Steering Technological Progress *

Anton Korinek
UVA Economics, Darden and NBER

Joseph E. Stiglitz
Columbia and NBER

April 2020

Abstract

Rapid progress in new technologies such as Artificial Intelligence has recently led to widespread anxiety about potential job losses. This paper asks how we should guide innovative efforts so as to increase labor demand and create better-paying jobs, especially for middle-class workers. It takes as a premise that it is desirable to offer well-paying jobs to all able-bodied workers – either because jobs are important to provide meaning or for political economy reasons. We develop a theoretical framework to analyze the factors that make an innovation desirable from the perspective of workers, including its technological complementarity to labor, the factor share of labor in producing the goods involved, and the relative income of the affected workers. Examples of labor-friendly innovations are intelligent assistants who enhance the productivity of human workers. The paper also delineates what policy measures may steer technological progress in a desirable direction for workers, ranging from nudges for entrepreneurs to changes in tax, labor market and intellectual property policies to direct subsidies and taxes on innovation.

Keywords: technological progress, AI, inequality, redistribution
JEL Codes: E64, D63, O3

*We would like to thank Daniel Harper for excellent research assistance. We also gratefully acknowledge financial support from the Institute for New Economic Thinking (INET).

1 Introduction

Technology has advanced rapidly in recent years and is leading to widespread anxiety that it will soon make an increasing number of human professions redundant. Over the next decade or two, Frey and Osborne (2013) predict that 47% of US jobs could be automated. A substantial number of technologists go even further than this and predict that artificial intelligence will reach and then surpass human levels of general intelligence within the next several decades (see e.g. Kurzweil, 2005; Bostrom, 2014), enabling them to perform all jobs more cheaply than the subsistence cost of human labor, and threatening to make humans redundant in all economic activities. Such alarmist predictions are of course speculative and subject to considerable uncertainty. Nonetheless, they suggest that it may be a good idea for economists to think more carefully about how the direction of technological progress affects human well-being.

Our perspective is that technological progress does not happen by itself but is driven by human decisions on what, where, and how to innovate. It would be mis-placed to view our fate as pre-determined by blind technological forces and market forces that are beyond our control, as some techno-fatalists suggest. Our material condition is shaped jointly by the technological innovations that we humans create and by the social and economic institutions that we collectively design and within which these innovations take place. We as a society have the power actively steer the path of technological progress in AI so as to confront the challenges posed by our technological possibilities.

The central topic of this paper is thus how to steer progress in AI so as to increase demand for labor rather than displacing labor. We identify what the labor market effects of a given innovation are and how to categorize AI-based innovations according to their effects on labor demand. For this, it is necessary to pinpoint what the key conceptual properties of an innovation are that increase labor demand and therefore raise wages and employment. To provide two simple examples, AI-based intelligent assistants, such as "Directions" from Google Maps, complement and augment human labor – they allow unskilled workers with little geographical knowledge to take up jobs as drivers.

On the other hand, technologies such as Autonomous Vehicles from Waymo may predominantly substitute for workers and may lower demand for human labor.

Our premise is that it is desirable for the economy to offer well-paying jobs to all able-bodied workers, for two complementary reasons: First, jobs offer income, and from a political economy perspective, it is very difficult to sustain the large transfers that would be required if a significant part of the work force is displaced by AI and could no longer earn a living from work. Secondly, from a psychological perspective, jobs offer not only income but also identity, pride and meaning to workers.

The technical model setup that we develop builds on the approach to public economics of Atkinson and Stiglitz (1980), which solves for an optimal public policy while recognizing that the private agents subject to public policy interventions also maximize their individual objective functions. However, we specifically focus on how to apply the tools of public policy to steering technological progress in AI. In doing so, we build on recent descriptions of progress with emphasis on information technologies and AI, such as Greenwald and Stiglitz (2014), Baqaee and Farhi (2018), Acemoglu and Restrepo (2018) and Korinek and Stiglitz (2017, 2018).

