

NBER WORKING PAPER SERIES

PREPARING FOR THE (NON-EXISTENT?) FUTURE OF WORK

Anton Korinek
Megan Juelfs

Working Paper 30172
<http://www.nber.org/papers/w30172>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2022

This article is an expanded version of a chapter prepared for the Oxford Handbook of AI Governance. The authors would like to acknowledge financial support from the Institute for Business in Society at Darden and helpful research assistance by Eleanor Giordano and Marcella Cartledge. We thank Daron Acemoglu, Zach Bethune, Justin Bullock, Bill Gale, Katya Klinova, Lee Lockwood, Joseph Stiglitz, Risto Uuk and the participants of an Oxford Handbook of AI Governance conference, the UVA Inequality Working Group, seminars at Tilburg and the University of Vienna, and a workshop at the Centre for the Governance of AI for their thoughtful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Anton Korinek and Megan Juelfs. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Preparing for the (Non-Existent?) Future of Work
Anton Korinek and Megan Juelfs
NBER Working Paper No. 30172
June 2022
JEL No. E6,J2,O3

ABSTRACT

This paper considers the labor market and distributional implications of a scenario of ever-more-intelligent autonomous machines that substitute for human labor and drive down wages. We lay out three concerns arising from such a scenario and evaluate recent predictions and objections to these concerns. Then we analyze how a utilitarian social planner would allocate work and income if these concerns start to materialize. As the income produced by autonomous machines rises and the value of labor declines, a utilitarian planner finds it optimal to phase out work, beginning with workers who have low labor productivity and job satisfaction, since they have comparative advantage in enjoying leisure. This is in stark contrast to welfare systems that force individuals with low labor productivity to work. If there are significant wage declines, avoiding mass misery will require other ways of distributing income than labor markets, whether via sufficiently well-distributed capital ownership or via benefits. Recipients could still engage in work for its own sake if they enjoy work amenities such as structure, purpose and meaning. If work gives rise to positive externalities such as social connections or political stability, or if individuals undervalue the benefits of work because of internalities, then a social planner would incentivize work. However, in the long run, the planner might be able to achieve a higher level of social welfare by adopting alternative ways of providing these benefits.

Anton Korinek
Department of Economics
University of Virginia
Monroe Hall 246
248 McCormick Rd
Charlottesville, VA 22904
and NBER
anton@korinek.com

Megan Juelfs
Office of Research Services
Darden School of Business
University of Virginia
Charlottesville, VA 22906
JuelfsM@darden.virginia.edu

1 Introduction

In modern societies, work is both the main source of income for working-age people and the main pursuit on which they spend their time. However, there are concerns that advances in artificial intelligence and related technologies may substitute for a growing fraction of workers, presenting significant challenges for the future of work. This paper analyzes how to optimally allocate – or phase out – labor and distribute income from the perspective of a utilitarian planner in a world in which intelligent machines may increasingly become substitutes for human labor and drive down wages.

The paper starts by laying out three distinct concerns about the future of work and analyzing them from an economic perspective. The first concern is that technological progress may be labor-saving and reduce wages. Since the beginning of the Industrial Revolution, technological progress has benefited capital and labor in roughly equal proportions, giving rise to large increases in real wages. However, recent advances in automation have substituted for unskilled labor and have saved on labor in some segments of the labor market, reducing real wages for the affected workers. Future technological progress may reduce the wages of an ever-larger fraction of the population.

The second concern is that machines may become perfect substitutes for labor. Throughout our history, technological progress has made production more efficient, but human labor has always remained essential for the production of output. The central role of labor in production is one of the main reasons why progress has led to increasing wages. By contrast, if advanced AI and robotics can fully substitute for any type of human labor in the future, then labor will no longer be essential. We summarize predictions that labor may become perfectly substitutable within the current century. This may accelerate labor-saving progress and may mark the beginning of the end of the Age of Labor. On the positive side, if the scarce factor labor is no longer essential for production, economic growth may increase significantly.

The third concern is that labor may become economically redundant, which we define as the cost of machines that substitute for human labor falling below the cost of human subsistence. In such a world, humans could no longer survive based on earning competitive market wages. Work would become obsolete and would cease to play the central role that it currently plays in our society.

We evaluate a number of objections to the three concerns we described. We discuss the notion, put forward by some, that human labor may inherently be superior to anything that machines could ever do. We describe how the articulated concerns relate to the lump-of-labor fallacy, to historical experience, and to the creation of new jobs when old jobs are displaced. We assess whether human demand is necessary for the functioning of the economy if labor is phased out. We identify a role for “nostalgic” jobs in which humans are hired merely for the fact of being human, even if machines could perform their work more cheaply and more effectively. Finally, we discuss how concerns about the economic redundancy of labor relate to the concepts of comparative and absolute advantage in international trade.

In the following section, we analyze how to optimally allocate work and income to

maximize utilitarian social welfare as a function of the prevailing state of technology. For given labor productivity and non-labor income, an individual worker can find herself in one of three different regimes. In the first regime, the worker's productivity and non-labor income are too low and she perishes. In the second regime, which arises for higher levels of labor productivity but sufficiently low non-labor income, the agent works, first out of need, then out of choice. As her non-labor income rises, it is optimal for her to spend more and more time on leisure rather than work. She enters the third regime if technological progress increases her non-labor income or reduces the marginal product of her labor sufficiently so that it becomes optimal for her to no longer work.

In an economy with multiple agents who differ in their labor productivity, a utilitarian planner finds it optimal to allocate work to individuals solely based on their labor productivity, with more work allocated to individuals with greater productivity. If output in the economy becomes sufficiently high or the marginal product of labor declines sufficiently, it becomes optimal to phase out work, beginning with workers who have low labor productivity and job satisfaction, since they have comparative advantage in enjoying leisure. Once the income that is produced by autonomous machines is sufficiently high, it becomes optimal for humans to stop working altogether.

When we account for job amenities that workers value in addition to their pay, for example identity, meaning, or structure, then the picture becomes more nuanced: for workers who value such amenities sufficiently, it may be desirable to continue to work even if their labor productivity and their competitive market wages fall to zero. Work may also give rise to externalities and public goods, such as social connections or political stability. Moreover, it may entail internalities, i.e. workers may derive benefits (or costs) from work that they do not rationally internalize, for example from the structure to their daily lives that work provides. Such instances may lead to suboptimal decisions by workers and firms, and public policy interventions may increase utilitarian welfare.

Next we discuss how a utilitarian planner would reform prevailing economic institutions to allocate work and income when technological progress reduces wages and when work becomes economically redundant. At present, one of the main institutions to allocate both is the market. However, market failures imply that individuals have very limited access to long-term insurance markets that would protect them against adverse shocks, including against the risk of labor-saving technological progress and economic redundancy. Social insurance systems may fill the gap and provide individuals with some protection. Moreover, as our economy grows wealthier or as market wages decline, social insurance may provide sufficient income to individuals so they are not forced into work that has low productivity and therefore creates little social value from a utilitarian perspective, while using up their valuable leisure time.

When progress is labor-saving, a utilitarian planner would engage in greater redistribution. Current mechanisms of social insurance frequently condition benefits on work so as to provide incentives to recipients. A utilitarian planner would unlink the two and avoid providing excessive work incentives – she finds it optimal for individuals to stop working if the marginal product of labor declines sufficiently and/or the

economy becomes sufficiently prosperous. We observe that there may be a multiplier effect for transfers given to workers: by reducing labor supply, transfers may both increase equilibrium wages and enable workers to demand better work amenities, which increases the incomes of workers and further raises their welfare. Moreover, as long as labor is an important source of income, a utilitarian planner actively steers technological progress in directions that are labor-using and increase competitive market wages so as to reduce the need for costly redistribution.

If labor becomes economically redundant, individuals will no longer be able to survive based on their labor income alone, and other sources of income – whether from sufficiently well-distributed capital ownership or benefit payments – are critical to avoid mass misery and any potential political instability that may result from it. Since machines that make labor redundant would be able to produce unprecedented levels of output, there would be significantly more resources available, making it easier to raise revenue that can be redistributed.

Since the vast majority of the population would require income support in such a world, the advantages of targeting benefits according to individual needs would decline and a utilitarian planner would provide a basic income for everyone. She would find an unconditional Universal Basic Income (UBI) preferable to benefits conditioned on work. Recipients of such a UBI could still engage in work for its own sake if they enjoy work amenities such as structure, purpose and meaning.

The planner would find it optimal to implement policies that actively encourage work only if work gives rise to positive externalities such as social connections or political stability, or if individuals undervalue the benefits of work because of internalities. However, these externalities and internalities may be difficult for society to identify and agree upon. In the long run, we conjecture that there would likely be alternative and more efficient ways of generating any positive externalities that traditionally derived from work.

Finally, we also observe that systems of taxation will have to adapt significantly if the role of labor in the economy declines. Taxes would have to be raised increasingly on factors other than labor, for example via Pigouvian taxes on activities that generate externalities, or via taxes on inelastic factors such as land that are not distorted by taxation. Moreover, there may be equity reasons to tax any factors that benefit from technological progress at the expense of labor.

Structure of the Paper In the ensuing section, we formally describe and evaluate three concerns about the future of work as well as objections to such concerns. Section 3 develops a benchmark to analyze how technological advances affect the optimal allocation of income and work from the perspective of a utilitarian planner, accounting for the potential role of work amenities as well as for externalities and internalities from work. Section 4 evaluates how economic institutions can be adjusted to reflect the allocations of a utilitarian planner, focusing on the role of markets and social insurance mechanisms. Section 5 concludes.




Figure 1: Varieties of technological progress

2 Evaluating Concerns About the End of Labor

This section formalizes and evaluates three concerns that capture the economic implications of a scenario of ever-more-intelligent autonomous machines that substitute for human labor and drive down wages. The first concern is that technological progress may be labor-saving and reduce wages in absolute terms. The second concern, frequently articulated by technologists, is that machines may become perfect substitutes for labor. We discuss recent technological developments and summarize predictions regarding this concern. Next we analyze under what conditions this leads to a third concern, that labor becomes economically redundant. We also evaluate the main objections to predictions about the redundancy of labor.

2.1 Labor-Saving Progress

Since the onset of the Industrial Revolution, technological progress has led to broad-based economic growth that has, on average, benefitted everybody in the economy. In particular, output and wages grew at approximately the same rate over that period – a phenomenon that economists call neutral technological progress.¹ Figure 1(a) illustrates neutral technological progress in a stylized diagram, in which output is depicted on the horizontal axis and labor compensation (wages w times employment ℓ) on the vertical axis. Neutral technological progress implies that labor compensation grows in lockstep with output so that the share of labor remains constant.

However, as Hicks (1932) already noted, technological progress need not always be

¹In technical terms, economists call technological progress Hicks-neutral if labor and capital benefit proportionately for given factor supplies (or for constant relative factor supplies) – named in honor of Hicks (1932) who first characterized the phenomenon. They call technological progress Harrod-neutral if labor and capital benefit proportionately after the capital stock has adjusted to the new technology (or for a constant capital/output ratio). We are referring to the latter concept since we are interested in the equilibrium effects of technological progress, which typically involve an adjustment in the capital stock.

neutral. It is well possible that future advances in technology substitute for human labor and benefit labor less than other factors, or that they even lead to declines in competitive market wages. Panel (b) illustrates technological progress that is biased against labor: after the kink, labor compensation still grows, but at a smaller rate than output, reducing labor’s share of income. Karabarbounis and Neiman (2014) document that this has been happening in the majority of countries around the world in recent decades.

Panel (c) of the figure depicts the possibility that technological progress may be labor-saving in absolute terms, so that the compensation of labor declines even as overall economic output rises. We formalize this concern as follows:

Concern 1 (*Labor-Saving Progress in Absolute Terms*). Technological progress reduces the demand for labor at given prices.

If the demand for labor declines at given prices and wages, then the equilibrium competitive wage will decline in absolute terms (as long as the supply of labor is constant or upward-sloping).

