology. An example of why this demographic emphasis is necessary can be found in what economists call the "luddite fallacy", which indicates the supposedly mistaken belief that technological innovation leads to structural unemployment (Jerome 1934). From a macroeconomic perspective labelling this belief a fallacy has thus far been proven correct, yet the 19th century British textile weavers after which the fallacy is named saw their well paying jobs disappear while their profession became structurally obsolete (Krugman 2013). The people who suffer most consequentially from new technologies, and in turn resist its adoption most fervently, are often those who the technology threatens to replace.

Well-being

With the scope established, the framework requires an evaluative mechanism informed by normative judgement. The governing mechanism of this paper is the maximisation of knowledge labour well-being. The concept of well-being is understood in the broadest possible sense, referring to all things that would be 'good for' an individual, in other words representing the prudential value of a thing to an individual (Dorsey 2021). This therefore accounts for physical, psychological and communal well-being as well as financial and political well-being. To avoid a vague analysis resulting from such a broad concept, this paper strictly defines two constitutive normative dimensions that make up the well-being of a knowledge worker. This allows for generative AI consequences to be analysed directly in terms of values that knowledge workers necessarily care about, while avoiding a reductionist economic productivity analysis. Note that such an analysis does not yield a complete account of the well-being experienced by a person employed as a knowledge worker, since the framework is limited to only identify aspects of well-being that are sensitive to changes in a person's relationship to the economy.

Flourishing of knowledge labour

The first normative dimension according to which the framework will analyse the impact of generative AI on knowledge labour well-being is its effect on the capacity of knowledge workers to flourish. This refers to aspects of well-being internal to an individual, who possesses agency to affect how these internal constituents of well-being are manifested. In part, this reflects the innate desire of any human to be presented with new opportunities through developing and expressing themselves. This paper argues that the awareness of, and desire for personal growth is a defining feature of people engaged in knowledge labour. This desire of knowledge workers to engage in self-improvement is revealed by the observation that the vast majority of knowledge workers finish at least 3 years of formal education post-high school (OpenAI 2023). Sometimes referred to as cognitive labour, knowledge workers are often engaged in tasks that rely on their cognition, thereby experiencing the benefits from greater cognition at a personal as well as economic level. The combination of a revealed proclivity towards self-improvement, and the possibility to extract elevated returns from it, shows that the desire of an individual knowledge worker to achieve the freedom and capacity to cultivate themselves is a value necessarily inherent to their demographic as explicitly economic participants. While the desire for self-expression and self-development is not exclusive to knowledge workers, this desire is effectively a prerequisite for a successful career in knowledge labour. Analysis of this dimension will explore the effects of generative AI on the capacity of knowledge labour for self-expression, self-development and their access to new opportunities.

Economic relevance of knowledge labour

The second normative dimension of the framework is the extent to which knowledge workers retain their economic relevance. Although it is a common human desire to feel recognized and included in a community, economic relevance is not a universal prerequisite to human well-being. Likewise, a person's relevance to society is not exclusively determined by their economic activity. The segment of humans that are categorised as knowledge workers however, are primarily defined by their relationship towards society, namely as labour. Though it is debatable if the human being who is occupied as a knowledge worker necessarily requires their ‘knowledge’ to be socio-economically relevant to achieve well-being, it is not debatable that without this relevance they would cease to be a worker. Consider the example of an early 19th century medical professional who was renowned for their knowledge and skill in attending to patients with Miasma. By the end of the century, even though they would still possess the same knowledge and skill, this knowledge is deemed entirely irrelevant by society (Susser and Stein 2009). It is unlikely he would continue to be compensated for this knowledge, indicating his exit from the category of knowledge labour, in the absence of his commitment to start developing from scratch in a new field of study. In contrast to the first dimension, which evaluates the consequences of generative AI in terms of the agency of an individual worker, the second dimension is concerned with the external context in which this agency is manifested. A worker can make their best guess to determine what training and subsequent knowledge will ensure their relevance, but it is never up to the worker to decide how relevant society considers them to be. It is thus a constitutive element of their well-being, but only indirectly within their control. The analysis of knowledge labour relevance is divided into an assessment of the relevance of capabilities possessed by knowledge labour when compared to generative AI, and an assessment of how generative AI might affect the ways in which knowledge labour could be applied.

Limitations and criticism of the framework

Before analysing generative AI in terms of the established framework, obvious limitations to such an approach need to be addressed. The normative aim of well-being maximisation is closely related to the utility maximisation aim found in utilitarianism, so it is important to note the consequences of significantly deviating from this framework (Bentham 1789). First, utilitarianism concerns itself with the consequences generative AI would have on the utility of all stakeholders affected by the technology (Schinkel 2008). As a result, the framework used in this paper is comparatively blind to changes in well-being for other segments of society as a consequence of this technology. More insidiously, this means the framework is also blind to 2nd-order well-being changes to other segments of society as a result of a change in the well-being of knowledge workers. It could for instance be the case that a decrease in overall knowledge labour well-being is welcomed by other segments of society, since knowledge workers have historically been under-exposed to the threat of technological disruption (Autor, Levy, and Murnane 2003). Society at large could benefit from a shrinking difference in the well-being of knowledge labour compared to other segments of the population, even if this only occurs at the expense of knowledge labour well-being.