Our main innovation over these existing works is to make technology endogenous and ask in which directions technological progress should be steered to make its effects on the factor owners (specifically workers) as beneficial as possible. We first describe conceptually which properties of a given innovation lead to increased labor demand. These include the innovation's complementarity to different types of human labor, the marginal utility of the affected workers compared to the rest of the population, and how much labor the workers are supplying. Then we assess the direction of technological progress in AI that markets are currently taking and are likely to take in the future in the absence of intervention, and we contrast these with what is desirable from a broader social perspective.

Our findings on how to steer progress in AI to maximize the positive impact on average workers is relevant in four specific domains: First, many entrepreneurs in the technology sector are eager to maximize the positive impact of their developments on mankind and will find it useful to obtain better guidance on the likely impact of their developments on income distribution. If such entrepreneurs put their minds to it, they can play an important role in guiding progress in a direction that is beneficial for the average worker. Second our findings are useful for unions and work councils that are interested in how to steer progress to the benefit of their members. Third, a significant part of AI research is either conducted or sponsored by government. Using our findings on the labor market implications of different types of innovations, such research can actively be steered in a direction that augments human labor rather than replacing it. Fourth, our work also highlights the important role that our tax system plays in steering technological progress: at present, labor is the most highly-taxed factor in our economic system, creating strong incentives for labor-saving innovation. One of the most natural public policy steps to steer progress in a direction that augments human labor is to reduce the burden of taxation on labor. Last but not least, our work also provides insights on how to actively provide economic incentives for innovative efforts to augment human labor.

2 Model

2.1 Setup

Consider an economy in which there are $i = 1, ...I$ agents, $j = 1, ..., J$ goods and $h = 1, ..., H$ factors of production. Each individual agent i has a utility function $u^i(c^i)$ over the vector of consumption $c^i = (c^{i1}, ..., c^{iJ})'$ of the J goods of the economy. Furthermore, each agent is born with a vector of factor endowments $\ell^i = (\ell^{i1}, ..., \ell^{iH})'$ that add up to a total factor endowment $\ell = \sum_i \ell^i$.

There is a representative firm that has access to a technology described by the production possibilities set $F(\ell; A)$ for a given vector of factor inputs ℓ and a vector of technological parameters $A = \left(A^1, \ldots, A^K\right)$, which capture in reduced form the state of technology in the economy. The firm's output vector $y = \left(y^1, ..., y^J\right)'$ thus satisfies

$$y \in F\left(\ell; A\right)$$

For now, we assume that the production technology exhibits constant-returns-to-scale in the factors ℓ and that the representative firm is competitive so that it earns zero profits in equilibrium and questions of ownership are irrelevant. (The case of decreasing returns can easily be subsumed by introducing a fixed factor "ownership" that earns any excess profits.) In the case of a single output good, we can denote the production technology using the more conventional format of a production function

$$y = F\left(\ell; A\right)$$

Finally, we assume that there is a planner in the economy who has a welfare function that weighs the utility of individual consumers with a set of weights $\{\theta^i\}$. W.l.o.g. we assume that the welfare weights are normalized so that $\sum_i \theta^i = 1$. This allows us to use the welfare weights to define a probability measure E_i. The planner's measure of social welfare can then be equivalently expressed either as a sum over all agents' utilities or as an expectation

$$\begin{aligned} W &= \sum_i \theta^i u^i\left(c^i\right) \\ &= E_i\left[u\left(c^i\right)\right] \end{aligned}$$

2.2 First Best

We start by analyzing what the first-best allocation in the described economy would look like. For this, we consider a social planner who maximizes social welfare under two important assumptions: First, the planner can directly choose the consumption allocations c^i for all the consumers i in the economy – this implicitly assumes that the planner has access to lump-sum transfers. Second, the planner can pick the technological parameters $A = \left(A^1, \ldots, A^K\right)$

in the economy, capturing in reduced form that the planner can, for example, make investments in basic research or provide other inducements to steer technological progress.