Autor (2019) documents that large categories of the US workforce, esp. lesser-educated workers, have already experienced stagnating or declining real wages in recent decades. Acemoglu and Restrepo (2021) find that the majority of these declines was due to automation. What has been the fate of unskilled lower-wage workers in recent decades may turn out to be the fate of high-skilled and high-wage workers in future decades. Korinek and Stiglitz (2021) provide several economic models that exemplify that it is well possible that technological advances may reduce competitive market wages in absolute terms even though output increases, and that this may hold even if automation is accompanied by the accumulation of additional capital. This captures what Wassily Leontief (1983) described as follows: *“the role of humans as the most important factor of production is bound to diminish – in the same way that the role of horses in agricultural production was first diminished and then eliminated by the introduction of tractors.”*

2.2 Perfect Substitutability of Labor

Many predictions about the redundancy of labor are based on the premise that the human brain is at its core a computing device that processes information by transforming inputs into outputs. This premise makes it plausible that advances in hardware and software may catapult the computing capabilities of machines to the point where they may rival the human brain. When combined with sufficiently advanced sensors and actuators, machines could then perform any kind of work that humans can perform.

To formalize this concern, consider a production function $F(\ell, m)$ that employs labor ℓ and other factors, represented by machines m here. (Both labor ℓ and the other factors m could be interpreted either as scalars or as vectors of different types of labor or machines.) The formal condition for the substitutability of labor can then be denoted as follows:

Concern 2 (*Perfect Substitutability of Labor*). There exist machines that can substitute for any type of labor $\ell > 0$ in production

$$\forall \ell > 0, m \exists \Delta m \text{ s.t. } F(0, m + \Delta m) \geq F(\ell, m)$$

When this condition is satisfied, labor is no longer an essential input for production – given enough machines, production at any desired scale can take place without labor.

Current State of Technology At present, artificial intelligence is clearly nowhere near possessing the ability to perform all human work. Humans cannot be substituted for and are still essential both in simple non-routine manual tasks and in higher-order cognitive tasks that require an understanding of the world, planning, and crucially, social intelligence. However, over the past decade, advances in deep learning have enabled artificial intelligence and AI-powered robots to perform a growing number of less-structured and higher-level cognitive tasks.²

Technological Predictions In terms of sheer computing power, the world’s most advanced computers are already roughly on par or superior to the human brain. One common measure of computing power are floating point operations per second (flops), corresponding to how many arithmetic operations on real numbers a computer can perform per second. Carlsmith (2020) estimates that the computing power of the human brain can be replicated with about 10^{15} flops, given the right software. At the time of writing, Fugaku, the world’s top publicly known supercomputer, was able to reach a peak performance exceeding 10^{18} flops, easily surpassing this estimate, albeit the system was reported to cost more than \$1bn. And computing capacity is expected to continue to grow for the foreseeable future.³

The models underlying cutting-edge AI applications are also experiencing rapid progress. Given the described hardware capabilities, the absence of a human-level general AI that we know of suggests that advances in software are somewhat lagging behind advances in hardware. However, the field is experiencing a rapid inflow of talent and funds (AI Index Report, 2021), suggesting that progress will continue unabated.

²AI can now accomplish many economically useful tasks at levels of accuracy that either already are or soon will be super-human. Speech recognition and image recognition tasks can already be performed at human levels by state-of-the-art AI. DeepMind’s AlphaFold can predict how the sequence of amino acids encoded in our genes will fold into proteins, opening new avenues for drug discovery and bioengineering. Transformers such as OpenAI’s GPT-3 can produce high-quality text that is frequently indistinguishable from human-written text and even displays creativity. And machine learning has also led to rapid advances in robot dexterity, surpassing humans in many applications. At present, many of these technologies are at the lab stage, with commercial applications that share their benefits with society at large only at the beginning.

³For more than sixty years now, the processing capabilities of cutting-edge processors have roughly doubled every two years, as predicted by a generalized version of Moore (1965)’s Law. Although there is some disagreement about the exact pace of future progress, few question that the growth in computing capabilities will continue to be exponential for some time to come.

The futurist Ray Kurzweil (2005) predicts that AI will achieve human levels of intelligence in 2029 and has reaffirmed his prediction in recent years. Bostrom (2014) and Grace et al. (2018) conduct surveys of AI experts on when human-level artificial intelligence may be reached. Bostrom reports a median prediction of 2040 in a sample that included many futurists. Grace et al. report that a broad sample of AI researchers assign a 50% chance that humans will be technologically redundant by the early 2060s.

2.3 Economic Redundancy of Labor

Perfect substitutability does not necessarily imply economic redundancy of labor since it does not consider how costly it is to replace human work with machines. If substituting for all human labor is technologically possible, but it would cost many times more than human workers at current market wages, then it is not economically efficient to do so.

The cost at which machines can perform a given human job imposes a ceiling, i.e. an upper bound on the competitive market wage of humans performing the job. As technology advances and machines become more and more efficient, this ceiling – and by implication the market wages of humans – will decline. At first, jobs that are easily automated are affected, leading humans to switch to work that is more difficult to automate and therefore pays higher wages, as they have been doing for centuries. However, in the past, there were always jobs left that only humans could perform. This will no longer be the case if machines become perfect substitutes for labor. And as machines continue to become more efficient, wages may fall below the subsistence cost necessary to support a human worker so that human labor is no longer economically worth its keep.

To formalize this concern, assume that we normalize the units of machines m such that their user cost is unity, and that the subsistence consumption of a worker is given by c_0 . Then we can define the economic redundancy of labor in the production function $F(\ell, m)$ as follows:

Concern 3 (*Strong Economic Redundancy of Labor*). Machines are able to perform any economically valuable task cheaper than humans, valued at their subsistence cost,

$$\forall \ell > 0, m \exists \Delta m \text{ s.t. } F(\ell, m) \leq F(0, m + \Delta m), \Delta m < c_0 \ell$$

Concerns 2 and 3 reflect the way in which the potential future redundancy of labor is frequently framed in technology circles. Observe that perfect substitutability of labor (Concern 2) is a necessary condition for strong economic redundancy of labor (Concern 3) but not a sufficient one. However, if Concern 2 is satisfied and a broad version of Moore’s Law holds whereby the cost of human-replacing machines falls by half roughly every two years, then even very expensive technologies to substitute for human labor may become affordable relatively quickly: if the cost of substituting for one unit of human labor in its most expensive use is c_m , then it will take $2 \log_2(c_m/c_0)$ years for

Concern 3 to materialize. For example, for $c_m = \$1\text{bn}$ and $c_0 = \$1000$, the process would take about 40 years.

If Concern 3 materializes, humanity would no longer add economic value – humans would require more economic resources than they are able to produce, and labor would be obsolete in the sense of becoming a dominated technology. Moreover, humans could no longer survive based on their competitive market wages alone – a stark departure from the way our societies have been organized since the onset of the Industrial Revolution – and we would have to choose between mass misery or providing a basic subsistence income, as we will explore in more detail below.

An analogous economic outcome may result even from a weaker condition. Assume that there are still some jobs that can only be done by humans (violating Concern 2) or in which humans are more cost-effective than machines (violating Concern 3). If the demand for those jobs is insufficient, then the market-clearing price of labor may still fall below the subsistence cost of humans. We formalize this as follows:

Concern 3' (*Weak Economic Redundancy of Labor*). The competitive wage of human labor falls below the subsistence cost of humans, or equivalently, for given economy-wide factor supplies of labor ℓ and other factors m ,

$$F_\ell(\ell, m) < c_0$$

Observe that Strong Economic Redundancy (Concern 3) implies Weak Economic Redundancy of Labor (Concern 3') in a competitive economy but not vice versa, and that Weak Economic Redundancy (Concern 3') can be satisfied even if the Perfect Substitutability of Labor (Concern 2) is violated. To prove the first statement, if Concern 3 holds for all ℓ, m , then it also holds for infinitesimally small ℓ and m , and we can re-express $F(\ell, m) = F(0, m) + \ell F_\ell(\ell, m)$ and $F(0, m + \Delta m) = F(0, m) + \Delta m F_m(\ell, m)$ where $F_m = 1$ because we are in a competitive economy and normalized the user cost of machines to unity. The first inequality in Concern 3 can then be written as $\ell F_\ell \leq \Delta m$, and combining it with the second inequality in the condition, $\Delta m < c_0 \ell$, yields Concern 3'. Conversely, for an example that satisfies Concern 3' but does not satisfy either Concern 2 or 3, consider an economy with a production function $F(\ell, m)$ in which a small amount of labor is essential – figuratively speaking, a worker who presses the “on”-switch of the machines every morning – but in which any additional labor exhibits productivity below c_0 . In such an economy, labor is still essential, and machines are not perfect substitutes, but the market clearing wage is below the subsistence level c_0 .

Whereas Concerns 2 and 3 described features of the technological environment, i.e. the supply side of the economy, Weak Economic Redundancy (Concern 3') characterizes the economy's equilibrium and therefore depends not only on technology/labor demand but also on labor supply. For example, if an economy in which labor was weakly redundant introduces a more generous welfare system and labor supply declines, it may lift the marginal product of labor so the condition no longer holds.

A critical component in both conditions for the redundancy of labor was the market price of human labor compared to workers' subsistence cost. Advances in technology

may well lead to declining nominal consumer prices as production becomes more efficient. Economic redundancy would only be reached if competitive market wages decline faster than the subsistence cost. This would likely be the case if machines become ever more efficient – compared to humans – at transforming factor inputs such as energy and raw materials into output.

2.4 Objections

Human Superiority One common objection to the described perspective is that human labor will never become fully redundant because humans are innately superior to machines in certain domains. This belief is held firmly by many. We acknowledge that this is a possibility, but we also observe that there are no physical or economic laws that would suggest that the intelligence and dexterity of machines cannot in principle surpass their human counterparts. Human intelligence is subject to significant natural limits, for example because of natural constraints on the size of our brains. At present, these seem difficult to overcome.

Advances in the intelligence of machines have far outstripped advances in human intelligence in recent decades. If this differential progress continues, machines will eventually surpass human levels of intelligence.⁴ We also note that for machines to be able to become perfect substitutes for all labor (Concern 2), they do not necessarily need to have anything corresponding to human consciousness, and they do not need to possess metaphysical attributes that many attribute to humans, such as a soul. They just need to be able to perform all tasks that have economic value in an effective manner – including tasks that involve social and emotional intelligence, which will require them to develop a sufficiently advanced theory of mind. A useful discussion of several related points is given in section 6 of Turing (1950)’s famous paper on “Computing Machinery and Intelligence,” in which he responds to objections to the notion of machine intelligence.

The Lump-of-Labor Fallacy and New Jobs There have been numerous predictions about the “end of labor” in the past, most famously predictions that were based on the so-called “lump-of-labor fallacy.” This is the notion that there is a fixed amount of work to be done in the economy and by extension a fixed number of jobs, so that automating jobs will generate persistent unemployment. The lump-of-labor fallacy is false because it fails to acknowledge the power of markets – specifically of the price mechanism – to clear demand and supply. In a well-functioning economy, surplus labor exerts downward pressure on wages, making it cheaper to create new jobs and

⁴Transhumanists observe that it may be possible to use human enhancement technology to enable humans to merge with machines and/or to keep up with rapid advances in machine intelligence (see e.g. Bostrom, 2014). However, such technology may be very costly and not widely accessible. Korinek and Stiglitz (2019) argue in section 6 of their paper that the long-term implications of human enhancement for inequality and for the fate of the median worker may not be very different from a scenario in which machine intelligence evolves separately from humans.

hire workers up until the point where the market clears again. Similarly, innovation that creates additional demand for workers exerts upward pressure on wages until the market clears. From a macroeconomic perspective, what matters is not the creation or destruction of specific jobs – these are symptoms of the economy’s adjustment process – but the effects of technology on overall labor demand.