The second important deviation from utilitarianism is the disaggregation of well-being into separate constitutive dimensions that explicitly reflect desired values of knowledge workers (Sen 2002). This act of disaggregation hinders a direct comparison between these values with the aim to formulate an overall measure of well-being. Instead, this paper argues that the previously defined constituents are irreducible, and necessarily incommensurable. By explicitly defining well-being, utilitarians would further critique that this framework is susceptible to generating an incomplete account of well-being. This is a legitimate concern that is mainly addressed by establishing a dualism of two specific, yet broadly defined concepts. However, this act of selection is itself based on an assumption that these constitutive dimensions are the primary mechanisms through which generative AI might affect the well-being of knowledge labour. Health and bodily integrity are crucial aspects of well-being, yet the framework omits these aspects since it is unclear if and how generative AI could promote or impede them.

The resemblance of the framework to utilitarianism, with its reliance on the concept of well-being also engenders its own criticism. Neither constitutive dimension is empirically measurable (Sen 1992), or a direct reflection of knowledge worker capabilities. The selected normative dimensions are inspired by the capabilities approach to reflect values which knowledge workers necessarily possess, however they only indirectly interact with the basic human capabilities outlined by Martha Nussbaum (Nussbaum 2006). Lastly, by interpreting well-being in terms of pre-defined, pre-existing values of knowledge workers, the framework is liable to disregard a change in these preferred values. The concept of adaptive-preferences suggests that the extent to which a person's true preferences are revealed is context dependent (Sen 1985). As such, if the progress or adoption of generative AI is sufficiently substantial, preferences of the knowledge labour demographic could change to no longer align with, or be accurately reflected by the dimensions of the framework. While this paper concedes that revealed preferences are context dependent, it argues that the selected normative dimensions are integral to the well-being of knowledge labour, such that if these preferences changed, the meaning of knowledge labour itself would change. This means that the framework is suitably constructed to assess the well-being of knowledge workers within the current occupational environment. However, if generative AI were to change hiring practices such that the subset of people identified as knowledge labour changes, the framework is only fully applicable to existing knowledge labour. Whenever possible, this paper will attempt to acknowledge the well-being of existing knowledge labour as well as future knowledge workers whose employment would be enabled by the adoption of generative AI, however only knowledge labour operating within existing business practices will be considered for the framework's formal implication.

Section 3: Flourishing of knowledge labour

Self-development of knowledge labour

As was indicated previously, a defining feature of knowledge workers is that they have at some point in their lives committed resources to develop specialised skills and knowledge. These resources consist of both time and money, indicating a revealed desire by these individuals to improve their own capabilities. Since the personal finances of an (aspiring) knowledge worker informs both the time they can afford to commit to themselves, and the quality of educational resources to which they gain access, economic inequality plays a significant role in determining the possibility and extent to which a person can develop themselves.

Generative Artificial Intelligence, and in particular Large Language Models, are technologies that look poised to dramatically reduce the barriers to knowledge. The primary enabler of this reduction is the almost non-existent skills or knowledge barrier for use of the technology itself (OpenAI 2023). In theory, the only skill required to interact with this technology is native language proficiency. While the best model today functions best in English, other competing models show that this is likely just the result of engineering priorities, and therefore not a fundamental aspect of Large Language Models (Google 2023). In turn, this implies the possibility for generative AI to dispense almost entirely with the language barrier to knowledge. Today, anyone who seeks to acquire highly specialised knowledge is required to be proficient in at least one of a handful of languages. Google's Bard shows that Indians from obscure regions of India might soon be able to acquire this specialised knowledge in their regional native languages.

The second fundamental aspect of generative AI that facilitates the personal development of a knowledge worker is the infinite iterability of the results it produces. If the user requires any clarification following the initially generated response, this response is infinitely adjustable by the user until the user determines that they understand, and are satisfied with the response. Over time, this would allow the system to generate content that optimally corresponds to their current competency, tailored by the feedback provided by the user. As such this technology is likely to accelerate the current push for institutions to facilitate personalised education at a large scale (Rouhiainen 2019). The educational non-profit Khan Academy has already demonstrated some of this potential by incorporating ChatGPT into its service (Fried 2023). Early indications are that the technology benefits the students as well as the teachers, by providing the students with immediate clarification and teachers with a digital teacher's assistant. The educative potential of generative AI is not restricted to its adoption by formal institutions, motivated individuals can also leverage the technology to self-educate in a new field from scratch. An example of such a motivated individual is Yohei Nakajima, the creator of BabyAGI (Arya 2023). This program acts as an autonomous agent that allows GPT-4 to prompt itself, thereby having it mimicking the independent agency to complete tasks. For the purposes of this paper it is not particularly significant that this capability was not intended or screened for by GPT-4's creators, it is instead very significant that Yohei had no coding experience whatsoever 3 days prior to BabyAGI's creation, prompting GPT-4 for anything he needed to know. While it is unlikely that he truly learned a lot about code writing in those few days, it clearly shows the learning potential these tools will provide in the hands of dedicated learners. With knowledge labour as a demographic defined by their specialised knowledge and skill, the adoption of generative AI will improve the well-being of knowledge labour in terms of their capacity for self-development.

Self-expression of knowledge labour

The need for accurate self-expression is almost universal amongst human beings. Many of those who experience this need the greatest organise their lives such that they can get paid to express themselves. Designers, journalists and even corporate consultants operate their jobs by finding ways to express the vision present in their mind. Generative AI could therefore have a substantial effect on knowledge labour well-being by augmenting their capacity for self-expression.

A significant augmentation of this capacity was already alluded to in this previous subsection, namely the effect of generative AI to reduce the barrier for skills or knowledge acquisition. Highly specialised knowledge labour might feel confined to express itself exclusively through the limited number of means within their specialisation. Instead, generative AI enables them to access a new range of modalities for self-expression. Some professional writers have already used Stable Diffusion Models to transcend their expertise of the written word by using the technology to generate images befitting a graphic novel (Edwards 2022).