The planner's optimization problem is thus

$$\max_{c^i, A} W = \sum_i \theta^i u^i \left(c^i\right) \quad \text{s.t.} \quad \sum_i c^i \in F\left(l; A\right)$$

We assume that the technology parameters are specified such that the resulting maximization problem is concave, allowing for an interior solution. We then find:

Proposition 1 (First-best allocation)**.** *The planner chooses the consumption allocations and technology parameters in the economy such that they satisfy the optimality conditions*

$$\begin{aligned} \theta^i u' \left(c^i\right) &= \lambda \quad \forall i \\ \lambda \cdot F_{A^k}\left(\ell; A\right) &= 0 \quad \forall k \end{aligned}$$

Proof. The proof follows directly from taking the optimality conditions of the Lagrangian of the planner's maximization problem. □

The first optimality condition reflects that the planner simply distributes resources among the consumers so that their weighted marginal utility of consumption of each good is equated to the shadow price on the economy's resource constraint. The second optimality condition captures that each technology parameter is chosen so as to maximize the value of output at the given shadow price. Note that production efficiency can be pursued independently of distributive concerns – the planner simply maximzies output and then uses lump sum transfers to allocate it to consumers in a desirable manner.

2.3 Laissez Faire Equilibrium

In the laissez faire equilibrium, each agent i rents out her factor endowments at the prevailing rental rates $w = (w_1, ..., w_L)$ to earn a total factor income of $w\ell^i$, which she spends on purchasing the consumption goods c^i at the prevailing market price $p = \left(p^1, ..., p^J\right)$ where we normalize $p^1 = 1$ so good 1 is the numeraire. The problem of an individual consumer is thus

$$\max_{c^i} u^i\left(c^i\right) \quad \text{s.t.} \quad pc^i = w\ell^i$$

The representative firm rents the factors of production ℓ from the agents in the economy and picks the technology parameters A so as to maximize total profits

$$\max_{\ell, A} \Pi = p \cdot F\left(\ell; A\right) - w \cdot \ell$$

The equilibrium in the economy consists of a set of consumption allocations $\{c^i\}$, factor allocations $\{\ell^i\}$ and technological parameters A together with prices p and rental rates w such that all agents and the representative firm satisfy their optimization problem and goods and factor markets clear, i.e. $\sum_i c^i \in F\left(\ell; A\right)$ and $\sum_i \ell^i = \ell$.

Proposition 2 (Laissez-faire equilibrium)**.** *Under laissez-faire, the consumption allocations and technology parameters in the economy satisfy the optimality conditions*

$$u'\left(c^i\right) = \mu^i p \quad \forall i$$

$$p' F_\ell\left(\ell; A\right) = w \tag{1}$$

$$p \cdot F_{A^k}\left(\ell; A\right) = 0 \quad \forall k \tag{2}$$

Proof. The proof follows from taking the optimality conditions of the Lagrangian of private agents' and the firm's maximization problems. □

The first optimality condition reflects that each agent allocates consumption efficiently across the different goods of the economy; however, the overall distribution of wealth is determined by each agent's factor endowment, reflected in the agent's shadow value of wealth μ^i, and stands in no relationship to the welfare weights θ^i. The last optimality condition reflects that a decentralized will also pursue production efficiency – just like the planner in the first best.

Equilibrium with technology regulation For use below, let us also analyze the case in which the firm faces a linear tax vector τ on the choice of the technological parameters A. (W.l.o.g. we can always parameterize technology such that this specification of taxes is meaningful). Then the firm's profit objective can be rewritten as

$$\Pi = p \cdot F(\ell; A) - w \cdot \ell - \tau \cdot A$$

and the firm's optimality condition on A becomes

$$p \cdot F_{A^k}(\ell; A) = \tau^k \quad \forall k \tag{3}$$

Compared to optimality condition (2), the tax implies that the firm deviates from production efficiency because of the tax.

2.4 Constrained Planner

Let us now analyze the case of a constrained planner who is completely unable to redistribute between the agents of the economy. This serves as a benchmark to contrast to the first-best setup in section 2.2 and illustrate our basic insights in as simple of a setting as possbile. The real-world setting faced by most policymakers can be interpreted as an in-between of what is described in this section and the first-best in section 2.2. In this setup, the consumption of agent i is simply given by

$$c^i = w \cdot \ell^i = F_\ell(\ell; A) \cdot \ell^i \tag{4}$$

For simplicity (to avoid tensor notation), we assume that there is now just a single consumption good in the economy, although we continue to allow for multiple factors of production that are differentially owned by the agents of the economy.