Note that Concerns 3 and 3’ above about the economic redundancy of labor did not state that jobs would not exist or that no new jobs could be created, just that they would not pay living wages. Or more technically speaking, that the marginal product of labor would be less than the cost of human subsistence. In other words, our concern about the economic redundancy of labor is very different from the (false) concerns articulated by the lump-of-labor fallacy.

Historic Extrapolation Concerns about the economic redundancy of labor are in stark contrast with the historical experience since the Industrial Revolution. During the 19th and early 20th century, machines replaced our brawn so humans instead focused on more brain-intensive cognitive tasks. Since the beginning of the computer age, machines have replaced dull and repetitive structured information processing tasks, while creating new jobs for humans that involved more varied cognitive tasks and leveraged our multi-faceted human intelligence. Autor et al. (2021) document that 63.5% of US employment in 2018 occurred in job titles that did not even exist in 1940. Given the additional wealth generated by technological advances in each of these cases, there was also greater demand for labor, including for the newly created jobs, leading to rising wages. Aghion et al. (2019) describe in an elegant model how wages and the economy can grow continually along a balanced growth path if a constant fraction of all the tasks in which human labor is employed is automated every period – assuming that labor remains essential for the remaining tasks.

However, there are no fundamental physical or economic laws that would say that these patterns will to continue to hold going forward – it is solely based on extrapolation from the 250 years since the Industrial Revolution. Moreover, before the Industrial Revolution, Malthusian forces implied that living standards were essentially stagnant for much of humanity. This serves as a reminder that past trends will not necessarily continue in the future.

Human Demand Another objection is that an economy could not operate without consumer demand, i.e. that humans need to earn wages so that they can afford to consume goods and services and keep the economy going (see e.g. Ford, 2015). This is a fallacy – it is true that all output produced needs to be demanded by someone or something. However, (1) that demand does not have to derive from humans, and (2) even to the extent that it derives from humans, it need not be financed with labor income. On the first point, it is perfectly feasible for a thriving economy to exist in which all output is used solely for investment purposes, i.e. machines producing output to serve as input for machines (see e.g. Korinek, 2019). The economy would have to

re-tool, for example by switching from agricultural farms to server farms that meet the demand created by machines, but there are no economic laws that would make this impossible. On the second point, human demand could also be financed by other sources of income than labor, for example by factor income such as the returns to capital, machines or land, by transfers and social benefits, or by government spending.

Nostalgic Jobs Another objection is that even if human labor can be perfectly substituted for (Concern 2), the economic redundancy of labor will be avoided because humans will always prefer to obtain certain services from other humans rather than from machines, for reasons that we may call “nostalgic.” For example, humans may prefer not to replace the services of human priests, judges, or lawmakers. In that case, strong economic redundancy (Concern 3) will not be satisfied in nostalgic jobs.

If human labor is indeed perfectly substitutable and the cost of substituting for it progressively declines, then reasoning that there will remain certain human-only nostalgic jobs relies on three assumptions: First, it requires that we can in fact tell the difference. A robot priest who has greater emotional intelligence than humans and has a more comprehensive understanding of the human psyche than a human priest may well be able to play the role of human priest quite perfectly, or intentionally slightly imperfectly so as to not give away that it is a robot. Second, it assumes that humans will still prefer services performed by other humans – even if machines have an equal or superior track record. For example, properly calibrated artificial judges may be able to make more accurate and humane judgments than humans, leaving behind the noise, discrimination and biases that have plagued our justice system (e.g. Kahneman et al., 2021). Sufficiently advanced autonomous vehicles may kill far fewer people on our roads than human drivers. It may seem brutish to insist on jobs being performed by humans in a sub-standard and inefficient way if machines can perform them better and cheaper. Third, it requires that humans earn sufficient income to spend on human services to support human jobs, i.e. that the human share of income (consisting of both their labor and capital income) is sufficiently large.

If these three assumptions are satisfied for a sufficiently large number of nostalgic jobs, labor demand will remain high enough to avoid weak economic redundancy (Concern 3') and keep wages above subsistence levels. Nostalgic jobs may therefore be an important way of keeping humans employed in jobs that pay living wages. However, if the number of nostalgic jobs that survive is too small, overall labor demand may still decline to the point of triggering weak economic redundancy and pushing wages below subsistence levels.

Comparative Advantage The concept of strong economic redundancy (Concern 3) corresponds to machines having absolute advantage in all tasks that can be performed by human labor. By contrast, the theory of international trade tells us that what matters for gainful exchange is comparative advantage, not absolute advantage – more developed nations that are technologically superior in the production of all goods

can still engage in gainful trade with less developed nations because they export the goods in which they have comparative advantage, i.e. in which they are relatively more productive, and import what their trading partner has comparative advantage in.

Some invoke the principle of comparative advantage to argue that humans cannot become economically redundant when they interact with ever more intelligent machines. However, this conflates two separate questions. In trade theory, countries are assumed to possess exogenous endowments of factors such as labor that they deploy to their most efficient use – even if these factors earn a pittance. Trade theory does not usually consider whether it is actually cost-effective to maintain these factors. In our setting, by contrast, factors (including labor) are costly to maintain and producers can choose which technology to pick – under strong economic redundancy, labor is simply a dominated technology that is not worth paying for from a purely economic perspective.

3 Optimally Allocating Work and Income

This section lays out an economic framework to analyze how to optimally allocate work and income from a utilitarian perspective given the state of technology. To do so we characterize the first-best allocation of the economy, i.e. how to allocate scarce resources in order to maximize utilitarian welfare. This reflects economic allocations under idealized circumstances – it assumes that we can start from an institutional blank slate, i.e. that we can design institutions or adjust existing institutions so as to achieve the described allocation of work and income. This implies that we do not restrict our thinking to conform to the institutional status quo – we just ask what is the best possible allocation that a utilitarian planner could aim for given the law of scarcity, which is formally captured by a resource constraint.

We start out by considering the case of a single economic agent – an individual who has a given level of labor productivity and receives a certain amount of non-labor income, e.g. income from capital or a benefit payment. At first we abstract from any non-monetary amenities that come with work and may affect the agent’s welfare, such as meaning or social connections – we describe work simply as a transaction that gives up valuable leisure time to produce output. This allows us to illustrate the fundamental forces at work in an intuitive figure. Next we examine the case of multiple individuals who differ in their labor productivity, and we analyze how work and consumption should be allocated between them in an idealized first-best world. Then we extend our framework to incorporate non-monetary amenities such as identity and meaning and analyze how this should affect the allocation of work and income. Finally, we also observe that some amenities – such as social connections – may be subject to externalities or internalities, implying that individuals may make inefficient choices on whether and how much to work.

3.1 Work and Income for an Individual Agent

Consider an individual consumer-worker who values consumption c and leisure $1 - \ell$, with preferences given by

$$U(c, \ell) = u(c - c_0) + v(1 - \ell) \quad \text{for } c \geq c_0 \text{ and } -\infty \text{ otherwise} \quad (1)$$

where ℓ captures the fraction of time worked, and $u(\cdot)$ and $v(\cdot)$ are both strictly increasing, concave, and satisfy the Inada conditions. We assume that the agent faces a subsistence level of consumption c_0 that is required for her survival. When $c < c_0$, the agent perishes, captured by a utility level of $-\infty$. This is not a significant constraint in today's advanced economies, but it is important to include for some of the adverse scenarios that we are investigating. Observe that our setup reflects that the disutility of the first marginal unit of labor supply at $\ell = 0$ is positive, given by $v'(1) > 0$, because it requires giving up valuable leisure time. We believe that this is an important feature of preferences that is missing in some neoclassical models that impose an Inada condition on the disutility of labor at $\ell = 0$ for analytic simplicity, baking in that agents always supply a positive amount of labor, even if they obtain vanishingly small wages.

The agent's labor is converted into consumption goods according to the production function $y = f(\ell) = w\ell$, where w reflects the agent's labor productivity. The units of ℓ are chosen such that $\ell = 1$ corresponds to the maximum amount of work possible; therefore w reflects the amount of output produced if the agent engages in maximum work. In a competitive market economy, the labor productivity would correspond to the agent's wage, but our analysis here is focused on the first-best and is more general than a specific market structure.

We assume that the agent also obtains non-labor income in the amount of T , which can be interpreted as income from other factors such as autonomous machines, land or capital, or as a transfer or benefit payment, and which we take as exogenous for the analysis of the single-agent case. Together, this implies a resource constraint that reflects that the agent's consumption needs to be covered by the sum of her labor and non-labor income,

$$c = w\ell + T$$

The first-best in this economy maximizes the agent's utility subject to the resource constraint or, combining the two,

$$\max_{\ell \geq 0} U(w\ell + T, \ell) \quad (2)$$

This following proposition describes the solution to the optimization problem formally:

Proposition 1 (Optimal Allocation of Labor). *The agent's optimum can be decomposed into the following regions:*

1. (*Perish*) If $T + w < c_0$, the agent cannot meet the subsistence level of consumption and perishes.




Figure 2: Regions for labor provision and iso-welfare curves

2. (Work) If $T + w \geq c_0$ and the agent's labor productivity is sufficiently high, $w > \bar{w}(T)$, the agent works. The reservation level of productivity $\bar{w}(T)$ is given by

$$\bar{w}(T) := \frac{v'(1)}{u'(T - c_0)} \text{ if } T > c_0 \text{ and } 0 \text{ otherwise} \quad (3)$$

The optimum amount of labor is determined by equating the marginal benefit of greater output to the marginal cost of forgoing leisure,

$$wu'(w\ell + T) = v'(1 - \ell) \quad (4)$$

An increase in non-labor income T raises utility, decreases optimal labor supply, and raises the threshold \bar{w} for $T > c_0$. An increase in w increases utility and has an ambiguous impact on the optimum amount of labor ℓ . If $T < c_0$, labor must satisfy $\ell > (c_0 - T)/w > 0$ to guarantee the agent's survival.

3. (Don't Work) If $T > c_0$ and $w \leq \bar{w}(T)$, it is optimal for the agent not to work.

Proof. See Appendix 1. □

The left panel of figure 2 illustrates the three different regions into which the solution to the agent's optimization problem can be decomposed as a function of her labor productivity w and the non-labor income T . The figure covers the whole range of possible economic scenarios that an individual agent may face as a result of the increasing automation of labor.

1. (Perish) The triangle at the bottom left of the figure is region 1, which reflects the most dystopian scenario: it occurs when the sum of the individual's potential labor income and non-labor income are insufficient to meet the subsistence level

of consumption. In this region, the individual will perish. This may materialize if the individual's labor productivity diminishes as a result of ever more automation, and if she does not have sufficient alternative non-labor sources of income to survive. The region is delimited by the minus-45 degree line representing the constraint $T + w = c_0$, which we may call the Malthusian frontier.

2. (Work) To the North-East of the Malthusian frontier is region 2, in which it is optimal for the individual to work. To the left of the dashed vertical line that captures $T = c_0$, it is necessary for the individual to work to cover her subsistence consumption. To the right of the vertical line, she could survive based on the non-labor income T , but it is optimal for her to work as long as her labor productivity is sufficiently high – above the threshold $\bar{w}(T)$, which corresponds to the reservation wage in a competitive economy. $\bar{w}(T)$ is an upward-sloping curve that reflects that a higher non-labor income raises the threshold that makes it worthwhile for an individual to work.
3. (Don't Work) The area to the right of the reservation wage curve is region 3. It captures that the individual's labor productivity is low relative to her non-labor income, making it optimal for her to enjoy her time on leisure instead of working. It would be socially wasteful for the individual to work in this region since the marginal benefit of extra consumption is declining in the level of consumption, and the extra income generated by work would not compensate the individual for her loss of valuable leisure time.

In the right panel, Figure 2 shows the iso-welfare lines of the agent, which are defined for regions 2 and 3. A given iso-welfare curve depicts all combinations of labor productivity w and non-labor income T for which the individual is indifferent. For example, the individual is indifferent if she moves upward or downward along the indifference curve that goes through point A_0 , as any change in her labor productivity is compensated by an offsetting change in the non-labor income. This is illustrated e.g. by a movement from A_0 to point A_2 . The arrows in the figure illustrate the direction in which welfare is increasing.