Another significant augmentation of the capacity for self-expression arises from use of the technology as a source of creative inspiration. The instant and infinitely generative capacity of the technology is a mighty tool in the face of any creative's greatest source of anguish; writer's block. Through the use of clever prompting, creatives can use the technology to change, adapt, criticise or simply continue any work they are currently stuck on. The prompt can be leveraged to produce a response in line with direction provided by the user, or the prompt may be designed to produce a wildly unique response with each iteration. Companies such as Microsoft and Adobe are soon expected to directly incorporate this feature into all of their productivity software, further improving the accessibility of artificial creative inspiration (Spataro 2023). Teachers might for instance ask the tool to generate new test questions or multiple choice answers based on a prompt that contains a curated list of questions developed by the teacher in years prior. Creative inspiration can also be attained through the conversational use of LLMs. The tool could be prompted to respond as though it were the personification of a particular book, or simply the book's long dead author (De Cremer, Bianzino, and Falk 2023). Although these exercises are rife with hallucinations, for the purposes of creative inspiration, correspondence of the generated response to a real world truth is not a definite necessity. All of these examples show that generative AI has the potential to significantly empower the self-expressive capacity of individuals. To the extent to which self-expression is a need experienced by knowledge labour, generative AI also appears to promote their well-being in this aspect.

Access or denial of economic opportunities

Gaining access to a new set of economic opportunities can be an important driver of knowledge labour well-being. Likewise, a receding level of financial freedom is likely to have an equally drastic inverse impact on knowledge labour well-being. This subsection will therefore first analyse the effect new economic opportunities arising from this technology will have on the well-being of knowledge labour, which is followed by analysing the likely distribution of economic benefits extracted by this technology.

A defining criteria of general purpose technologies is that they enable widespread innovation in complementary technologies (Jovanovic and Rousseau 2005). Generative AI is already showing signs of this, with businesses in diverse industries announcing their plans to use the technology to augment their existing products. As technological innovation moves the product possibilities frontier further out, entrepreneurially minded knowledge workers are presented with a fresh canvas of unclaimed potential economic opportunities. It is notable that successfully actualising such opportunities would categorically shift some labourers into becoming owners of capital. Yet the financial potential accompanying this shift is likely to be perceived preferentially by knowledge workers, while not necessarily interfering with their capacity to exercise their profession should they desire to keep contributing in the form of labour. This shows that the capacity of generative AI to enable new technological innovations exploitable by knowledge labour is beneficial to their well-being.

Absent possible future technological innovations, generative AI has already proven its ability to enhance the personal efficiency of its users. This improved efficiency gives rise to several new economic opportunities available to knowledge workers. The first is to simply leverage the efficiency gain to become a more productive worker. This could come from using LLMs to draft emails, public statements or reports, while using Stable Diffusion Models to cycle through prototype ideas in minutes, where the production of a single visual prototype used to take two weeks (De Cremer, Bianzino, and Falk 2023). Alternatively, knowledge workers could leverage the new cost structures arising from this newfound efficiency to set up new businesses. Costly services required to set up a business in any industry, such as rudimentary legal or accounting aid can be generated at negligible marginal cost (Greene 2022). This also extends to services that are not vital, but instead help the business thrive, such as marketing or strategy consultants (Sharma 2023). Both these uses of the efficiency provided by generative AI, whether to maximise personal productivity or set up a business, contribute to the well-being of knowledge labour by improving their financial situation. On the other hand, knowledge labour could elect to leverage the efficiency gain to commit fewer hours to work while achieving similar levels of productivity. While this prediction has been made several times throughout the history of technological progress, some argue that this time might truly be different (Keynes 2010). This argument is based on the fact that many of the productivity benefits arise from efficiencies that require the worker to take initiative. If a worker used to take 8 hours a week reading and replying to emails, it is unclear to the employer how much of this time the worker continues to spend on this task with the availability of generative AI. This information asymmetry would suggest the worker has some power to decide how they will allocate the time saved (Arrow 1963). Efficiencies resulting from generative AI therefore not only improves the well-being of knowledge workers in terms of their availability to financial resources, but also time.

However, while the technology itself might contribute to the availability of new economic benefits, the distribution of these extracted benefits is determined by society. The next section analysing the economic relevance of knowledge labour will elaborate on whether, and to what extent generative AI might interfere with the economic value of their labour. The remainder of the current subsection will instead investigate the effect adverse distribution of extracted economic benefits might have on the capacity of knowledge workers to flourish.

The predominant mechanism at play in this context is the mathematical relationship of capital ownership versus labour. While this paper rejects Karl Marx's concept of surplus value, knowledge labour is already becoming increasingly dependent on paying capital owners for use of their technology simply to do their work. Marx defined surplus value as the value generated by labour in excess of what labour would get in compensation, thus equal to the profit collected by capital owners (Marx 1867). This paper argues that instead of framing this profit as purely exploitative, capital owners are being compensated for arranging the infrastructure labour requires to optimally generate economic value, benefitting labour by maximising the economic return they can demand for their time. Having said that, if knowledge labour grows more dependent on paying for new services or technologies to remain economically competitive, these workers will see their net compensation decrease over time. This decrease inherently interferes with the capacity for knowledge labour to optimise its well-being.

