Steering Progress in Artificial Intelligence*

Anton Korinek†
Joseph E. Stiglitz‡

May 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. We develop a theoretical framework to identify the properties 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 discusses measures to 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. In the future, we find that progress should increasingly be steered to provide workers with utility from the non-monetary aspects of their jobs.

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).
†University of Virginia and Darden, NBER and Oxford Future of Humanity Institute
‡Columbia University, NBER and Roosevelt Institute

1 Introduction

Technological advances in recent years have led to widespread anxiety that progress 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 may be 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 ask how and in which directions to actively steer technological progress to make its effects on worker as beneficial as possible.

In our baseline model, we consider a framework of endogenous technological progress and assume a set of agents who differ in their factor endowments (e.g. capital and labor, or labor of different skill levels). We compare how a laissez-faire economy determines under which technologies the economy operates with what a social planner would choose who values the welfare of the differen agents according to defined weights. We assume the planner is unable to perform transfers between the agents of the economy but can shape the economy's technology as a second-best way to affect the factor earnings of the agents. The planner increases societal welfare by raising the demand for factors that are owned by relatively poor agents in the economy. Her optimal choice of technology depends on an 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 each of the workers is supplying. In performing this analysis, we also contribute to a long literature on endogenous and directed technological progress, going back to Ahmad (1966); Drandakis and Phelps (1966); Kennedy (1964) and Samuelson (1965). More recent works include Acemoglu (1998, 2002, 2010) and Acemoglu and Restrepo (2018).

We provide two examples for our general findings. First, we analyze the desirability of labor-augmenting versus capital-augmenting technological progress. For the empirically more plausible case that capital and labor are gross complements, a planner who attempts to raise wages finds it desirable to choose innovations that augment capital rather than labor. Second, we study an economy in which private agents or the planner need to determine what fraction of tasks to automate. The greater a planner's weight on workers, the less we find that she will automate.

We expand our basic findings to economies with multiple goods and identify two additional effects. First, a planner can increase social welfare by focusing technological progress on goods that are disproportionately consumed by relatively poorer agents, raising their real income. Second, the desirability of progress in a given sector for the owners of factors employed in the sector depends on the demand elasticity of goods in the sector. If demand for a good is relatively inelastic, then progress reduces factor demand in the sector.

Next we study how technological choices interact with market power and compare market outcomes with the choices of a social planner. When workers have market power, we find that profit-maximizing firms will pursue innovations that erode their market power by making workers more easily replaceable, even if this comes at the expense of lower production efficiency. A social planner who place sufficient weight on worker welfare will employ technologies that grant workers more market power but keeps them more efficient. Moreover, when employers have monopsony power over workers, they choose technologies that expand their monopsony power compared to what a social planner would do.

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 broader policy framework (including our tax system) plays in steering technological progress: at present, labor is the most highly-taxed factor in our economic system whereas the cost of capital has been kept low – perhaps artificially low – by more than a decade of expansionary monetary policy, creating strong incentives for labor-saving and capital-using 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 benefit workers.

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

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

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

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

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

$$\max_{c^i, A} W = \sum_i \theta^i u^i \left(c^i\right) \quad \text{s.t.} \quad \sum_i c^i \in F\left(\ell; A\right) \tag{1}$$

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 maximizes output and then uses lump sum transfers to allocate it to consumers in a desirable manner. However, there is by now a large literature explaining why the second welfare

theorem is not in general true. This paper can be thought of as expanding on those discussions in the context of endogenous technology.

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

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

where we define μ^i as the Lagrangian on the agent's budget constraint.

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*

$$u'\left(c^i\right) = \mu^i p \quad \forall i$$
$$p' F_\ell\left(\ell; A\right) = w \tag{2}$$
$$p \cdot F_{A^k}\left(\ell; A\right) = 0 \quad \forall k \tag{3}$$

The laissez-faire allocation is Pareto efficient.