The indifference curves in the figure allow us to illustrate how technological changes affect welfare. In region 2, there are two distinct channels through which technological change affects individuals' welfare: changes in labor productivity, captured by vertical movements, and changes in non-labor income, captured by horizontal movements in the figure. Starting from an initial point A_0 in the figure, we illustrate several possibilities:

- Progress that reduces the labor productivity of the agent unambiguously reduces her utility, as reflected in point A_1 . As we discussed above in Section 2, this is a scenario that is frequently discussed by technologists who focus on the labor-displacing properties of new technologies.
- By contrast, technological progress that increases the non-labor income T of the agent unambiguously increase her utility, as illustrated in the movement to point

A_3 . An increase in T arises most directly if the individual holds factors such as capital or land that become more productive as a result of the technological progress. Alternatively, it could also result from institutional changes that provide transfers to the the agent.

- Predictions of future technological progress frequently involve a combination of the two points discussed so far – a diminished role for labor but greater non-labor output. The first effect reduces utility whereas the second effect increases utility, and the overall impact can be either. In our figure, we have illustrated the knife-edge case in which lower labor productivity is precisely offset by higher non-labor income in point A_2 . Whenever the increase in non-labor income is sufficient to offset the decline in labor productivity, technological progress leaves the individual better off, i.e. on a higher indifference curve in the figure.
- The best-case scenario is illustrated in point A_4 , in which both the agent’s labor productivity and her non-labor income go up, increasing the agent’s welfare on both counts.

3.2 Work and Income Distribution

Let us now extend our analysis to an economy in which there are many different types of agents to analyze how to optimally allocate work and income among them from a utilitarian perspective. We use this framework to analyze which agents should work and how much income the different agents should receive depending on their labor productivity. This goes to the heart of the question of how to optimally distribute income and work. We continue to focus on the first-best allocation that describes the best possible allocation that a utilitarian planner could aim for, given the law of scarcity.

Analytically, we consider I different categories of consumer-workers with utility function (1) indexed by $i = 1, \dots, I$ who differ in their labor productivity w^i . For simplicity we assume that there is a unit mass of each type of agent. A planner assigns welfare weight θ^i to each agent type i and values social welfare according to the function

$$\sum_i \theta^i U(c^i, \ell^i) = \sum_i \theta^i [u(c^i - c_0) + v(1 - \ell^i)]$$

Given the varying productivity levels among agents, the aggregate amount of effective labor in the economy is given by $\sum_i w^i \ell^i$. We assume that there is a given amount of the other factors of the economy, here captured by machines m , and embed the two types of factors in our earlier production function

$$Y = F(\ell, m) = F\left(\sum_i w^i \ell^i, m\right) \quad (5)$$

We denote by autonomous output Y_0 the amount of output when nobody is working, $Y_0 = F(0, m)$. When machines are perfect substitutes for labor (Concern 2), then $Y_0 > 0$ for $m > 0$. Moreover, we denote the amount of output when everybody works full steam at $\ell^i = 1$ by $\bar{Y} = F(\sum_i w^i, m)$. For the following analysis, we assume that $\bar{Y} > Ic_0$, which guarantees that it is feasible to cover the subsistence consumption level of all agents in the economy.

A planner in this economy maximizes welfare subject to the resource constraint

$$\max_{\{c^i, \ell^i \geq 0\}_i} \sum_i \theta^i U(c^i, \ell^i) \quad \text{s.t.} \quad \sum_i c^i = Y = F\left(\sum_i w^i \ell^i, m\right) \quad (6)$$

The ensuing proposition first lays out the case of a utilitarian planner who assigns equal weight $\theta^i = 1/I$ to all agents, followed by the case of more general welfare weights.

Proposition 2 (Optimal Labor Allocation With Heterogeneous Productivity). *A. The first-best under a utilitarian planner is characterized as follows:*

1. *There is a reservation level of productivity \bar{w} such that the planner assigns no work $\ell^i = 0$ to all agents with labor productivity $w^i \leq \bar{w}$, and a positive amount of work*

$$\ell^i(w^i) = 1 - (v')^{-1} [w^i \cdot F_\ell(\ell, m) u'(\cdot)] \quad (7)$$

to each agent with $w^i > \bar{w}$, where $\ell^i(\cdot)$ is increasing in the individual's labor productivity w^i and the aggregate marginal product of labor F_ℓ and decreasing in aggregate income Y .

2. *The planner assigns everyone to work as long as the resulting output is below a threshold $Y < \underline{Y}$ and starts to phase out workers thereafter. The planner completely phases out human labor if autonomous output rises above a threshold $Y_0 \geq \bar{Y}_0$.*
3. *The distribution of consumption is independent of the distribution of productivity and is given by*

$$c^i = \bar{c} = \frac{Y}{I} \quad \forall i \quad (8)$$

B. In the more general case with heterogeneous welfare weights, a greater welfare weight θ^i assigned to agent i implies that the agent is assigned less work and more consumption. Specifically, the productivity threshold $w(\theta^i)$ below which agents do not work is an increasing function of the welfare weight θ^i and labor supply $\ell^i(\theta^i)$ is a declining function of the welfare weight θ^i . Consumption $c(\theta^i)$ satisfies $c(\theta^i) > c_0$ and is an increasing function of the welfare weight θ^i . There is a threshold \bar{Y}_0 above which all human labor is phased out.

Proof. See Appendix A. □




Figure 3: Optimal allocation of labor ℓ^i as a function of labor productivity w^i

Figure 3 illustrates the optimal labor allocation ℓ^i as a function of an agent's labor productivity w^i for a utilitarian planner. The bold bar on the horizontal axis indicates the range of productivity levels that the agents of the economy possess, going from w^{\min} to w^{\max} in the example we depicted. The curves in the figure depict the optimal labor schedule of individuals as a function of their productivity. If the economy becomes wealthier and/or if the role of labor in production as reflected in the marginal product of labor declines, the optimal labor schedule shifts to the right, implying that it is optimal for individuals with given labor productivity to work less.

The figure illustrates three scenarios:

1. In the baseline scenario, indicated by the solid line, the reservation threshold for work \bar{w}_1 is below the labor productivity of the least productive agent w^{\min} . This implies that it is optimal for everybody in the economy to work. As the upward slope of the line illustrates, agents with greater productivity w^i are assigned a higher optimum amount of labor since their work produces more socially valuable resources. The curve is concave since the individuals' disutility of labor is rising as they work more and more.
2. The second scenario, indicated by the dashed line, corresponds to an economy that has become wealthier and/or in which labor plays a less important role in production as reflected in a lower marginal product. The reservation threshold for work \bar{w}_2 is now between w^{\min} and w^{\max} , implying that work is phased out for agents with relatively low productivity, $w^i \leq \bar{w}_2$. Intuitively, it is optimal for those agents to enjoy their leisure – they have a comparative advantage in enjoying leisure and can best contribute to overall welfare by enjoying all of their time off work. For agents with productivity above the threshold \bar{w}_2 , it is still optimal to work an amount of time that is increasing in their labor productivity; however, they work less than in the baseline scenario (solid line).

3. In the third scenario, indicated by the dotted line, the optimal labor schedule has shifted even further to the right. The scenario corresponds to an economy in which the autonomous output Y_0 (i.e. the output when nobody is working) is so high that it is no longer worthwhile to employ human labor, reflected by a threshold \bar{w}_3 that is above the productivity level of the most productive worker in the economy. Even for that worker, it is best to spend all of her time on leisure – given the abundance of material resources, this contributes more to social welfare than her comparatively miserly labor productivity. The economy has reached full unemployment.

Regarding income distribution, a utilitarian planner places equal weights on individuals and thus allocates income equally among all agents in the first best. By implication there is no income and consumption inequality. This highlights that market allocations which tie people’s income to the marginal product of their labor are sub-optimal compared to our first-best benchmark — although this is the outcome of competitive market forces, it is highly inefficient from the perspective of maximizing utilitarian social welfare. Moreover, even though such a planner would distribute income equally, there is inequality in utility – a utilitarian planner assigns more work and thus less leisure to individuals with higher labor productivity. The intuition is that the planner recognizes that individuals with low labor productivity still have the same valuation of leisure as everyone else, and it would be inefficient to make them give up their valuable leisure for work that does not produce sufficient social value.⁵

Justifying the levels of income inequality that we observe in practice requires that a planner places starkly different weights on the welfare of different individuals. To illustrate this, let us employ the commonly used CES utility function with coefficient of relative risk aversion of $\sigma = 2$ to estimate the relative marginal valuation of consumption of individuals at different consumption levels and, to be as conservative as possible, disregard the subsistence income by setting $c_0 = 0$. Meyer and Sullivan (2017) estimate that 90/10 consumption inequality is conservatively approximately by a ratio of 4/1 in the US in recent decades, i.e. consumption at the 90th percentile was four times that at the tenth percentile. This would correspond to the first-best choices of a planner who values individuals at the 90th percentile 16 times more than those at the tenth percentile. Comparisons across countries reveal even starker differences. For example, in several Central African countries, more than half of the population live below the World Bank’s threshold for extreme poverty of \$1.90/day, less than 1/71 of the average consumption of US households, corresponding to a welfare weight that is about 5000 times greater.

⁵As a result, in a stark reversal to the outcome of our current economic institutions, low productivity agents would actually enjoy greater overall utility than high productivity agents in the first best of a utilitarian planner.

3.3 Work Amenities

Work is not only an exchange of time against income, but it provides a bundle of non-monetary amenities that directly affect workers' welfare. Traditional job amenities include positive factors such as benefits, flexibility, and learning opportunities as well as negative factors like risk exposure, pollution, and arduous working conditions (see e.g. Rosen, 1986). As workers become more prosperous, they also increasingly care about factors such as social connections, identity, meaning, power, and self-actualization (see e.g. Danaher, 2017). This can be interpreted through the lens of Maslow (1943)'s hierarchy of needs as going from deficiency needs to growth needs.

Much of the economic discussion of the future of work focuses on how to provide income to people when they can no longer earn a living wage. However, sociologists, anthropologists, and philosophers have long observed that work also affects workers' welfare through these non-monetary amenities. This becomes especially salient when working-age people lose their jobs (see e.g. Brand, 2015). A comprehensive discussion of economic policy for a work-less future needs to take into account the role of these factors.

We expand our analysis to account for the role of job amenities that affect an individual's welfare through channels other than consumption utility or the disutility of providing labor effort. We capture the amenities that an individual receives from work in the scalar variable a^i , which depends on the amount worked by the individual and enters the utility function as

$$U(c^i, \ell^i, a^i) \tag{9}$$

where $U(\cdot)$ is quasiconcave and satisfies $U_{a^i} \geq 0$. We assume that the work amenities are linear in the amount worked by the individual, $a^i = \alpha^i \ell^i$, where the coefficient $\alpha^i \gtrless 0$ reflects that work may give rise to either amenities or disamenities.

Let us define the range of the utility function for zero work as $\mathcal{U}_0 = \{u : u = U(c, 0, 0), c > c_0\}$, which indicates all possible utility values that can be obtained by an individual who is meeting the subsistence consumption level but not working. For the following, let us consider a worker i with labor supply $\ell^i > 0$, labor income $w^i \ell^i$, consumption $c^i = w^i \ell^i + T^i$, and amenities $a^i = \alpha^i \ell^i$

Lemma 1 (Compensating Differential for Work). *(a) A compensating differential z^i that makes worker i indifferent between working and not working exists if and only if $U(c^i, \ell^i, a^i) \in \mathcal{U}_0$. The differential z^i is defined by*

$$U(T^i + z^i, 0, 0) = U(T^i + w^i \ell^i, \ell^i, a^i)$$

(b) The compensating differential is always positive $z^i > 0$ but can be more or less than labor income, $z^i \gtrless w^i \ell^i$. If $\alpha^i \leq 0$, then $z^i < w^i \ell^i$. If $\alpha^i > 0$, it is possible that $z^i > w^i \ell^i$. The compensating differential is strictly increasing in the labor productivity w^i and the amenity coefficient α^i of the agent.