The constrained planner with weights $\{\theta^i\}$ on individual utilities substitutes the implementability constraint (4) into her objective function and solves

$$\max_A W = \sum_i \theta^i u^i \left(F_\ell(\ell; A) \cdot \ell^i\right)$$

Proposition 3 (Constrained Optimum; No Redistribution)**.** *The constrained planner chooses the technology parameters of the economy such that they satisfy*

$$\sum_i \theta^i u^{i\prime}\left(c^i\right) F_{\ell A^k}(\ell; A) \cdot \ell^i = 0 \quad \forall k \tag{5}$$

Proof. The proof follows from taking the optimality conditions to the constrained planner's objective. □

Intuitively, the planner's sets the technological parameters such that she weighs the marginal effect of each technology parameter on the factor earnings of agent i, captured by $F_{\ell A^k}(\ell; A) \cdot \ell^i$, at the welfare weight and marginal utility of each agent i.

Implementation of Constrained Planner's Allocation To see how the constrained planner can implement this allocation, we identify the tax rates τ^k necessary so that expression (3) replicates the constrained planner's optimality condition (5). We find

Corollary 1 (Implementation of Constrained Optimum)**.** *To decentralize the constrained social optimum, a planner would set impose on the technological parameters the tax rates*

$$\tau^k = -\sum_h F_{\ell^h A^k}(\ell; A)\, E_i\left\{\ell^{hi} \cdot \left[u^{i\prime}\left(c^i\right) - E_i u^{i\prime}\left(c^i\right)\right]\right\} \quad \forall k \tag{6}$$

Intuitively, this tax rate takes into account how much the technological parameter k benefits each factor h, captured by the cross-derivative $F_{\ell^h A^k}$, how much of factor h a given agent i owns, and what the relative marginal utility of agent i is compared to the other agents. The planner will subsidize technological progress if it benefits factors that are owned by agents who have comparatively high marginal utility.

Proof. To obtain the tax formula above, we rewrite expression (3) as

$$F_{A^k}(\ell; A) = F_{\ell A^k}(\ell; A) \cdot \ell = \tau^k \quad \forall k$$

where we employed Euler's theorem in the first step. We then subtract equation (5) from the resulting expression to obtain

$$\tau^k = -\left(E_i\left[u^{i\prime}\left(c^i\right) F_{\ell A^k}(\ell; A) \cdot \ell^i\right] - E_i\left[u^{i\prime}\left(c^i\right)\right] F_{\ell A^k}(\ell; A) \cdot \ell\right)$$

Rearranging this expression and writing the vector product over the different factors h as a sum results in the tax formula (6). □

2.5 Discussion of Implementation

The tax formula presented in Proposition offers a sharp analytic description of how to steer technological progress when distribution is a concern, but it does so in an analytic framework that is relatively abstract. Let us thus discuss first a few tangible examples and then how to implement the proposed policies in practice. But before we do so a caveat: technological progress is by definition a step into the unknown, and the more fundamental an innovation, the more unknowns there will be in practice, and the more difficult it will be to apply the proposed policies. Nonetheless, for a great deal of innovative activity, we do have a sense of which factors will benefit and which factors will be hurt by it.

Examples of labor-using technologies

To make our description more concrete, let us discuss two specific examples of labor-using innovations in the space of artificial intelligence:

First, intelligent assistants are AI-powered devices that assist human workers and increase their productivity by complementing their cognitive capabilities. A specific application for which such assistants have recently been proposed are Augmented-Reality devices that help upskill lesser-skilled workers by providing them with specific instructions on how to perform cognitively intensive jobs. For example, such devices can assist factory workers perform complicated workflows that would otherwise require significant training. Another application are AI systems that provide call center workers with additional information about the callers in question, e.g. by analyzing the emotional content of voices. Even systems such as Google Maps can be interpreted as intelligent assistants that provide driving instructions to human drivers and thereby allow them to navigate routes in areas that they are not familiar with. (We also note a potential downside of intelligent assistants: they may actually lower the skill levels of workers because they make them dependent on the assistants, they may thus turn human workers that used to think for themselves more and more into "robots" that mechanically follow the instructions given by the assistant.)