Proof. The proof follows from taking the optimality conditions of the Lagrangian of private agents' and the firm's maximization problems. The decentralized optimality conditions replicate the conditions of the first-best for appropriately chosen welfare weights $\theta^i = 1/\mu^i$ and satisfy the same constraints; therefore the allocation is Pareto efficient. □

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 firm 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

$$p \cdot F_{A^k}(\ell; A) = \tau^k \quad \forall k \tag{4}$$

Compared to optimality condition (3), 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 unable to transfer resources between the agents of the economy – the only way to affect the

income distribution is via competitive factor returns, which depend on the chosen technology.[1] This setup serves as a benchmark to contrast to the first-best setup in section 2.2 and illustrate our basic insights in as simple a setting as possible. (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

$$c^i = w \cdot \ell^i = F_\ell\left(\ell; A\right) \cdot \ell^i \tag{5}$$

For simplicity (to avoid tensor notation), we consider the case in which there is 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 (5) into her objective function and solves

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

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

$$\sum_i \theta^i u^{i\prime}\left(c^i\right) F_{\ell A^k}\left(\ell; A\right) \cdot \ell^i = 0 \quad \forall k \tag{6}$$

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 $F_{\ell A^k}\left(\ell; A\right) \cdot \ell^i$, at the welfare weight and marginal utility of each agent i.

[1] The constrained planner's problem described below is isomorphic to the problem of a Ramsey planner who sets taxes or subsidies on the described choice variables and rebates (or raises) the associated revenue with lump sum transfers.

2.5 Implementation of Constrained Optimum

The proposition above offers a sharp analytic description of how to steer technological progress when distribution is a concern, but it does so in a framework that is quite abstract. In the following, we will therefore discuss how to implement the described solution. For analytic clarity, we will first describe – still at an abstract level – how to implement the constrained optimum if taxes/subsidies on innovative activity are available. Then we will discuss alternative avenues for how to make technological progress beneficial for workers. In the sections below, we will also discuss a few more tangible examples and 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. Even if policymakers can't ascertaion this, innovators might be able to. Moreover, it may be possible to guide innovation by committing to implement some of the proposed polices with ex post measures that are taken once the impact of an innovation is clear.

Implementation via Taxes/Subsidies To see how to implement the constrained optimal allocation via taxes/subsidies, we identify the tax rates τ^k necessary so that expression (4) replicates the constrained planner's optimality condition (6). 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*

$$\tau^k = -\sum_h F_{\ell^h A^k}\left(\ell; A\right) E_i\left\{\ell^{hi} \cdot \left[u^{i\prime}\left(c^i\right) - E_i u^{i\prime}\left(c^i\right)\right]\right\} \quad \forall k \tag{7}$$

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 (4) 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 (6) 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 (7). □

Alternative Avenues of Implementation Aside from such explicit economic incentives for innovative efforts, there are at least four additional avenues through which technological progress can be steered to benefit workers:

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 in-

vestments 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.

3 Applications

3.1 Application 1: Factor-Augmenting Progress

Our first application applies the intuition of the general model to a setting of factor-augmenting technological progress. For simplicity, assume an economy in which there is a single good and 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

$$W = \theta^K \log c^K + \theta^L \log c^L \tag{8}$$

where $\theta^K + \theta^L = 1$. 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 $\ell = (1,1)$. In a slight abuse of notation, we also label the two factors by the subscripts $h = K, L$ and assume one final good produced using a CES production function with $\rho \in (-\infty, 1) \{0\}$[2]

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

where a scalar technological parameter A captures how much progress augments capital as opposed to labor so $a'_K(A) > 0 > a'_L(A)$, i.e. an increase in A augments capital and diminishes labor. The elasticity of substitution in this specification is $\frac{1}{1-\rho}$. We observe that for $\rho < 0$ (or equivalently $\frac{1}{1-\rho} < 1$), the two factors are gross complements in production; conversely for $\rho > 0$, the two goods are gross substitutes in production.

The competitive factor prices of $h = \{K, L\}$ are

$$w_h(\ell; A) = \partial F(\ell; A)/\partial \ell_h = a_h(A)^\rho (y/\ell_h)^{1-\rho}$$

and the ratio of wages to capital rents is

$$\frac{w_L}{w_K} = \left(\frac{a_L(A)}{a_K(A)}\right)^\rho$$

A technological choice A that is relatively more capital-augmenting increases wages relative to capital rents if the two factors are gross complements ($\rho < 0$) but reduces wages relative to capital rents if the two factors are gross substitutes ($\rho > 0$).

The following assumption ensures that the set of innovation possibilites $\{a_K(A), a_L(A)\}$ is concave in the space of log factor prices so that the planner's optimization problem for the given social welfare function has a unique solution.

Assumption 1. *The two factor-augmenting technology functions $a_K(A)$ and*

[2] We rule out the Cobb-Douglas limit case $\rho = 0$ since labor and capital are always affected proportionately by factor-augmenting progress and no interesting distributive effects arise in that case.