Proof. See Appendix A. □

The lemma reflects that the total utility impact of work — including the earnings, the disutility of labor, and the benefits or costs of the individual amenities obtained — can alternatively be derived from a compensating differential of z^i units of consumption goods if the stated condition on the range of the utility function is satisfied. The principle is based on Rosen (1986)’s equalizing differences, although, critically, the lemma here looks at the *total effects* of work rather than solely at job amenities a^i .

The range condition on the utility function $U(c^i, \ell^i, a^i) \in \mathcal{U}_0$ captures that variations in consumption need to have sufficient impact on utility to make it possible to compensate individuals for their job amenities. This is an important condition, and it may not always be satisfied: consider an individual who values the meaning derived from her job so greatly that there is no amount of consumption goods that could compensate her for that experience – in that case, a compensating differential for the individual’s work does not exist. A stark example of preferences for which a compensating differential does not exist are lexicographic preferences for certain job amenities (which technically also cannot be captured by a utility function). Whenever the individual derives so much utility value from work amenities that she cannot be compensated, it is optimal for her to continue to work, no matter how rich society becomes and no matter how low her labor productivity.

The compensating differential for work is always positive when it exists, $z^i > 0$, since not working is in the choice set of the individual. By revealed preference, if the individual is working, it is because this makes her better off and allows her to earn a utility surplus compared to not working. $z^i > 0$ compensates the individual for this surplus.

However, the compensating differential may be more or less than a worker’s labor income, $z^i \gtrless w^i \ell^i$, depending on the value of the amenities. In traditional descriptions of the labor market that disregard amenities ($a^i = 0$), the compensating differential is always strictly less than the labor income, $z^i < w^i \ell^i$, since work requires giving up valuable leisure time. In the right panel of Figure 2, for example, the compensating differential for giving up work at, say, point A_0 can be obtained by following the indifference curve to the point where it intersects with the reservation wage curve \bar{w} , i.e. to point B , at which the individual ceases to work. The horizontal difference between A_0 and B reflects the transfer necessary to make the individual indifferent between working at A_0 and not working at B . If work involves disamenities to the individual, $a^i < 0$, then the individual needs to receive even less compensation to give up work, and $z^i < w^i \ell^i$ holds a fortiori.

Whenever $z^i < w^i \ell^i$, an individual whose job is displaced can be fully compensated with a benefit that is less than the wage received. Gallup (2022) reports that nearly 85% of employees worldwide and 65% of employees in the US are not engaged or actively disengaged at work, suggesting that they fall into this category. If the job could be automated at sufficiently low cost and the surplus could be shared with the individual, overall welfare would increase.




Figure 4: Optimal labor allocation with amenities

On the other hand, for workers who receive sufficiently positive work amenities, $z^i > w^i \ell^i$ may be the case, i.e. they would need to receive a compensation that is greater than their wage earnings to give up work, since the positive amenities associated with their work more than make up for the loss of valuable leisure time.

In the following, we assume that the condition of the lemma is satisfied, and that agents' utility function has the linearly separable form

$$U(c^i, \ell^i, a^i) = u(c^i - c_0) + v(1 - \ell^i) + x(a^i) \quad (10)$$

where $x(\cdot)$ is increasing and concave with elasticity $\epsilon_x = -ax''/x' < 1$.⁶

We now revisit the question of how to allocate work and income for the case that workers are heterogeneous in both productivity w^i and amenities α^i , focusing on the problem of a utilitarian planner and assuming there is sufficient output for all agents to reach their subsistence level of consumption.

Proposition 3 (Optimal Labor Allocation With Amenities). *A. The first-best of a utilitarian planner is characterized as follows:*

1. *There is a frontier of productivity levels and amenity values (w^i, α^i) such that the planner only assigns work $\ell^i > 0$ to agent i if w^i or α^i are sufficiently large, i.e. if $w^i > \bar{w}(\alpha^i)$, where the reservation productivity $\bar{w}(\alpha^i)$ is declining in α^i . Labor ℓ^i is increasing in both w^i and α^i .*
2. *If $\alpha^i > \alpha_0^i = v'(1)/x'(0)$, then the planner will never phase out work for agent i . This is always the case when the compensating differential for the agent's work satisfies $z^i > w^i \ell^i$ or is infinity.*

Proof. See Appendix A. □

⁶The latter technical condition ensures that $d[\alpha x(\alpha \ell)]/d\alpha > 0$, i.e. that the marginal amenities of work are increasing in the strength α of these amenities, generating well-behaved solutions.

Intuitively, expanding our earlier analysis to the case of work amenities introduces an additional trade-off. Besides considering how much an individual’s labor productivity contributes to the production of output, it also evaluates how much an agent enjoys (or dislikes) her work.

Figure 4 illustrates when it is optimal for an agent to work in the space of the amenity value α^i and labor productivity w^i . The downward-sloping curve represents the agent’s reservation productivity as a function of her amenity value of work α^i . If the agent’s labor productivity and amenities are low, she is in the region below the curve, and it is optimal for her not to work. Conversely, if either her labor productivity or amenity values are sufficiently high, then it is optimal for the agent to work. When the amenity value is above a threshold α_0^i , then an individual’s productivity no longer matters, and it is desirable for her to work even if $w^i = 0$, i.e. even if the worker is useless in the production of output.

3.4 Externalities and Internalities from Work

Individual decisions on how much to work may not always be in the best interest of society, or even of the individual making the decision, since work may generate rich externalities, or public amenities and disamenities, as well as internalities. Such externalities and internalities are likely to grow in relative importance if the role of work in generating income diminishes in the future.

Externalities imply that the work of one individual affects others in society not only through how much marketable output it produces but also through additional effects that are not mediated by the market. For example, the social connections engendered by work entail positive externalities as they arise not from the work of one individual in isolation but from the interactions of multiple workers. An economy with satisfactory jobs also creates positive externalities in the form of social and political stability (see e.g. Boix, 2022). Conversely, for examples of negative externalities, the cost of commuting becomes greater the more workers are commuting because of congestion externalities. Likewise, working hard to gain greater status imposes negative positional externalities on those whose status diminishes (see e.g. Frank, 1991).

Internalities of work occur when individuals deviate from perfect rationality in their decision-making about labor choices because they do not internalize the full effect of their choices on their own welfare. Such internalities can also be either positive or negative. For example, a hard-working individual who tells himself that work is a critical component of his happiness even though he would in fact be happier spending more time with family is experiencing negative internalities from work. Conversely, an example of positive internalities from work would be an individual who lacks structure and meaning and chooses to let his days pass by in loneliness but would in fact be happier accepting an offer for a job that provides a regular schedule and social interactions.

To formally account for the externalities generated by work, we assume that agents obtain not only individual amenities a^i from their work but also collective amenities \bar{a} that are generated by the work of all the agents in the economy, $\bar{a} = \sum_j \bar{\alpha}^j \ell^j / I$,

with $\bar{\alpha}^j \geq 0$ a parameter that captures how much type j contributes to the collective amenities, normalized by $1/I$. For simplicity, we consider a representative agent of type i who is atomistic so that her impact on aggregate \bar{a} is negligible. We expand the utility function of consumer-worker i to capture \bar{a} in the following linearly separable manner,⁷

$$U(c^i, \ell^i, a^i, \bar{a}) = u(c^i - c_0) + v(1 - \ell^i) + x(a^i) + \bar{x}(\bar{a}) \quad (11)$$

where $\bar{x}(\bar{a})$ is also increasing and concave with elasticity $\epsilon_{\bar{x}} < 1$.

To formalize the role of internalities, we assume that the individual who is choosing how much to work mis-perceives the amenities derived from his work – instead of the term $a^i = \alpha^i \ell^i$ in utility function (11), the individual perceives the amenities as $\hat{a}^i = \hat{\alpha}^i \ell^i$ with $\hat{\alpha}^i \neq \alpha^i$. One example is the case $\hat{\alpha}^i = 0$, i.e. that the individual disregards the amenity value of work and only focuses on the monetary payoffs and disutility of labor.

We then find the following results:

Proposition 4 (Allocation of Labor with Externalities and Internalities). *1. There is a frontier of $(w^i, \alpha^i, \bar{\alpha}^i)$ such that the planner only assigns work $\ell^i > 0$ to agent i if w^i , α^i and $\bar{\alpha}^i$ are sufficiently large, and ℓ^i is increasing in all three parameters.*

2. If $\bar{\alpha}^i \neq 0$ or $\alpha^i \neq \hat{\alpha}^i$, then individual labor supply decisions are inefficient. The planner's optimal subsidy (tax) to individual i 's labor is

$$s^i = \frac{\bar{\alpha}^i \bar{x}'(\bar{a}) + [\alpha^i x'(a^i) - \hat{\alpha}^i x'(\hat{a}^i)]}{u'(c^i)} \quad (12)$$

Proof. See Appendix A. □

Intuitively, this framework captures that there are two additional considerations in determining the optimum amount of work for individuals. The planner only assigns work to an agent if their labor productivity, their true job amenities, and their contribution to aggregate amenities are sufficiently large, and the optimum amount of work for the individual is increasing in all three parameters.

Both the externalities from aggregate amenities and the internalities from mis-perceived individual amenities make the choices of individuals inefficient. To the extent that they are positive, i.e. $\bar{\alpha}^i > 0$ for the externalities and $\alpha^i > \hat{\alpha}^i$ for the internalities, workers drop out of the labor market too early and supply too little labor, calling for public policy interventions such as job subsidies to incentivize work. The opposite holds for negative externalities and internalities. The optimal subsidy rate to labor provided by a planner, captured by equation (12), consists of a marginal externality term $\bar{\alpha}^i \bar{x}'$ as well as a marginal internality term in square brackets that subtracts the perceived from the actual marginal amenity value of labor.

⁷A richer version of our framework could also account for more individual-specific externalities and for the interactions of individual and aggregate amenities and disamenities as well as their effects on productivity and on the disutility of labor.

One additional aspect of externalities and internalities from work that is noteworthy is that they are not properly reflected in market prices, implying that market forces may give rise to inefficient decisions by both workers and firms in the absence of regulation. For example, firms may find it profitable to automate certain jobs that generate important public goods or positive internalities because they do not internalize the losses that society will experience as a result. Similarly, workers may happily accept a payout and stop working, even though they may end up suffering welfare losses because of the externalities or internalities from work.

4 Economic Institutions to Allocate Work and Income

The previous section analyzed how to allocate work and income in a “Panglossian” first-best world. This was instructive because it approaches the question from an idealized perspective that does not limit our thinking to what is possible under a specific set of institutions. However, it also risks disregarding the reasons why our institutions are set up the way they are. For example, it ignores the important role of incentives and the resulting limitations to raising government revenue and distributing income; it also ignores information problems and the associated reasons why certain markets do not exist and cannot be created in practice.

This section assesses the complementary roles played by the two main types of institutions to allocate work and income in our economy and society – by markets and by social insurance systems – when the concerns we spelled out in Section 2 materialize.

4.1 Markets

Markets are one of the main institutions that shape the distribution of income in today’s world. Economists typically start their analysis of market allocations by considering the theoretical benchmark of a free market economy in which prices are determined solely by the forces of demand and supply in competitive markets and no government intervention takes place. Just like our first-best analysis before, this is an instructive benchmark but not directly applicable to the real world – as already Karl Popper (1945) observed, “the idea of a free market is paradoxical,” since a functioning market requires a rich set of governmental institutions to support it, starting with institutions that protect our safety from internal and external threats to functioning judiciary, antitrust, and consumer protection institutions. It also disregards all the market imperfections that exist in the real world, ranging from externalities to other forms of missing and restricted markets.⁸

⁸In particular, AI and related information goods – which are playing a crucial role in reshaping the (potentially non-existent) future of work – are non-rivalrous and therefore fundamentally inconsistent with competitive markets. Perfect competition implies that goods are traded at the marginal cost of producing an extra unit, which is zero for information goods. But this makes it impossible for

Complete markets economy Putting aside these difficulties, we start our analysis with the benchmark case of a free market economy in which markets are complete, i.e. there are markets to trade all goods and insure against all risks, including the risks inherent in technological progress. Comparing this benchmark both with the real world and with the idealized first-best allows us to pinpoint the shortcomings of markets and by extension the potential role for public policy to complement markets.