Second, platforms that match labor services are another example of labor-using innovations. A number of high-tech corporations specialize in matching demand and supply for labor in the economy. Specific examples include Uber or Lyft, which match demand and supply for drivers, or Etsy, which matches demand and supply for artisan goods. Although there are justified concerns about the specifics of the jobs created e.g. by ride-sharing companies, these concerns do not change the fact that the platforms in question are labor-using and increase overall demand for labor in the economy. The specific concerns could be addressed by appropriate regulation.

Practical implementation Our findings on how to steer technological progress in the directions described by the "tax" formulas provide useful guidance in a number of distinct domains:

First, many innovators and entrepreneurs in the technology sector are eager to maximize the positive impact of their developments on society. At present, however, the impact of technological progress on labor markets and income distribution is all too often an afterthought for innovators. Publicly-spirited innovators will find it useful to be reminded of and obtain better guidance on the likely impact of their inventions. If the world's most creative innovators put their minds to it, they can play an important positive role in guiding progress in a direction that is beneficial for the average worker. Furthermore, innovators are perhaps also best-suited to predict the potential implications of their innovations and make better-informed decisions on what innovations to pursue to further the interests of workers.

Second, unions and works councils may have a say in which types of investments and innovations to pursue in their companies, and they may also be well-suited in judging the effects of specific innovations on workers. If they have the right to participate in the decision-making process, they will steer technological progress in a direction that is positive for their members. This is the precise opposite of the efforts of some corporations to make their workers as replaceable as possible in order to reduce workers' bargaining power.

Third, a significant part of AI research is either conducted or sponsored by government. Although this type of research is funded by the tax dollars of all workers, the government typically pays little attention to how the resulting innovations affect the livelihoods of workers. A natural public policy is thus to evaluate the likely effect of innovations on labor markets when determining what type of research the government should pursue.

Fourth, whether intentionally or unintentionally, our tax system plays an important in affecting the direction of technological progress: at present, labor is the most highly-taxed factor in our economic system, creating strong incentives for labor-saving innovation. One of the most natural public policy measures to steer progress in a direction that augments human labor is to reduce the burden of taxation on labor and instead subsidize human labor.

Last but not least, it may also be desirable to provide explicit economic incentives for innovative efforts to benefit human labor. If we can identify the labor market impact of innovations, we can directly tax or subsidize them as described in tax formula (6).

2.6 A Simple Example

Let us provide an example that captures the main intuition of the general model above in the simplest possible setting. Assume an economy in which there are only two types of individuals, capitalists and workers, who have log utility function and are labeled by $i = K, L$. The planner's social welfare function values the utility of the two agents according to

$$W = \theta^K \log c^K + \theta^L \log c^L$$

The two types are endowed with one unit of capital and labor, respectively. In the notation of our general model, this implies $\ell^K = (1,0)'$ and $\ell^L = (0,1)'$ so the economy's total factor endowment is $\ell = (1,1)$. Furthermore, there is only one final good produced using the two factors so $y = F(\ell; A)$. We also label the two factors by the subscripts $h = K, L$ and assume a CES production function

$$F(\ell; A) = \left[\left(a_K(A)\,\ell_K \right)^\rho + \left(a_L(A)\,\ell_L \right)^\rho \right]^{\frac{1}{\rho}}$$

with elasticity of substitution $\frac{1}{1-\rho} > 1$ so that the two factors are gross substitutes. In this formulation, A determines weight on capital- vs labor augmenting progress. The competitive factor rents for the two factors in this economy are

$$w_h = A_h^\rho \ell_h^{1-\rho} \left[(A_1 \ell_1)^\rho + (A_2 \ell_2)^\rho \right]^{1-\frac{1}{\rho}} = A_h^\rho \left(\ell_h / y \right)^{1-\rho}$$