Depending on the competitive landscape that will form between the providers of new AI services, these losses experienced by knowledge workers might be exacerbated by the natural monopolies that often form amongst data-dependent technology companies as a result of network-effect economies of scale (Levine and Jain 2023). There is however room for cautious optimism that the struggle between large AI companies is not a winner-takes-most competition. Cautious, not only due to the unpredictability of the technology, but also because it would contradict an established trend for powerful corporations to suppress any competition. In a supposedly leaked internal memo, a Google employee outlines the many advantages open-sourced AI models have over the closely guarded models developed by institutions like Google (Patel and Ahmad 2023). The prime example of this is Stanford's Alpaca, which is a LLM trained using Meta's recently leaked LLM which has since become the backbone of open-source AI developers (Taori, Gulrajani, and Zhang 2023). While Google, Meta and OpenAI have spent billions of dollars to train their models independently, Alpaca cost the students only $600 to train by leveraging the commercially available generative AI models to provide training data to itself, while matching or occasionally exceeding its much more expensive counterparts in several industry benchmarks. In the unlikely scenario that open-source AI development has fundamental advantages, knowledge labour could become the beneficiary of highly competitive, low margin AI service providers.

Section 4: Economic relevance of knowledge labour

Capabilities of knowledge labour

Economic relevance is the dimension of knowledge labour well-being that lies outside the labourer's control. In contrast to internal aspects of well-being, relevance of a workers' skill is determined externally through the aggregate preferences held by society. If these capabilities are deemed irrelevant by society, the people possessing them involuntarily cease to be categorised as labour. It is important to reiterate that people maintain their relevance in society by contributing in ways that are not captured by their economic activity, however the ramifications of economic irrelevance on their financial security and associated social recognition can have a severe detrimental impact on their overall well-being. The general mechanism through which technology generates this irrelevance is by automating tasks into competitive obsolescence (Autor, Levy, and Murnane 2003). When new innovations become capable of automatically completing productive tasks, human labour often finds itself outcompeted by the machine in terms of either cost or consistency.

Generative AI uniquely threatens tasks completed by knowledge labour due to its generative capacity, allowing it to act as a reasoning or creativity machine (OpenAI 2023). Although this ontology is rightfully disputed by philosophers, the pragmatic truth is that the technology is already capable of completing certain reasoning and creativity tasks at a level surpassing average humans. When Boris Eldagsen won the creative category of the 2023 Sony World Photography Awards, he rejected the award upon revealing that the image was generated using AI (Glynn and Vallance 2023). This has since been far from a unique occurrence, with many AI creations winning unsuspecting art competitions (Roose 2022). Naturally, this brings into question the relevance of artistic skill in the creation of art. While society still displays a preference for human created art versus machine generated art (Hong and Curran 2019), the unit economics look very unfavourable for human labour. Where a design agency would traditionally hire 10 artists to produce one prototype every two weeks, the agency could now increase its total output while firing most of its artists, by paying for the services of a capable generative AI system to equip the remaining artists. This is likewise true for labour intensive reasoning based tasks like call-centre work, which is a well paying, highly developed industry in some developing countries.

The highest paid knowledge workers are frequently those whose specific knowledge and expertise is exceedingly rare. While generative AI cannot guarantee consistent reliability, it significantly reduces the barrier to acquire this rare knowledge. This can be beneficial to the education and productivity of some knowledge workers, but the workers whose professional career has been based on the rarity of this knowledge find themselves competing with a chatbot (Sharma 2023). Though generative AI is unlikely to render these experts entirely irrelevant, demand for their services will only be required in highly unique and consequential circumstances. All of this conspires to lower the pricing power knowledge labour has to charge for its expertise. As generative AI continues to progress, increasing segments of knowledge labour will see its economic relevance challenged.

Having said all of that, is generative AI truly capable of making knowledge labour irrelevant? In the context of creativity, while Stable Diffusion generates new images one pixel at a time, it is trained to do so based on its database of human created images. There are currently lawsuits ongoing, filed by artists, claiming that AI is generating images by copying their work (Escalante 2023). It is indeed fair to ask if any creativity is actually taking place, or if the machine is simply mixing and matching existing images without actually creating something new. After all, its parameters are determined by a training set of existing images, and the machine is entirely unaware of any meaning present in its images beyond what their expected loss function value would be in relation to the prompt. However, if pure originality became the benchmark for creativity, many creative workers would in fact fail to be recognised as such. Many famous creatives can be quoted as saying “good artists copy, great artists steal” (Naudus 2021). The main problem with this sentiment is that machines are much more efficient and productive at stealing than humans. The relevance that remains from this for knowledge labour is reserved for the true ‘creatives’, and those who successfully curate art generated by AI.

In a domain where the value of goods or services is as entirely subjective as it is in art, it remains unclear how preferences of society might adjust in the context of mass produced artificial creativity and reasoning. Twentieth century philosophers such as Walter Benjamin and Jean Baudrillard argue that the mass production and mass proliferation of art overwhelms us into becoming alienated from experiencing art or our reality for what it truly is (Cazeaux 2017). However, this does not preclude our aesthetic preferences from changing over time. History shows that our artistic preferences have often evolved alongside technology. Painters used to question the artistic authenticity of photographs, while photographers showed similar distrust toward the prevalent use of Adobe Photoshop. Over time, society has nuanced its preferences such that each method of visual expression has been recognised for its particular contribution to the totality of art. It is conceivable, albeit unlikely, that through the oversaturation of mass produced artificial art, society develops a distinct preference for art with human providence. This is unlikely since it would contradict a century-long trend of increasing mechanisation in the production of art, however generative AI appears set to become the first technology to truly threaten the primacy of human creativity itself. If the trend were to break, the assumption would be that our aesthetic preferences are derived from more than just our immediate sensations, to include contextual information such as an empathic connection to the artist. The recent surge in the popularity of chess is an example of this dynamic. More people than ever watch other humans playing chess matches, despite none of these humans being remotely capable of defeating an AI (Chess.com 2023).