An economy with complete markets would – surprisingly – replicate the first-best allocation of an egalitarian planner that we described in Section 3. Given complete markets, such an economy would also have complete risk markets, including risk markets that insure individuals and their descendants against the redistributions generated by technological progress.⁹ Obviously, such insurance markets do not exist in practice. The reasons why are the same as why it is impossible to implement the first-best allocation by government fiat – perfect insurance is unattainable largely because of limits to information and incentive problems.¹⁰ See Korinek and Stiglitz (2019) for a fuller discussion of this point.

Missing markets for long-term risk-sharing Since the real world departs from the complete markets benchmark and markets for long-term risk-sharing are largely non-existent, individuals in a free market economy have to live from their labor income and prior wealth, even when they experience stark declines in labor income. Their situation corresponds to Figure 2, with non-labor income T determined by the asset holdings of each individual, which derive either from inheritances or from the returns that individuals have earned on their labor and asset holdings in the past. The absence of markets for long-term risk-sharing implies that there are no incentive problems, but it assigns a large role to forces outside of the individual’s control and random chance, starting from the wealth of the family an individual is born into to the luck of the draw of how much the labor market values the skill set of an individual and what returns are realized on the individual’s risky investments.

Compared to the first-best allocation of an egalitarian planner, the described free market allocation without long-term insurance markets thus features large inefficiencies

producers to recoup the fixed costs that they have incurred in creating information goods (Arrow, 1962). To address this problem, our society has created intellectual property rights such as copyrights and patents (see e.g. Shapiro, 2008). Intellectual property rights confer a state-sanctioned monopoly on the creators of information goods, which introduces one inefficiency (monopoly distortions) to mitigate another inefficiency (free-riding on information goods) and to provide incentives for the creation of information goods. However, it is important to remember that intellectual property rights are the result of government intervention and thus represent a deviation from free and competitive markets.

⁹The argument requires that such risk markets existed since the beginning of time. Figuratively speaking, they would have allowed Adam and Eve to insure all their descendants against the potentially adverse effects of technological progress, creating perfect risk sharing.

¹⁰In practice, the main instrument to share the benefits and risks of new technologies are equity markets. However, in the US economy for example, 47% of families do not have any exposure to the equity markets, and only 15% of families have direct stock holdings, with a median portfolio size of only about \$25k (Survey of Consumer Finances, 2019, reported in Bhutta, 2020).

in the allocation of income and work: it implies that some unlucky individuals have low net worth and are forced to work even if their labor productivity is low and their work generates little economic value. Conversely, others have high net worth and no need to work even if they have high labor productivity.

If labor becomes economically redundant (Concerns 3 and 3' in Section 2), then those individuals in a free market economy with insufficient asset holdings would no longer have the means to cover their subsistence cost – even if they worked as hard as they can. This would place them in the bottom left region of Figure 2 and leave them to perish. In the US, the net worth of the median family covers only about two years of income (Bhutta, 2020), and even less for traditionally disadvantaged groups. This implies that the majority would not survive for long in a free market economy when labor becomes redundant, leading to widespread misery and starvation. The picture would be even more dire in developing countries where wages already are closer to subsistence levels than in advanced countries and the savings of the average household are far smaller (Korinek and Stiglitz, 2021a). The described allocation would be efficient in purely economic terms – it would reflect the first-best allocation of a planner who places zero weight on individuals who succumb to starvation – it would just be incredibly cruel.

4.2 Social Insurance

Since individuals have no other source of insurance against the long-term risks of technological progress, it is natural that society has developed social insurance mechanism that provide protection against such shocks.¹¹ The benefit of insurance is greater the larger the shocks that individuals are exposed to, and the consequences of labor becoming economically redundant (Concerns 3 or 3') would be among the largest conceivable economic shocks for workers.

As social insurance partially substitutes for the missing markets for long-term risk-sharing, a by-product is that it also generates an allocation in labor markets that is more efficient from a utilitarian perspective: it implies that individuals with low net worth are no longer forced into work just to survive. The more equal the distribution of income, the closer the labor supply decisions of individual workers are to the first-best benchmark of a utilitarian planner that we described in Section 3.2, in which an individual's labor supply depends only on her labor productivity.

In the following, we first discuss the implications of labor-saving progress for social insurance. Then we examine how institutions that provide a basic income can prepare society for the economic redundancy of labor. We evaluate the pros and cons of conditioning social benefits on work and argue that unconditional benefits are likely

¹¹We employ a broad definition of social insurance here that includes all types of transfers and social protection programs in order to emphasize that all of these are designed to bring us closer to the first-best utilitarian benchmark with perfect risk markets that we analyzed earlier. This definition also includes many goods or services provided in-kind, for example for healthcare, safety, or education, which are by their nature targeted to the needs of the individual recipient.

preferable in such a scenario. Finally, we also evaluate how to adapt our system of taxation to a world of redundant labor in order to raise the revenue necessary for public spending.

4.2.1 Social Insurance Under Labor-Saving Progress

Much traditional economic work on optimal redistribution policy has focused on redistributing labor income from individuals with high incomes to those with low or no incomes, as labor income taxation has traditionally been the largest source of revenue for most governments (see e.g. Piketty and Saez, 2013, for a summary). This literature centers on the classical trade-off between equity and efficiency, weighing the benefit of providing to the needy against the efficiency costs of raising revenue and of distorting incentives.

A central insight of this approach is that greater inequality in market outcomes calls for more redistribution. If labor-saving progress reduces the wages earned by certain types of workers, an optimal social insurance system would increase the transfers to them. If technological progress increases overall output, an optimal social insurance system would ensure that all members of society benefit from the progress.

The social insurance systems that exist in most advanced countries are designed to target benefits to those who have the greatest needs, and under present circumstances, they do so more efficiently than a basic income for everyone would (see e.g. Hoynes and Rothstein, 2019).¹² However, to prepare for a future in which the role of labor may decline, they should be reformed to ensure that they do not condition benefits on work, as we will argue in more detail below. In fact, transfers that are not conditioned on work also have a multiplier effect that benefits workers.

Multiplier Effect of Non-Labor Income As long as individuals are still working, providing transfers to workers has a multiplier effect on their income by raising equilibrium wages. Greater transfers or other forms of non-labor income induce agents to work less and, once their non-labor income is sufficiently high, to stop working altogether. In a competitive equilibrium, lower labor supply raises the equilibrium wage.

We show this result formally by building on the model of Section 3.2 but focusing, for simplicity, on the case of a representative agent, dropping the superscript i from all variables. We assume that there is technological progress that leaves the marginal product of labor $F_\ell(\cdot)$ unaffected but produces extra output $dY > 0$ that is distributed to the agent as additional non-labor income $dT = dY$.

Proposition 5 (Multiplier Effect of Social Insurance). *An increase in the non-labor income of agents raises the equilibrium wage in a competitive equilibrium.*

Proof. See Appendix A. □

¹²To provide a numerical example, at the time of writing, paying every American a monthly UBI of \$1000 would use up almost the entire revenue of the US federal government.

The described multiplier effect works in both directions – if social benefits or other forms of non-labor income decline, individuals are forced to work more, which reduces equilibrium wages and lowers their incomes further. Since greater supply of labor and lower wages benefit the owners of capital, this may explain part of the political appeal to capitalists of cutting social benefits (in addition to potentially lowering their tax burden).

Role of Work Amenities Greater income from social insurance also has a multiplier effect on the provision of work amenities. As the income obtained by workers rises – or, technically speaking, as the marginal utility of consumption of market goods goes down – the non-monetary aspects of their work become comparatively more important (see e.g. Korinek and Stiglitz, 2021b). Individuals with higher income care comparatively more about work amenities. By implication, the multiplier effect of social insurance applies to market-clearing wages as well as market-clearing amenity values.

The implications for a world with labor-saving progress are twofold: on the one hand, if individuals are not compensated for income losses that result from technological progress, they will also be forced to accept lower amenities in order to make ends meet, exacerbating the adverse effects on their utility; on the other hand, if the increase in output is shared with workers, they will demand and, in equilibrium, obtain more positive work amenities, multiplying the positive effects of social insurance on social welfare.

Steering Technological Progress As long as labor income plays an important role in the distribution of income, Korinek and Stiglitz (2021b) make the case that a planner would also complement traditional forms of redistribution with a policy of intentionally steering technological progress towards labor-using rather than labor-saving innovations in order to increase market wages. Both traditional redistribution and the steering of technological progress achieve first-order gains in income distribution at the expense of second-order efficiency costs, making it desirable to employ both tools in the spirit of the theory of the second-best.

Klinova and Korinek (2021) and Korinek (2022) develop an economic framework to assess how specific technological innovations affect labor markets and inequality to determine how desirable they are from a distributive perspective. They lay out five different channels that together track the full general equilibrium effects of an innovation on labor markets. This framework enables innovators, companies and civil society to analyze the distributive effects of their research and technology choices. It also enables public policy to consider how the policy environment for innovation affects the distribution of benefits and to actively steer innovation in a direction that benefits labor.

4.2.2 Economic Redundancy of Labor and Basic Income

If labor becomes economically redundant and is replaced by autonomous machines that are cheaper than human subsistence, it will affect the distribution of income in two opposing ways. First, automating labor carries the potential for vastly greater economic output. Economic growth may increase by orders of magnitude (Aghion et al., 2019; Trammell and Korinek, 2020). In principle, this should make it easy to distribute some of the surplus generated by autonomous machines to former workers to compensate them for their losses. Second, however, the ownership of capital and capital-like factors is distributed much more unequally than labor. If income from labor becomes almost negligible and no counteracting measures are taken, inequality will increase sharply. In the absence of redistribution, humans without significant financial net worth could no longer survive based on their labor income alone.

A basic income for everyone would become very appealing from a utilitarian perspective if labor is made economically redundant. As we observed before, in our present economic environment, targeting social benefits to recipients with low incomes makes sense because it allows our social insurance systems to focus scarce resources on where they are most needed. However, if labor becomes redundant, conditioning transfers on individuals' incomes would become significantly less useful, both because the vast majority of the population would require income support and because capital income would be more difficult to ascertain than labor income. Moreover, providing a basic income to everyone would also avoid the stigmatizing effects of income tests.

A basic income would require collective action but would not necessarily need to come in the form of transfers that are directly controlled by government. For example, Altman (2021) proposes that large corporations be required to provide equity stakes to a public fund, which would make annual distributions to all citizens. One benefit of such a scheme would be that it shares the risks from technological progress widely among society, allowing everyone to benefit from economic growth and bringing us closer to the optimal insurance that would prevail under complete markets for long-term risk-sharing.

Conditioning Benefits on Work? Many proposals for how to design a basic income when work becomes redundant suggest that benefits should be conditioned on recipients contributing to the social good, for example by participating in public works programs or earning social credits (see e.g. Lee, 2018; Susskind, 2020). Variants of transfers that are conditional on work include both work requirements and work subsidies for individuals, such as the Earned Income Tax Credit (EITC) in the US. An alternative way of implementing work subsidies that relies more on the private sector is a hiring mandate for companies – for example, to require that companies need to hire a certain number of workers per million dollars in profits or otherwise pay a penalty. (If the mandate is binding, the equilibrium wage of a given job would be determined by the penalty together with the amenity value of the job, making it desirable for companies to offer enjoyable jobs.)