We assume that the factor-augmenting productivity functions $a_h(A)$ are parameterized such that they define a locus w_K/w_L is increasing in A that is concave in the space of $(\log w_K, \log w_L)$. The constrained planner then chooses

$$\max_A \theta^K \log w_K + \theta^L \log w_L$$

Proposition 4. *The planner's optimal choice of A is a strictly increasing function of the planner's relative weight on capitalists versus workers θ^K/θ^L.*

Normalizing the welfare

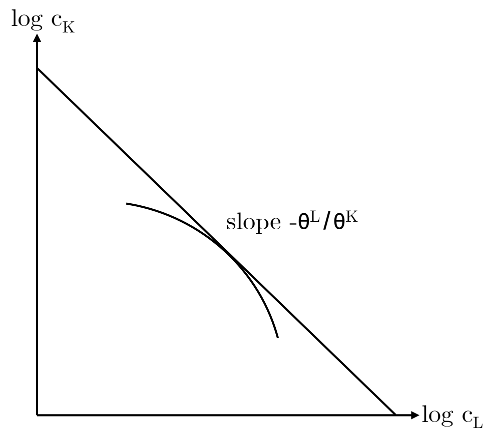

Figure 1: Innovation possibilities frontier and welfare isoquants

Proof. [to be written up] □

Intuitively, the more weight the planner places on the welfare of capitalists versus workers, the more she wants to invest in technological progress that makes the economy rely more on capital rather than labor. This pushes up demand for capital and the rent on capital, generating a redistribution from workers to capitalists.

The result is illustrate graphically in Figure 1. The concave line represents the innovation possibilities frontier as parameterized by A. Higher A corresponds to moving up and to the left along the frontier. The social planner's welfare isoquant at the optimum is represented by the downward-sloping straight line with slope $-\theta^L/\theta^K$. The optimum occurs where the innovation possibilities frontier is a tangent to the planner's highest possible welfare isoquant.

An increase in the planner's welfare weight on capitalists θ^K in this figure would correspond to a flattening of the welfare isoquants and would result in rotating the optimum along the innovation possibilities frontier such that the economy ends up with higher consumption for capitalists and lower consumption for workers.

3 Non-Monetary Benefits of Work

Work not only provides income but also imposes a number of other non-monetary benefits and costs. These include providing workers with a sense of identity and meaning as well as giving them status and social connections (see e.g. Korinek and Juelfs, 2020). If we are concerned with how to steer technological progress, then it makes sense to take into account these non-monetary factors as well. We expand our definition of agent's i utility to include an additional term for the non-monetary benefits of work,

$$U^i = u^i\left(c^i\right) + d^i \qquad \textbf{where} \qquad c^i = F_\ell\left(\ell; A\right) \cdot \ell^i,\ d^i = v\left(A\right) \cdot \ell^i$$

where $v(A)$ is a vector function that reflects how much the different technologies $A = \left(A^1, \ldots, A^k\right)$ affect the utility or disutility of providing the different factors contained in vector ℓ^i.

Laissez-Faire Equilibrium In a competitive market equilibrium, private firms will pay workers the marginal product of their labor as captured by optimality condition (1). The non-monetary benefits and costs of providing labor will add or subtract to the utility of a given workers but are not reflected in equilibrium wages as long as factors are supplied inelastically.

Constrained Planner's Solution The constrained planner by contrast recognizes the non-monetary effects and solves the optimization problem

$$\max_A \sum_i \theta^i \left[u^i\left(F_\ell\left(\ell; A\right) \cdot \ell^i\right) + v\left(A\right) \cdot \ell^i\right]$$

Her optimality condition on A^k can be expressed as

$$\underbrace{E_i\left[u^{i\prime}\left(c^i\right) F_{\ell A^k}\left(\ell; A\right) \cdot \ell^i\right]}_{\text{factor compensation}} + \underbrace{E_i\left[v_{A^k} \cdot \ell^i\right]}_{\text{non-monetary}} = 0 \tag{7}$$