Economic relevance for the use of knowledge labour will also remain in capabilities for which AI is fundamentally ill-suited. One such capability was alluded to earlier, namely the discernment of aesthetic value. The framework does not require a precise ontology of aesthetics, beyond stipulating that taste is a prerequisite to aesthetic value (Kant 1987). Although the technology is capable of rapidly producing more images than any human can in a lifetime, the aesthetic value of these images is subjective and entirely determined through human interpretation. Even if the technology outperforms any human in every artistic category, there will be a person, or team of people who determine which generations produced by the technology correspond best to human sensibilities and therefore optimal for human consumption. At the extreme this could result in personalised entertainment where the person who decides is the same person who consumes, however it's likely that most people would wish to outsource this burden to a person whose aesthetic judgement they admire.

Technology is also fundamentally incapable of being held responsible or accountable for any of the consequences it causes. Relevance for the employment of knowledge labour will remain in tasks where the cost of failure is very high. At the time of writing, the main flaw of generative AI is its unreliability as a result of the probabilistic nature of its output and reasoning, in contrast to traditionally coded, deductive systems. Although the system has demonstrated the capability to diagnose exceedingly rare diseases, very few people would confidently act on this information without receiving a second opinion from a human doctor (Brueck 2023). Likewise, planes have been capable of fully autonomous flight for decades, but it is unlikely we will see pilot-free commercial flights any time soon.

In summary, the adoption of generative AI is expected to have a severe adverse effect on the average economic relevance of employing the capabilities of knowledge labour. Some will lose their jobs entirely, while others will earn less by losing the exclusivity of their knowledge. There are factors that mitigate against the complete obsolescence of knowledge labour, but these factors generally arise in less labour intensive-roles such as oversight.

Application of knowledge labour

Generative AI not only influences the relevance of knowledge labour through direct competition with labour in terms of capabilities, the technology will also generate the availability of new business practices. Although the average relevance of existing knowledge labour capabilities is likely to decrease, generative AI is a technology particularly suited to enable greater participation and flexibility in the knowledge labour force. This section will first examine features of generative AI that might enable greater participation and workforce flexibility, followed by a reflection on the consequent effect on knowledge labour relevance and their respective well-being.

Greater knowledge labour participation would imply that a larger segment of the population has invested resources to acquire the necessary skills to be employable as effective knowledge labour. This paper has already alluded several times to the feature of generative AI enabling this, namely its capacity to reduce the resources required to acquire new knowledge. This is particularly significant for people aspiring to become a knowledge worker, but do not have the means to access the required training. Since knowledge workers generally earn more on average, owing to their training, generative AI could function as an enabler of upward mobility

(Glassdoor 2022). This mechanism could further be amplified by the apparent characteristic of generative AI to be most effective at improving the performance of the worst performing employees. Early studies conducted in workplace environments find that generative AI makes all workers more productive, but that low performing employees benefited much more than high performers (Noy and Zhang 2023). For businesses, this means hiring new workers is a lower risk proposition, since the cost of poor performance is lower as a result of a higher performance floor. Generative AI also makes it more attractive for businesses to hire labour because of the greater level of productivity each worker can contribute when incorporated with the technology. Knowledge labour participation could also somewhat increase purely on the basis of long neglected expertise gaining renewed technological applicability. Given LLMs are operated through natural language, humanities expertise is suddenly being valued for its potential technological contributions (Roemmele 2023). The technology is essentially designed to mimic humans to the best of its ability, so linguists can help identify which words to use in the prompt, psychologists can interrogate the system to reveal the depth of its reasoning, while the critical eye of a philosopher can validate whether the generated responses are logically sound.

Other than potentially increasing the size of the knowledge labour force, generative AI also has the capacity to make existing knowledge labour more flexible. Broadly speaking, the features that enable greater levels of participation also make it easier for specialised knowledge labour to become economically relevant in a wider range of contexts. By reducing the barrier to specialised knowledge, less specialised, lower paid knowledge workers could be empowered with more responsibility. For example with the aid of an AI assistant trained on all of society's medical knowledge, nurses could be empowered to act with more discretion. This could increase the economic relevance of nurses, while resulting in expedited medical care for patients. Alternatively, by automating routine tasks such as data entry and financial reporting, financial analysts might be granted the latitude to analyse financial factors that balance sheets omit. In contrast to the ever-present fears of technological task displacement, this would be an example of task reinstatement (Acemoglu and Restrepo 2019). Reinstatement occurs when the automation of automatable tasks leads to greater freedom for labour to fulfil the ultimate aim of the job. For the financial analyst this ultimate aim could be to provide the client with an exhaustive yet comprehensible account of a business's risks and opportunities. If generative AI leads to an increasing flexibility of knowledge labour, it would signal the relevance of inherently human capabilities exceeding the relevance of formally trained capabilities.

The concept of flexibility is fundamentally related to the prospect of change. As a dynamic categorization of people, it risks becoming meaningless to assess the relevance and well-being of knowledge labour when its constitution faces radical change. The knowledge workers who have built a successful career prior to the adoption of generative AI will watch their familiar work practices around them change, as well as risk becoming irrelevant unless they commit to change themselves. Change is inherently uncomfortable, and it is sensible to assume that these knowledge workers would experience a decline in well-being. On the other hand, by effectively democratising the acquisition of knowledge, generative AI might make knowledge labour a more relevant occupational choice for underprivileged individuals. This would lead to greater economic relevance for those individuals and it is in turn sensible to assume that this would be beneficial for the well-being of knowledge labour. On balance, it is challenging to conclusively determine how the use of generative AI in the application of knowledge labour will affect its relevance and well-being. Some knowledge workers will become more economically relevant alongside improved well-being, but it remains unclear what fraction of knowledge labourers these workers will represent.