However, our earlier results suggest that a planner would only condition transfers on work if individual decisions do not deliver socially optimal outcomes because of externalities or internalities arising from work (Section 3.4). This brings up two important questions:

1. How would a planner determine the most significant externalities and internalities from work? We observed that there may be both positive and negative externalities and internalities. Moreover, people differ widely in their ideas of what social benefits are delivered by work. Such disagreements make the implementation of conditional transfer schemes problematic. Letting some members of society impose their choices on others runs counter to the ideas of liberalism and freedom and may reduce welfare, counter to the intentions of the program.
2. Even if a planner can correctly identify them, are there more efficient ways of providing the public goods and positive externalities arising from work? Similarly, are there better ways to account for the internalities of work and provide individuals with the positive private amenities from work that they themselves do not rationally internalize, other than requiring them to work?

One way of learning about the social value of work is to look at the consequences of job loss. In a comprehensive survey, Brand (2015) reports that the negative effects of job loss are far-reaching and include several potential internalities and externalities – they range from declines in psychological and physical well-being, social withdrawal and family disruption to reductions in the educational attainment and well-being of the children of affected workers. Crucially, however, the social context matters. Whereas unemployment fills working-age adults with dread, retirement on average leads to increases in happiness and life satisfaction (see e.g. Fonseca et al., 2017). Well-being in retirement is also supported (i) by the narrative that workers have earned their well-deserved retirement after decades of work and (ii) by a range of social customs and institutions.

This suggests two dimensions along which public policy may be able to improve outcomes. First, policymakers may need to craft the right public narrative as we approach a work-less future, that work is phased out because humanity has accomplished its economic mission and can finally transition past the Age of Labor and retire into a well-deserved new epoch in which people no longer need to work. Second, it may require that we create social customs and institutions that provide well-being, identity, meaning, social connections, and any other amenities that workers currently derive from work.

This transition may take time – parents have taught their children for generations that “working hard” is an important value (IACS, 2012). Even after work has become economically redundant, it may take a new generation of citizens to fully embrace a future of human flourishing without work, and to ensure that the positive amenities and public goods provided by work can be found in new ways. During the transition, some workers might benefit from public policies to support the existence of jobs for

those who want them. However, in the long run, there are reasons to believe that if labor becomes economically redundant, utilitarian social welfare may be maximized if humans obtain a basic income without being required to work.

A Universal Basic Income (UBI) is the unconditional transfer program that has perhaps garnered the most public attention. A UBI is an unconditional transfer to a defined group of recipients that allows them to meet their basic needs and that is obtained no matter what economic decisions the recipients make. The unconditional nature makes it a very efficient way of distributing income.¹³

In particular, a UBI neither requires recipients to work nor actively discourages them from working – unlike social programs that are conditional on people maintaining low levels of income. If someone decides to stop working in response to a UBI, economists do not consider this a distortion or an inefficiency – in fact quite the opposite: ceasing to work would simply reflect that the recipient does not consider her work worthwhile at the given income level and market wage, and that she only performed her job out of financial duress before the UBI was introduced.

In Figure 2, introducing or raising a UBI would increase non-labor income T and shift the equilibrium to the right, potentially into the no-work region. If the individual obtains sufficient positive amenities from work, she can still choose to continue working, even if her market wage goes to zero.

The basic idea of a UBI is by no means new – in 1797, Thomas Paine wrote that American citizens should receive a lump sum payment upon reaching adulthood in order to compensate for the unequal distribution of land and wealth. In 1968, 1200 economists signed a letter to Congress recommending a form of government assistance akin to universal basic income (Brynjolfsson and McAfee, 2014). Given the non-distortionary nature of a UBI, the proposal was supported by economists of all ideological orientations, including e.g. conservative icon Milton Friedman. President Nixon tried to implement a form of universal basic income under his Family Assistance Plan, but ultimately the plan fell apart because of concerns about its poor targeting.

In recent years, there has been a resurgence of interest in UBI as a potential solution to combat the economic effects of job loss stemming from automation and AI (Marinescu, 2017; Lee, 2018).

Preparing for the Redundancy of Labor Developing institutions that will be able to distribute a UBI that insures all members of society against the risks of labor-saving progress in the future may help to future-proof our society. At present, we have social insurance programs that are relatively effective at targeting benefits to those who need them most, and there are few distributional benefits to a UBI. However, a

¹³To be sure, there are two types of distortions that may be generated by UBI schemes. First, they require tax revenue which needs to be raised. Second, distortions may also arise from their eligibility criteria. For example, the Alaska Permanent Fund distributes some of the state’s oil revenue to Alaska residents only. This may create incentives to acquire residency or citizenship in order to become eligible for a UBI. It may also increase opportunities for fraud.

modest UBI program that is designed to automatically scale up if labor-saving progress continues and the earnings opportunities of workers decline in the future would provide insurance to workers. This could be achieved e.g. by indexing the UBI to the growth of the non-labor part of national income. As we noted earlier, output growth is likely to go up significantly if labor becomes economically redundant. Moreover, if the labor share of income declines, the non-labor share would mechanically rise. A UBI that automatically grows in line with non-labor income would therefore constitute an insurance mechanism for the redundancy of labor.

Making a Basic Income Truly Universal It is worth emphasizing that for a UBI to maximize utilitarian welfare for all of humanity, the recipient pool would have to be universal in the truest sense of the word – including every citizen of the world. By contrast, most current proposals for a UBI restrict eligibility to the residents of certain geographic areas or to the citizens of certain countries.

Living up to the ambition of a UBI in the truest sense of the word “universal” would be extremely difficult to implement in practice. It would require developing significant administrative capabilities that currently do not exist, including the development of registers and record-keeping systems that minimize fraud and abuse, and payment systems that distribute the scheme. It would also require mechanisms to ensure that the citizens of autocratic states can obtain some of the benefits, for example by receiving benefits in-kind. Finally, a true UBI would also be enormously expensive, and our world does not currently have any ways of raising significant amounts of revenue at the supernational level. Still, it is not unthinkable.¹⁴

If autonomous machines that fully substitute for human labor are developed, they would generate large economic gains. The companies or research labs developing such machines as well as the countries in which they are developed would obtain large wind-fall gains. They would be able to share some of their gains with the rest of the world, via philanthropy and via development assistance, to ensure that everyone benefits from the technological advances created by human ingenuity. These may be good reasons for the world’s top charities and global organizations to develop institutions that prepare our world for a future in which labor may be economically redundant.

4.2.3 Taxation

As the role of labor income in the economy diminishes, it would become necessary to focus taxation on alternatives to labor if current relative or absolute spending levels are to be maintained (Korinek, 2020). This is in stark contrast to a trend that has taken

¹⁴At present, providing e.g. a \$100 UBI to every single one of the 7.9bn citizens of the world would cost \$790bn per year and would make a significant difference to the lives of the poorest. For comparison, the combined 2021 earnings of the top-10 corporations worldwide were \$787bn (CompaniesMarketCap.com, 2022). The combined net worth of the world’s 2755 billionaires as of 2021 was \$13.1tn and would provide an annual income of \$786bn at a 6% rate of return (Forbes, 2022). All these numbers would rise significantly if labor becomes fully substitutable.

place in most countries in recent decades and that has reduced the burden of taxation on capital income. As long as labor is the most important factor in the economy, labor taxation raises substantial revenue and is a relatively efficient way of taxation (see e.g. Atkinson and Stiglitz, 1980). However, when labor income falls as a share of total output (as articulated in Concern 1), other sources of revenue will need to grow in importance. Moreover, if labor becomes economically redundant (as articulated in Concern 3), all taxation will have to rely on sources other than labor.

Henry George (1879) proposed that taxes on factors in inelastic supply, such as land, are excellent measures to raise revenue, since an inelastic supply implies that the taxes do not introduce distortions – the amount of land in the economy is essentially fixed, no matter what tax rate is imposed on it. Arthur Pigou (1920) observed that externalities such as pollution or congestion are ideal targets for taxation, since taxes on them raise revenue while simultaneously making the economy more efficient.

Another way of identifying candidates for taxation is to look for the winners of technological change. If technological progress reduces the returns on labor, then other factors in the economy must gain disproportionately and earn quasi-rents that were likely unexpected and can therefore be taxed without efficiency losses (Korinek and Stiglitz, 2019). If the same technological advances that devalued labor lead to large productivity gains throughout the economy, then they would generate a lot of income, making it easier to tax some of it to distribute to individuals whose role as workers was devalued.

5 Conclusions

Human society successfully adapted to past technological revolutions – we went from forager to farmer during the Neolithic revolution, and from farmer to worker during the Industrial Revolution. We hope that our society will also be able to adapt to the end of the Age of Labor if and when it occurs, enabling humans to enjoy their lives freed from the drudgery of having to work.

Crucially, however, the adaptation may require new economic institutions to distribute the output that is produced by autonomous machines, posing significant challenges to our existing institutions. The Neolithic Revolution led humans to establish property rights. Still, the majority of the population lived self-sufficient, largely autarkic lives. The ensuing specialization and division of labor, turbo-charged by the Industrial Revolution, implied that markets became essential for our daily needs and for the survival of the masses – an artisan could not survive without the farmer and baker who produced her daily bread. Still, this institutional arrangement allowed people to maintain an illusion of self-sufficiency, that they survived based on the market value of their own labor.

If labor becomes economically redundant, this would no longer be possible. Moreover, it would be wasteful to force humans to remain in low-productivity jobs when autonomous machines become ever more productive. We would need new economic

institutions to share the output that our economy produces – institutions that do not depend on the market value of labor. At present, the share of output earned by labor is still larger than that of capital and other factors but is declining. Moreover, many experts predict that advances in AI may make it possible to build machines that are perfect substitutes for labor and that may make labor economically redundant within the current century. Building institutions that provide workers facing economic redundancy with alternative sources of income would insure and future-proof our society against such a scenario.

Ultimately, with the right economic institutions to distribute the abundant output in a world in which labor is made economically redundant by autonomous machines, we may be able to implement the dream of Arthur C. Clarke that “the goal of the future is full unemployment, so we can play.”

The figures that appear in this article are licensed by Anton Korinek under the Creative Commons Attribution 4.0 International License (CC BY 4.0). To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aghion, P., Jones, B., and Jones, C. (2019). Artificial intelligence and economic growth. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*. NBER and University of Chicago Press.
- Altman, S. (2021). Moore’s Law for everything. <https://moores.samaltman.com/>.
- Arrow, K. J. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 29(3):155–173.
- Atkinson, A. B. and Stiglitz, J. E. (1980). *Lectures on Public Economics*. McGraw-Hill, New York.
- Autor, D., Salomons, A., and Seegmiller, B. (2021). New frontiers: The origins and content of new work, 1940-2018. *Working Paper, MIT*.
- Bhutta, N. e. a. (2020). Changes in U.S. family finances from 2016 to 2019: Evidence from the survey of consumer finances. *Federal Reserve Bulletin*, 106(5).
- Boix, C. (2022). AI and the economic and informational foundations of democracy. In Bullock, J. et al., editors, *Oxford Handbook of AI Governance*. Oxford University Press.
- Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
- Brand, J. E. (2015). The far-reaching impact of job loss and unemployment. *American Review of Sociology*, 41:359–75.

- Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.
- CompaniesMarketCap.com (2022). Top public companies by earnings. <https://companiesmarketcap.com/most-profitable-companies/> [Last accessed: Feb 15th, 2022].
- Danaher, J. (2017). Will life be worth living in a world without work? technological unemployment and the meaning of life. *Science and Engineering Ethics*, 23(1):41–64.
- Fonseca, R., Kapteyn, A., Lee, J., and Zamarro, G. (2017). Does retirement make you happy? A simultaneous equations approach. In Wise, D. A., editor, *Insights in the Economics of Aging*, pages 339–372. University of Chicago Press.
- Forbes (2022). World’s billionaires list 2021. <https://www.forbes.com/billionaires/> [Last accessed: Feb 15th, 2022].
- Ford, M. (2015). *The Rise of the Robots*. Oneworld Publications, UK.
- Frank, R. H. (1991). Positional externalities. In Zeckhauser, R., editor, *Strategy and Choice: Essays in Honor of Thomas C. Schelling*, pages 25–47. MIT Press.
- Gallup (2022). How to improve employee engagement in the workplace. Technical report.
- George, H. (1879). *Progress and Poverty*. Appleton.
- Grace, K., Salvatier, J., Dafoe, A., Zhang, B., and Evans, O. (2018). Viewpoint: When will ai exceed human performance? evidence from ai experts. *Journal of Artificial Intelligence Research*, 62:729–754.
- Hicks, J. (1932). *The Theory of Wages*. Macmillan, London.
- Hoynes, H. and Rothstein, J. (2019). Universal basic income in the united states and advanced countries. *Annual Review of Economics*, 11(1):929–958.
- IACS (2012). Culture of american families: Tables from the national survey. Technical report, Institute for Advanced Studies in Culture.
- Kahneman, D., Sibony, O., and Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment*. Little, Brown Spark.
- Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103.
- Klinova, K. and Korinek, A. (2021). AI and shared prosperity. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES 2021)*.