$a_L(A)$ are such that the marginal rate of log factor price substitution satsisfies

$$MRLFPS = \frac{d\log w_K}{d\log w_L} < 0 \qquad and \qquad \frac{d^2\log w_K}{(d\log w_L)^2} < 0$$

The two inequalities imply that changes in A trace out a declining and concave log factor price frontier in the space of $(\log w_L, \log w_K)$, as illustrated in Figure 1.

The planner then solves

$$\max_A \theta^K \log w_K(\ell; A) + \theta^L \log w_L(\ell; A)$$

Proposition 4. *If factors are gross complements (substitutes), then the planner's optimal choice of A is a strictly increasing (decreasing) function of the planner's relative weight on workers versus capitalists θ^L/θ^K.*

Proof. The planner's optimality condition can be expressed as

$$MRLFPS = \frac{d\log w_K/dA}{d\log w_L/dA} = -\frac{\theta^L}{\theta^K}$$

Given the concavity of the log factor price locus, this first-order optimality condition describes the optimum. An increase in the planner's relative welfare weight θ^L on workers implies that the optimality condition is satisfied for a lower capital rent to wage ratio. If factors are gross complements (the empirically more plausible case), then a relatively higher wage can be achieved by a higher choice of A, i.e. by relatively more capital-augmenting technology; if factors are gross substitutes, then a relatively higher wage can be achieved by a lower choice of A, i.e. by relatively more labor-augmenting and less capital-augmenting technology. □

Intuitively, the more weight the planner places on the welfare of workers versus capitalists, the more she wants to gear technological progress in a direction that raises wages relative to capital rents and generates a redistribution from capitalists to workers. If the factors are gross complements, this can be

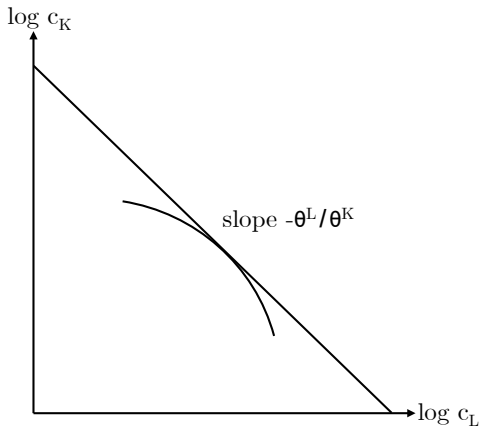


Figure 1: Innovation possibilities frontier and welfare isoquants

done by augmenting capital relatively more than labor; if they are complements, it requires augmenting labor relative to capital.

The result is illustrated graphically in Figure 1. The concave line represents the log factor rent frontier as parameterized by A. If the factors are gross complements, then higher A corresponds to moving down and to the right along the frontier; if they are gross substitutes, higher A corresponds to the opposite movement. 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^L in this figure would correspond to a steepening of the welfare isoquants and would rotate the optimum clock-wise along the log factor rent frontier such that the economy ends up with higher consumption for workers and lower consumption for capitalists.

Multi-sector economy In a multi-sector economy, an additional consideration concerns the elasticity of demand of the goods produced by the different sectors of the economy. If the demand elasticity for a good is less than unity, then innovation that makes production of the good more efficient will in fact reduce overall demand for the factors producing the good. Conversely, if the

demand elasticity for the good is sufficiently high, then greater efficiency in producing it will raise factor demand in the sector.

Examples of factor-augmenting technologies An example of labor-augmenting technologies are intelligent assistants, which are frequently cited in the AI space as a product that holds promise for improving the productivity of workers. They are AI-powered devices that assist human workers and increase their productivity by complementing their cognitive capabilities. A specific application of such assistants are Augmented-Reality devices that help to upskill lesser-skilled workers by providing them with specific instructions on how to perform cognitively intensive jobs. 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.[3]

Platforms that match labor services can be interpreted as another example of labor-augmenting innovations. A number of high-tech corporations specialize in matching demand and supply for labor in the economy. Important examples are Uber or Lyft, which match demand and supply for drivers.[4] Taxi drivers used to spend significant amounts of time looking for jobs, and the matching efficiency of these platforms has enabled them to deliver more driving services in the same amount of time. Other examples include MTurk, which matches demand and supply for tasks that human workers can perform digitally, and Etsy, which matches demand and supply for artisan goods.