Section 5: Evaluation

Implications of the framework

The analysis of generative AI through the framework established by this paper has yielded several implications. Generative AI is shown to assist the self-development of knowledge labour by reducing the barriers to specialised knowledge and enabling personalised education. Additionally, generative AI is shown to increase an individual's capacity for self-expression by enabling new means of expression, and acting as a source of creative inspiration. Insofar as knowledge workers value self-improvement and self-expression, the technology is shown to increase the well-being of knowledge labour. Generative AI will provide knowledge labour with new economic opportunities, whether through newly available business opportunities or by increasing their personal productivity. While technological capital is likely to encroach on the productive output of knowledge labour, widespread adoption would signal that this cost is outweighed by the productivity gain enabled by the technology. In sum, the framework indicates that generative AI promotes the well-being of knowledge labour by enhancing their capacity to flourish.

Regarding the economic relevance of knowledge labour capabilities, generative AI threatens to replace knowledge labour in some cognitive and creative tasks, while reducing the pricing power of highly specialised knowledge labour. There exist exclusively human capabilities that mitigate against obsolescence, but having to directly compete with generative AI will reduce the economic relevance of knowledge labour. Conversely, the technology could promote this relevance by enabling new approaches to the application of knowledge labour. Knowledge labour force participation would increase stemming from a higher productivity floor of labour, while the greater flexibility of knowledge workers raises the possibility for meaningful task-reinstatement. These new approaches imply changing work practices, which has a severe adverse effect on the relevance and well-being of existing knowledge workers. These divergent experiences and changes in group composition leave an assessment of knowledge labour relevance, resulting from business practices enabled by generative AI, indeterminable. On balance, the framework shows generative AI impedes the well-being of knowledge labour by decreasing their economic relevance.

Bearing in mind the intrinsic unpredictability associated with any novel technology, the final implication following from the framework established in this paper is: *Generative AI reduces the economic relevance of knowledge labour while enhancing the capacity of knowledge workers to flourish.*

Discussion

The implication that results from the framework when regarding the effect generative AI will have on well-being is contradictory, inviting the possibility for discussion. This paper has assumed a largely deterministic relationship between economic relevance, the employment potential of knowledge labour and the well-being of the underlying human. Having now considered the potential consequences of generative AI, this relationship can break down in two significant ways.

As the example of commercial airline pilots shows, technology might make almost all the capabilities of labour irrelevant, yet the benefits derived from the remaining capabilities outweigh the cost of employing a pilot. However, the economic relevance of labour capabilities can be completely disregarded by industries when isolated from market dynamics. Governments are not exclusively motivated by productivity growth, they also aim for minimal unemployment figures. As an institution tasked with representing the political will of its people, governments will likewise modulate their attitude towards generative AI in accordance with the political needs of its constituents. Policy makers could therefore be encouraged to demand arbitrarily restrictive safety standards for the adoption of the technology in workplaces, such that its use is effectively disallowed and consequently protecting the employment of labour in spite of their economic irrelevance. Alternatively, policy makers could use this legislative power to protect forms of labour particularly exploited and marginalised by the adoption of generative AI. The well-being of society is primary, so the economic relevance of labour should ironically be irrelevant to their employment potential if the well-being of society benefits from their mere employment. Consider the thesis you have been reading. Viewed exclusively in terms of productivity and efficiency, generative AI could make a human writer economically irrelevant for this task, though the loss of human employment in this activity would be detrimental for the well-being of society. Not only does it undermine the signalling and selection mechanisms underlying formal education, the well-being of the student is eroded by removing an opportunity for self-development in the face of a purposefully difficult yet attainable challenge. As such, the relationship between the economic relevance of labour and their employment possibilities is not entirely deterministic, since society's preference can modify the effective relevance.

Likewise, this paper treats the loss of employment as logically equivalent to a loss in well-being. This is appropriate when considering a large subset of the population defined by their economic relationship, but the logic breaks down on the level of a human being, particularly in a society capable of moulding its preferences to maximise overall well-being. Employment is not a fundamental prerequisite to human well-being, instead it acts as a means to access some prerequisites to well-being: self-expression, social inclusion and financial independence. Those who perform their labour in order to express themselves will indeed experience a fundamental loss of well-being upon losing employment. On the other hand, social inclusion and the distribution of economic resources are expressions of society's preferences. It just so happens that in much of the developed world, people are primarily identified by their relationship to the economy and unemployment is socially stigmatised. If labour finds itself becoming obsolete due to an amazingly productive technological innovation, society should be capable of distributing this additional wealth such that the unemployed retain some semblance of financial freedom while leaving the incentive to innovate intact. Additionally, as technology drives more labour into obsolescence, the well-being of society would benefit from greater social inclusion of the voluntarily or involuntarily unemployed. Naturally, this view merely represents a simplistic economic idealisation of factors that ought to determine society's preferred distribution of economic resources. In practice, this distribution is largely determined by the vested interests of powerful economic participants rationally maximising their own self-interest. This does not, however, preclude the normative and rational claim that the collective well-being of society would benefit from a distribution of economic resources that is somewhat agnostic towards the precise economic contributions of an individual. The value human beings contribute to society is not contained by their relationship to the economy.