- Korinek, A. (2019). The rise of artificially intelligent agents. *Working Paper, University of Virginia*.
- Korinek, A. (2020). Taxation and the vanishing labor market in the Age of AI. *Ohio State Technology Law Journal*, 16(1):244–257.
- Korinek, A. (2022). How new technologies affect labor markets and inequality: An impact assessment. *Working Paper, Washington, DC: Brookings*.
- Korinek, A. and Stiglitz, J. E. (2019). Artificial intelligence and its implications for income distribution and unemployment. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*. NBER and University of Chicago Press.
- Korinek, A. and Stiglitz, J. E. (2021a). Artificial intelligence, globalization, and strategies for economic development. *NBER Working Paper*, 28453.
- Korinek, A. and Stiglitz, J. E. (2021b). Steering technological progress. *Working Paper*.
- Kurzweil, R. (2005). *The Singularity Is Near*. Viking.
- Lee, K.-F. (2018). *AI Super-Powers: China, Silicon Valley, and the New World Order*. Houghton Mifflin Harcourt.
- Leontief, W. (1983). National perspective: The definition of problem and opportunity. In *The Long-term Impact of Technology on Employment and Unemployment: A National Academy of Engineering Symposium*. National Academies.
- Marinescu, I. (2017). No strings attached the behavioral effects of us unconditional cash transfer programs. Technical report, Roosevelt Institute.
- Maslow, A. (1943). A theory of human motivation. *Psychological Review*, 50(4):370–396.
- Meyer, B. D. and Sullivan, J. X. (2017). Consumption and income inequality in the u.s. since the 1960s. NBER Working Paper 23655.
- Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics Magazine*, 38(8).
- Pigou, A. C. (1920). *The Economics of Welfare*. McMillan.
- Piketty, T. and Saez, E. (2013). Optimal labor income taxation. In Auerbach, A. J., Chetty, R., Feldstein, M., and Saez, E., editors, *Handbook of Public Economics*, volume 5, pages 391–474.
- Popper, K. (1945). *The Open Society and Its Enemies*. Routledge.

- Rosen, S. (1986). The theory of equalizing differences. *Handbook of Labor Economics*, pages 641–692.
- Shapiro, C. (2008). Patent reform: Aligning reward and contribution. In Jaffe, A. B., Lerner, J., and Stern, S., editors, *Innovation Policy and the Economy, Volume 8*, pages 111–156. University of Chicago Press.
- Susskind, D. (2020). *A World Without Work: Technology, Automation, and How We Should Respond*. Metropolitan Books.
- Trammell, P. and Korinek, A. (2020). Economic growth under transformative AI. *GPI Working Paper 8-2020*.
- Turing, A. (1950). Computing machinery and intelligence. *Mind*, LIX(236):433–460.

A Proofs

Proof of Proposition 1. If $T + w < c_0$, meeting the subsistence level is not feasible since $c = w\ell + T < c_0$ even for the maximum amount of labor $\ell = 1$, as reflected in point 1. Otherwise, the optimality condition for labor is

$$wu'(w\ell + T) \leq v'(1 - \ell) \quad (13)$$

When $T \leq c_0$, this equation always holds with equality for an optimal level of ℓ that satisfies $\ell > (c_0 - T)/w > 0$. If $T > c_0$ and condition (13) holds as a strict inequality for $\ell = 0$, the extra income from work is not worth the effort, as reflected by the second part of the reservation wage function, and it is optimal for the agent not to work so $\ell = 0$, proving point 3. Otherwise, the optimality condition holds as an equality and pins down a positive optimal level of labor supply $\ell > 0$.

The effects of increases in T and w on utility can be obtained from differentiating the utility function (2) and applying the envelope theorem. The findings on the threshold and on labor supply result from differentiating (3) and implicitly differentiating (4). \square

Proof of Proposition 2. Assigning shadow value λ to the resource constraint in maximization problem (6) and assuming that a subsistence level of consumption for all agents is feasible ($\bar{Y} > Ic_0$), the planner’s first-order optimality conditions are

$$\theta^i u'(\cdot) = \lambda \quad (14)$$

$$\theta^i v'(\cdot) \geq \lambda F_\ell(\cdot) w^i \quad (15)$$

For part A. with $\theta^i \equiv 1 \forall i$, optimality condition (14) requires that all agents obtain the same level of consumption, which satisfies (8) according to the resource constraint,

and which pins down $\lambda = u'(\bar{c} - c_0)$. In combination with optimality condition (15), this implies that $\ell^i = 0$ as long as $v'(1) \geq u'(\bar{c} - c_0) F_\ell(\cdot) w^i$ or, equivalently,

$$w^i \leq \bar{w} := \frac{v'(1)}{u'(\bar{c} - c_0) F_\ell(\cdot)} \quad (16)$$

which takes a similar form as (3). Conversely, for $w^i > \bar{w}$, the optimality condition can be reformulated to (7), which is increasing in w^i , F_ℓ and $u'(\bar{c} - c_0)$, with the latter declining in aggregate income according to (8).

The inequality $w^i > \bar{w}$ is satisfied for all agents if it holds for the lowest-productivity value $w^i = w^{\min} = \min_i \{w^i\}$ at given output level Y and consumption level $\bar{c} = Y/I$. Substituting these two values in equation (16) and rearranging defines the threshold \underline{Y} as follows:

$$Y < \underline{Y} := I(u')^{-1} \left[\frac{v'(1)}{w^{\min} F_\ell(\cdot)} \right] + I c_0 \quad (17)$$

Conversely, $w^i \leq \bar{w}$ is satisfied for all i if it holds with equality for the highest-productivity agent with $w^i = w^{\max} = \max_i \{w^i\}$ when nobody is working so that $Y = Y_0$. Substituting these two equalities into (16) and rearranging delivers the threshold \bar{Y}_0 as follows,

$$Y_0 \geq \bar{Y}_0 := I(u')^{-1} \left[\frac{v'(1)}{w^{\max} F_\ell(0, m)} \right] + I c_0 \quad (18)$$

For the proof of part B., we use a variational argument $d\theta^i$ to account for how changes in each agent's valuation affect the described thresholds and the optimum amount worked. \square

Proof of Lemma 1. If $U(c^i, \ell^i, a^i) \in \mathcal{U}_0$, then by the definition of the range, there exists a z^i such that $U(T^i + z^i, 0, 0) = U(c^i, \ell^i, a^i)$ and vice versa. If $a^i \leq 0$ then observe that $U(c^i, 0, 0) > U(c^i, \ell^i, a^i)$ since utility is strictly decreasing in labor effort and increasing in amenities, which are non-positive here. Given that utility is also strictly increasing in consumption, this requires that $z^i < w^i \ell^i$. The compensating differential is increasing in labor productivity since the agent's utility is increasing in w^i . \square

Proof of Proposition 3. We expand the planner's optimization problem (6) by the extra amenity term defined in utility function (10) where $a^i = \alpha^i \ell^i$ and observe that the planner's optimality condition for labor is

$$v'(1 - \ell^i) \geq \alpha^i x'(\alpha^i \ell^i) + w^i \lambda F_\ell(\ell, m) \quad (19)$$

Setting $\ell^i = 0$, the left-hand side captures the marginal disutility of the first unit of labor, and the right-hand side defines a linear frontier in the space of (w^i, α^i) that is increasing in both parameters and reflects the marginal amenity value and productivity of work. The condition can also be expressed as $w^i \leq \bar{w}(\alpha^i) := [v'(1) - \alpha^i x'(0)] / \lambda F_\ell(\cdot)$, which is declining in α^i . When the inequality is violated for $\ell^i = 0$, equation (19)

holds as an equality and defines the optimum amount of work of individual i , which is increasing in both ℓ^i and α^i .

For part 2. of the proposition, observe that $\alpha^i > v'(1)/x'(0)$ implies that $\bar{w}(\alpha^i) < 0$, so it is optimal for individual i to work even at zero productivity because she derives sufficient amenity value from work. If this is the case, the individual also requires a compensating differential for work that is greater than her labor income. \square

Proof of Proposition 4. The optimization problem of an egalitarian planner is

$$\max_{\{c^i, \ell^i \geq 0\}_i} \sum_i U \left(c^i, \ell^i, \alpha^i \ell^i, \sum_j \bar{\alpha}^j \ell^j / I \right) \quad \text{s.t.} \quad \sum_i c^i = Y = F \left(\sum_i w^i \ell^i, m \right) \quad (20)$$

with associated optimality condition for work

$$v'(1 - \ell^i) \geq \alpha^i x'(a^i) + \bar{\alpha}^i \bar{x}'(\bar{a}) + w^i \lambda F_\ell(\cdot) \quad (21)$$

As in the proof of Proposition 3, by setting $\ell^i = 0$, the condition defines the frontier above which it is desirable for the agent to work, which is increasing in all three parameters $(w^i, \alpha^i, \bar{\alpha}^i)$.

By contrast, a private agent who maximizes utility (9) with perceived $\hat{a}^i = \hat{\alpha}^i \ell^i$ and obtains a competitive market wage $w^i F_\ell$ would arrive at the optimality condition for labor ℓ^i of

$$v'(1 - \ell^i) \geq w^i F_\ell u'(c^i) + \hat{\alpha}^i x'(\hat{a}^i)$$

Compared to the planner's optimality condition (21), the private agent's condition is missing the term $\bar{\alpha}^i \bar{x}'(\bar{a})$, which reflects the externalities that individual i 's work imposes on others, and is mis-perceiving the term $\alpha^i x'(a^i)$ as $\hat{\alpha}^i x'(\hat{a}^i)$. If $\bar{\alpha}^i > 0$ or $\alpha^i > \hat{\alpha}^i$, then the externalities or internalities imply that agent i faces insufficient incentives to work. If a specific subsidy s is added to the market wage, the resulting individual optimality condition is

$$v'(1 - \ell^i) \geq \hat{\alpha}^i x'(\hat{a}^i) + u'(\cdot) [w^i F_\ell(\cdot) + s]$$

The optimality conditions of the planner and the private agent coincide for $\lambda = u'(c^i)$ if the subsidy is set to the level indicated in the proposition, inducing the agent to internalize the externalities and internalities of her labor choice. \square

Proof of Proposition 5. As shown in the proof of Proposition 2, we find the following optimality condition when the optimal amount worked of the representative agent is interior, where $c = Y$ follows from the representative agent's resource constraint,

$$w = \frac{v'(1 - \ell)}{u'(Y - c_0) F_\ell(\ell, m)}$$

We apply the implicit function theorem to the agent's optimality condition to find

$$\frac{d\ell}{dY} = - \frac{w^i u''(\cdot) F_\ell(\ell, m)}{w^i u'(\cdot) F_{\ell\ell}(\ell, m) + v''(1 - \ell)} < 0$$

which is negative because all three second derivatives in the expression are negative. The effect on the marginal product of labor is

$$\frac{dF_\ell(\ell, m)}{dY} = F_{\ell\ell}(\ell, m) \cdot \frac{d\ell}{dY} > 0$$

In a competitive economy, the wage equals the marginal product, proving the proposition. \square