[3] We also note an important 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.

[4] There are also justified concerns about other aspects of the jobs created by ride-sharing companies; these concerns are in addition to the effects of these platforms on labor demand and could be addressed separately by appropriate regulation.

3.2 Application 2: Automation of Tasks

Our next application considers a setup that centers on the question of task automation in a framework that is inspired by Zeira (1998) and Acemoglu and Restrepo (2019). We assume that final output is produced using a unit mass $j \in [0,1]$ of intermediate goods or tasks according to the production function $\log y = \int_0^1 \log y(j)\, dj$, where each individual task is performed using either capital K or labor L. Specifically, we assume the variable $A \in [0,1]$ captures what fraction of the tasks is automated so they are performed using capital and $y(j) = K(j)$ for $j \leq A$. The remaining fraction $(1-A)$ reflects all the tasks that are not automated and are performed using labor so $y(j) = L(j)$ for $j > A$. As in our previous application, we assume that capital is owned exclusively by capitalists and labor by workers. We denote the fractions of the economy's factor endowments of capital and labor by $K/L = \alpha/(1-\alpha)$.

Within the set of automated and non-automated tasks, it is optimal to allocate capital and labor symmetrically. For given K and L, this implies that $y(j) = {}^{K}/_{A}$ for $j \leq A$ and $y(j) = {}^{L}/_{1-A}$ for $j > A$. The aggregate production of the economy can then be expressed as

$$F(K, L; A) = \left(\frac{K}{A}\right)^{A} \left(\frac{L}{1-A}\right)^{1-A} \tag{9}$$

Lemma 1. *Production efficiency implies that a fraction $A = \alpha$ of tasks is automated.*

Proof. Taking the log and maximizing the production function (9) delivers the optimality condition $\frac{A}{\alpha} = \frac{1-A}{1-\alpha}$ which is satisfied for $A = \alpha$. □

The result of the lemma will hold both in the first-best allocation and in the decentralized equilibrium of the economy. By contrast, production efficiency will generally no longer hold if we are concerned with the distributive implications of automation and if direct transfers are not available. To see this, we solve the problem of a second-best planner who maximizes the welfare function (8).

Proposition 5. *The second-best planner chooses a degree of automation A strictly between the welfare weight θ^K on capitalists and the fraction α of the factor endowment that is capital. An increase in the welfare weight on workers reduces the optimal degree of automation.*

Proof. See appendix. □

Intuitively, the fraction A of automated tasks also represents the share of output that is earned by capitalists in a decentralized setting. A planner who places greater weight on workers will reduce automation, which increases the fraction of tasks available for workers and raises their share of output. However, deviating from production efficiency reduces the total amount of output. At the optimum, the planner weighs off the desired redistribution with the associated loss in production efficiency.

4 Steering Progress under Imperfect Competition

We now consider the effects of market power on incentives for steering technological progress.

Market power in factor supplies Consider an economy with a single final good in which consumer-worker i supplies factors elastically subject to a convex disutility $d\left(\ell^i\right)$. Moreover, assume consumer i has market power in some of the factor markets $h = 1 \dots H$, captured by the fraction $\alpha^i = \left(\alpha^{i1}, \dots, \alpha^{iH}\right)$ that the consumer-worker is supplying of each factor h, with $\alpha^{ih} > 0$ for some h. The consumer's maximization problem is then

$$\max_{c^i} u^i\left(c^i\right) - d\left(\ell^i\right) \quad \text{s.t.} \quad c^i = w\left(\alpha^i \circ \ell^i + (1-\alpha) \circ \ell^{\setminus i}\right) \cdot \ell^i$$

where $\alpha^i \circ \ell^i$ captures element-by-element multiplication and $\ell^{\setminus i}$ reflects the factor supply by all agents other then agent i. The agent's optimality condition

is then

$$w\left(1+\alpha^{i}\circ\epsilon_{w,\ell}^{i}\right)=\frac{d'\left(\ell^{i}\right)}{u'\left(c^{i}\right)}$$

which defines a

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 6 (Steering Progress and Employee Market Power)**.** *The greater the weight θ^{L} placed on workers, the less the planner will erode the market power of workers.*

Proof. See appendix. □

5 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). Moreover, factors such as how much autonomy and control workers have over their working conditions and schedule matter greatly for their life satisfaction. 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,

$$U^i = u^i\left(c^i\right) + d^i \qquad \textbf{with} \qquad d^i = v\left(A\right) \cdot \ell^i \tag{10}$$

where $v\left(A\right)$ 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 factors contained in vector ℓ^i.