This reflects that the planner combines the monetary effects of factor income on the different individuals of the economy with the non-monetary utility effects to find the optimum level of the technology parameters. A tax formula that is analogous to expression (6) can easily be derived,

$$\tau^k = -\sum_h F_{\ell^h A^k}(\ell; A)\, E_i\left\{\ell^{hi}\cdot\left[u^{i\prime}\left(c^i\right) - E_i u^{i\prime}\left(c^i\right)\right]\right\} - v_{A^k}\cdot E_i\left[\ell^i\right] \quad \forall k \tag{8}$$

The second term in this tax formula reflects that the planner would like to encourage a technology (imposing negative taxes, i.e. subsidies) the more non-monetary utility it provides to factor owners, where the weights on each agent's factor endowment ℓ^i is independent of agents' marginal utilities and is determined solely by the planner's welfare weights as $E_i\left[\ell^i\right] = \sum_i \theta^i \ell^i$.

3.1 Balance of Monetary and Non-Monetary Considerations

One question that is of particular interest is how the planner should balance the monetary and non-monetary effects of work. From equation (7) it can be seen that the monetary effects will carry greater relative weight the higher the marginal utility of the agents who are earning returns from a given factor – this is natural: the poorer an agent, the more the planner values greater resources for her.

Let us now push this observation a step further and consider a thought experiment in which each agent receives a homogenous lump sum transfer T in addition to her factor earnings so that $c^i = w\ell^i + T$, for example because a universal basic income is put in place. We can then observe the following:

Corollary. *The larger the lump sum transfer T, the more steering progress should focus on non-monetary factors.*

The result follows because the transfer raises the incomes of all agents and therefore reduces the marginal utility in the first term of expression (7). By implication, the second term becomes more and more important — the better we have addressed the material needs of all agents, the more we should focus on providing utility from non-monetary sources.

4 Conclusions

In recent decades, our economy has experienced a growing number of labor-saving innovations, and recent progress in AI risks accelerating the trend. Our systems of redistribution are only partially effective in countering this trend. Faced with these developments, this paper analyzes how to actively steer technological progress to have desirable distributive effects.

We discussed the basic economic properties of innovations that matter for their distributive desirability – the factor bias of innovations as well as the income levels and factor supply of the factor owners involved. But we have also left many interesting questions for future research. Among these are questions of how imperfect redistributive systems and steering technological progress should best interact with each other and what constitutes the optimal mix of the two. Moreover, it is important for income distribution how an innovation is distributed across the economy, e.g. whether it will be freely available or restricted by intellectual property rights.

Finally, we have also ventured into the question of how to consider the non-monetary factors of work – ultimately steering technological progress in a direction that maximizes social welfare should also focus on making work more fun, especially for lesser-paid workers for whom the market undervalues the non-monetary rewards of labor.

References

Acemoglu, D. and Restrepo, P. (2018). The race between machine and man: Implications of technology for growth, factor shares and employment. *American Economic Review*, forthcoming.

Atkinson, A. B. and Stiglitz, J. E. (1980). *Lectures on Public Economics.* McGraw-Hill, New York.

Baqaee, D. R. and Farhi, E. (2018). The microeconomic foundations of aggregate production functions. *NBER Working Paper No. 25293.*

Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies.* Oxford University Press.

Frey, C. B. and Osborne, M. (2013). The future of employment. *manuscript, Oxford Martin School.*

Greenwald, B. C. and Stiglitz, J. E. (2014). *Creating a Learning Society: A New Approach to Growth, Development, and Social Progress.* New York, NY: Columbia University Press.

Korinek, A. and Juelfs, M. (2020). Preparing for the (non-existent?) future of work. *UVA Darden Institute for Business in Society, Working Paper.*

Korinek, A. and Stiglitz, J. E. (2017). Artificial intelligence and its implications for income distribution and unemployment. *NBER Working Paper.*

Korinek, A. and Stiglitz, J. E. (2018). Artificial intelligence, worker-replacing technological progress and income distribution. *Working Paper.*

Kurzweil, R. (2005). *The Singularity Is Near.* Viking.

A Mathematical Proofs

(to be completed)