Conclusion

In conclusion, this paper evaluated the consequences of generative AI through the construction of a normative framework that accounts for the well-being of knowledge labour. The technology is shown to increase the well-being of knowledge labour by promoting their capacity for self-development, self-expression and new economic opportunities. On the other hand, the economic relevance and well-being of knowledge labour is challenged by generative AI through direct competition in terms of capabilities, and disruption of existing business particles. The framework thereby arrives at the following implication, and thesis of this paper: *Generative AI reduces the economic relevance of knowledge labour while enhancing the capacity of knowledge workers to flourish.*

While this thesis is shown to be the appropriate implication resulting from the framework as constructed, it leaves the effect of relevance on well-being somewhat unresolved. This is due to the composition of knowledge labour changing as a result of generative AI, but also because the relationship between economic relevance, employment potential, and well-being is not entirely deterministic, and susceptible to changes in societal preferences. Generative AI is a tool with an incredible capacity to promote well-being, even for the group of people most economically threatened by its adoption. It is subsequently the responsibility of society to express preferences that optimise our collective well-being.

Bibliography

Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3-30. 10.1257/jep.33.2.3.

Altman, Sam. 2021. “Moore's Law for Everything.” Samaltman.com. https://moores.samaltman.com/.

Arrow, Kenneth J. 1963. “Uncertainty and the Welfare Economics of Medical Care.” *The American Economic Review* 53, no. 5 (December): 941-973.

Arya, Nisha. 2023. “Baby AGI: The Birth of a Fully Autonomous AI.” KDnuggets. https://www.kdnuggets.com/2023/04/baby-agi-birth-fully-autonomous-ai.html.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *The Quarterly Journal of Economics* 118, no. 4 (November): 1279-1333. https://doi.org/10.1162/003355303322552801.

Baumol, William J. 2012. *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't*. New Haven: Yale University Press.

Bentham, Jeremy. 1789. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press. http://dx.doi.org/10.1093/oseo/instance.00077240.

Brueck, Hilary. 2023. “Can ChatGPT Be a Doctor? Bot Passes Medical Exam, Diagnoses Conditions.” Insider. https://www.insider.com/chatgpt-passes-medical-exam-diagnoses-rare-condition-2023-4.

Brynjolfsson, Erik. 2023. “Your job is (probably) safe from artificial intelligence.” The Economist. https://www.economist.com/finance-and-economics/2023/05/07/your-job-is-probably-safe-from-artificial-intelligence.

Cazeaux, Clive. 2017. *The Continental Aesthetics Reader*. New York: Taylor & Francis Group.

Chess.com. 2023. “Chess Is Booming! And Our Servers Are Struggling.” Chess.com. https://www.chess.com/blog/CHESScom/chess-is-booming-and-our-servers-are-struggling.

Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values*. N.p.: WW Norton.

De Cremer, David, Nicola M. Bianzino, and Ben Falk. 2023. "How Generative AI Could Disrupt Creative Work." Harvard Business Review. https://hbr.org/2023/04/how-generative-ai-could-disrupt-creative-work.

Dorsey, Dale. 2021. *A Theory of Prudence*. Oxford: Oxford University Press.

Edwards, Benj. 2022. "Artist receives first known US copyright registration for latent diffusion AI art." Ars Technica. https://arstechnica.com/information-technology/2022/09/artist-receives-first-known-us-copyright-registration-for-generative-ai-art/.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." (March). https://doi.org/10.48550/arXiv.2303.10130.

Escalante, Shanti. 2023. "Artists Sue Midjourney, Stability AI — The Case Could Change Art." ARTnews.com. https://www.artnews.com/art-in-america/features/midjourney-ai-art-image-generators-lawsuit-1234665579/.

Fried, Ina. 2023. "Sal Khan explains why GPT-4 is ready to be a tutor." Axios. https://www.axios.com/2023/04/07/sal-khan-chatgpt-gpt4-tutor.

Glassdoor. 2022. "Salary: Knowledge Worker in United States 2023." Glassdoor. https://www.glassdoor.com/Salaries/knowledge-worker-salary-SRCH_KO0,16.htm.

Glynn, Paul, and Chris Vallance. 2023. "Sony World Photography Award 2023: Winner refuses award after revealing AI creation." BBC. https://www.bbc.com/news/entertainment-arts-65296763.

Goldman Sachs. 2023. "Generative AI Could Raise Global GDP by 7%." Goldman Sachs. https://www.goldmansachs.com/intelligence/pages/generative-ai-could-raise-global-gdp-by-7-percent.html.

Google. 2023. "PaLM 2 Technical Report." (May). https://doi.org/10.48550/arXiv.2305.10403.

Greene, Jenna. 2022. "Will ChatGPT make lawyers obsolete? (Hint: be afraid)." Reuters. https://www.reuters.com/legal/transactional/will-chatgpt-make-lawyers-obsolete-hint-be-afraid-2022-12-09/.

Hong, Joo-Wha, and Nathaniel M. Curran. 2019. "Artificial Intelligence, Artists, and Art: Attitudes Toward Artwork Produced By Humans vs. Artificial Intelligence." *ACM*

Transactions on Multimedia Computing, Communications, and Applications 15, no. 2 (July): 1-16. https://doi.org/10.1145/3326337.

Jerome, Harry. 1934. *Mechanization in Industry*. New York: National Bureau of Economic Research.

Jovanovic, Boyan, and Peter Rousseau. 2005. "Chapter 18." In *Handbook of Economic Growth*, 1181-1224. Part B ed. Vol. 1. Amsterdam: Elsevier.

Kalliamvakou, Eirini. 2022. "Research: quantifying GitHub Copilot's impact on developer productivity and happiness." The GitHub Blog. https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/.