First-best In the first-best, a planner maximizes the weighted sum of utilities $E_i\left[U^i\right]$ as specified above in (10) subject to the resource constraint $\sum_i c^i = F\left(\ell; A\right)$, where we assume a single consumption good. The optimality condition for the choice of technology parameter A^k is

$$E_i\left[v_{A^k}\left(A\right) \cdot \ell^i\right] = -\lambda F_{A^k}$$

It is optimal to modify the choice of technology such that the marginal non-monetary benefits of employing labor are weighed against the marginal cost in terms of productivity, where the non-monetary benefits are calculated taking into account each agent's welfare weight. In other words, it is optimal to forgo some productivity gains and employ less productive productive technologies if they provide non-monetary benefits to workers.

Laissez-Faire Equilibrium In a competitive market equilibrium, on the other hand, private firms will pay workers solely based on the marginal (monetary) product of the factors supplied, as captured by optimality condition (2). The non-monetary benefits and costs of providing factors will add to or subtract from the utility of factor owners but are not reflected in equilibrium wages in our framework (since factors are supplied inelastically). As a result, the privately optimal choice of technology in the laissez-faire equilibrium completely disregards the non-monetary effects on factor owners, including workers.

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

$$\max_{A} \sum_{i} \theta^i \left[u^i \left(F_\ell \left(\ell; A \right) \cdot \ell^i \right) + v \left(A \right) \cdot \ell^i \right]$$

We readily find the following result:

Proposition 7 (Constrained Optimum with Non-monetary benefits)**.** *The constrained planner chooses the technology parameters of the economy to balance the effects on factor compensation and non-monetary benefits such that*

$$\underbrace{E_i \left[u^{i\prime} \left(c^i \right) F_{\ell A^k} \left(\ell; A \right) \cdot \ell^i \right]}_{\textit{factor compensation}} + \underbrace{E_i \left[v_{A^k} \cdot \ell^i \right]}_{\textit{non-monetary}} = 0 \tag{11}$$

Proof. The proof follows from taking the optimality condition to the constrained planner's maximization problem. □

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 analogous to expression (7) can easily be derived,

$$\tau^k = -\sum_{h} F_{\ell^h A^k} \left(\ell; A \right) E_i \left\{ \ell^{hi} \cdot \left[u^{i\prime} \left(c^i \right) - E_i u^{i\prime} \left(c^i \right) \right] \right\} - v_{A^k} \cdot E_i \left[\ell^i \right] \quad \forall k \tag{12}$$

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 $E_i \left[\ell^i \right] = \sum_i \theta^i \ell^i$.

In summary, a planner who considers the non-monetary benefits and costs of technology on factor owners, esp. workers, will give weight to these considerations in her choice of technology in a way that will be disregarded by

private markets, wasting the potential for large welfare gains.

5.1 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 (11) 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 in place. We can then observe the following:

Corollary. *The larger the monetary 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 (11). 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.

6 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. With multiple goods, additional considerations include the demand elasticity for goods that are produced more efficiently as well as their share in the consumption basket of poor versus rich households. 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. (1998). Why do new technologies complement skills? directed technical change and wage inequality. *Quarterly Journal of Economics*, 113(4):1055–1089.

Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies*, 69(4):781–809.

Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy*, 118(6):1037–1078.

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.

Acemoglu, D. and Restrepo, P. (2019). Artificial intelligence, automation and

work. In Agrawal, A. e. a., editor, *The Economics of Artificial Intelligence*. NBER and University of Chicago Press.

Ahmad, S. (1966). On the theory of induced invention. *Economic Journal*, 76(302):344–357.

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.

Drandakis, E. M. and Phelps, E. S. (1966). A model of induced invention, growth, and distribution. *Economic Journal*, 76(304):832–840.

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.

Kennedy, C. (1964). Induced bias in innovation and the theory of distribution. *Economic Journal*, 74(295):541–547.

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.

Samuelson, P. (1965). A theory of induced innovation along kennedy-weisacker lines. *Review of Economics and Statistics*, 47(4):343–356.

Zeira, J. (1998). Workers, machines, and economic growth. *Quarterly Journal of Economics*, 113(4):1091–1117.

A Mathematical Proofs

(to be completed)