Kant, Immanuel. 1987. *Critique of judgment*. Translated by James C. Meredith. Cambridge: Clarendon Press.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107, no. 1 (February): 35-78. https://doi.org/10.2307/2118323.

Keynes, John M. 2010. "Economic Possibilities for Our Grandchildren." In *Essays in Persuasion*, 321-332. London: Palgrave Macmillan.

Krugman, Paul. 2013. "Sympathy for the Luddites." *The New York Times*, June 13, 2013. https://www.nytimes.com/2013/06/14/opinion/krugman-sympathy-for-the-luddites.html?_r=0.

Kumichev, Gleb. 2022. "The Inductive Bias of ML Models, and Why You Should Care About It." Towards Data Science. https://towardsdatascience.com/the-inductive-bias-of-ml-models-and-why-you-should-care-about-it-979fe02a1a56.

Levine, Sheen S., and Dinkar Jain. 2023. "How Network Effects Make AI Smarter." Harvard Business Review. https://hbr.org/2023/03/how-network-effects-make-ai-smarter.

Manyika, James, Jake Silberg, and Brittany Presten. 2019. "What Do We Do About the Biases in AI?" Harvard Business Review. https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai.

Marx, Karl. 1867. "The Production of Absolute Surplus-Value." In *Capital: The Process of Production of Capital*, 125-130. Moscow: Progress Publishers.

Merchant, Brian. 2023. “Merchant: The writers' strike and the rebellion against AI.” Los Angeles Times. https://www.latimes.com/business/technology/story/2023-05-11/column-the-writers-strike-is-only-the-beginning-a-rebellion-against-ai-is-underway.

Naudus, Philip S. 2021. “Steve Jobs: “Good Artists Copy, Great Artists Steal” | by Philip S. Naudus.” The Writing Cooperative. https://writingcooperative.com/steve-jobs-good-artists-copy-great-artists-steal-4fc6593ac09.

Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” (March).

Nussbaum, Martha C. 2006. *Frontiers of Justice: Disability, Nationality, Species Membership*. Cambridge: Harvard University Press.

OpenAI. 2021. “DALL·E: Creating images from text.” OpenAI. https://openai.com/research/dall-e.

OpenAI. 2023. “GPT-4 Technical Report.” (March). https://doi.org/10.48550/arXiv.2303.08774.

Ortiz, Sabrina, and Alyson Windsor. 2023. “What is Auto-GPT? Everything to know about the next powerful AI tool.” ZDNET. https://www.zdnet.com/article/what-is-auto-gpt-everything-to-know-about-the-next-powerful-ai-tool/.

Patel, Dylan, and Afzal Ahmad. 2023. “Google "We Have No Moat, And Neither Does OpenAI."” SemiAnalysis. https://www.semianalysis.com/p/google-we-have-no-moat-and-neither.

Reinhardt, Wolfgang, Benedikt Schmidt, Peter Sloep, and Hendrik Drachsler. 2011. “Knowledge Worker Roles and Actions—Results of Two Empirical Studies.” *Knowledge and Process Management* 18, no. 3 (August): 150-174. https://doi.org/10.1002/kpm.378.

Roemmele, Brian. 2023. “ChatGPT and the Dawn of Computerized Hyper-Intelligence | The Jordan B. Peterson Podcast • Podcast Notes.” Podcast Notes. https://podcastnotes.org/jordan-b-peterson-podcast/brian-roemmele-chatgpt-and-the-dawn-of-computerized-hyper-intelligence-the-jordan-b-peterson-podcast/.

Roose, Kevin. 2022. “An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy.” The New York Times. https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html.

Rouhiainen, Lasse. 2019. “How AI and Data Could Personalize Higher Education.” Harvard Business Review. https://hbr.org/2019/10/how-ai-and-data-could-personalize-higher-education.

Schinkel, Anders. 2008. “Martha Nussbaum on Animal Rights.” *Ethics and the Environment* 13 (1): 41-69. https://www.jstor.org/stable/40339148.

Schmidt, Eric, and Dan Wang. 2023. “Eric Schmidt: Why Technology Will Define the Future of Geopolitics.” Foreign Affairs. https://www.foreignaffairs.com/united-states/eric-schmidt-innovation-power-technology-geopolitics.

Sen, Amartya. 1985. *Commodities and Capabilities*. Amsterdam: North-Holland.

Sen, Amartya. 1992. *Inequality Reexamined*. Oxford: Clarendon Press.

Sen, Amartya. 2002. *Rationality and Freedom*. Cambridge: Harvard University Press.

Sharma, Divyanshi. 2023. “ChatGPT as management consultant? Bain says it will use OpenAI tool for its analysis.” India Today. https://www.indiatoday.in/technology/news/story/bain-says-it-will-use-chatgpt-ai-for-its-analysis-management-consultant-jobs-may-get-replaced-by-ai-2338153-2023-02-22.

Spataro, Jared. 2023. “Introducing Microsoft 365 Copilot – your copilot for work - The Official Microsoft Blog.” The Official Microsoft Blog. https://blogs.microsoft.com/blog/2023/03/16/introducing-microsoft-365-copilot-your-copilot-for-work/.

Susser, Mervyn, and Zena Stein. 2009. “Germ Theory, Infection, and Bacteriology.” In *Eras in Epidemiology: The Evolution of Ideas*, 107-122. New York: Oxford University Press, USA. https://doi.org/10.1093/acprof:oso/9780195300666.003.0010.

Taori, Rohan, Ishaan Gulrajani, and Tianyi Zhang. 2023. “Alpaca: A Strong, Replicable Instruction-Following Model.” Stanford CRFM. https://crfm.stanford.edu/2023/03/13/alpaca.html.