

Steering Technological Progress*

Anton Korinek[†] and Joseph E. Stiglitz[‡]

August 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 to 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.

*We would like to thank Ken Lipartito and participants in the CIFAR Workshop on Innovation, Equity, and the Future of Prosperity and the 2nd INET/IMF Conference on Macroeconomics in the Age of AI for helpful comments. Daniel Harper provided excellent research assistance. We also gratefully acknowledge financial support from the Institute for New Economic Thinking (INET).

[†]University of Virginia, Economics and Darden, NBER, CEPR and Oxford Future of Humanity Institute GovAI

[‡]Columbia University, NBER and Roosevelt Institute

Keywords: technological progress, AI, inequality, redistribution

JEL Codes: E64, D63, O3

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 human labor economically redundant. 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. 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. Moreover, 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.

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 complement and augment human labor – for example, navigation systems allow unskilled workers with little geographical knowledge to take up jobs as drivers. On the other hand, technologies such as Autonomous Vehicles 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 different 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, how much labor each of the workers is supplying, and how costly it is to deviate from the first-best choice of technology. 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, the desirability of advancing technology in a given sector also depends on the demand elasticity for goods in the sector. If demand for a good is relatively inelastic, then progress reduces factor demand in the sector. Second, a planner can increase social welfare by focusing technological progress on goods that are disproportionately consumed by relatively poorer agents, raising their real income.

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.

Lastly, we consider how to balance the monetary and non-monetary costs and benefits of work in steering technological progress. We find that firms do not sufficiently account for the non-monetary aspects of work, including how much meaning, satisfaction, identity, fulfillment or social engagement jobs provide because these factors are not efficiently priced by the market. A planner would find it desirable to include such non-monetary considerations in steering technological progress. Moreover, we show that, as the monetary income of workers rises, the planner would increasingly shift her attention to non-monetary considerations.

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

There is also 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 $A = (A^1, \dots, A^K) \in \mathbb{R}^K$, which capture in reduced form the state of technology in the economy, what investments in R&D have been made, etc. The firm's output vector $y = (y^1, \dots, y^J)'$ 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 remainder of Section 2, we consider the case of a single consumption good $J = 1$. We will investigate the additional considerations that arise with multiple goods below in Section 4. In the case of a single good, we can denote the production technology using the more conventional format of a production function

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

2.2 First Best

We start by analyzing the first-best allocation in the described economy. We consider a social planner who maximizes social welfare in the economy, given by the weighted sum of utility of individual consumers, with an exogenous 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 and an associated expectations operator E_i . Social welfare can then be equivalently expressed as a sum over all agents' utilities or as an expectation

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

The planner's choices are (i) to pick the technological parameters $A = (A^1, \dots, A^K)$ in the economy and (ii) to directly choose the consumption allocations c^i for all the consumers i in the economy – equivalent to the capacity to perform lump-sum transfers.

The planner's optimization problem is thus

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

This formulation highlights that the planner's choice of technology and the consumption allocation can be performed in two separate steps. The first step is the following.

Definition 1 (Production Efficiency). For given factor endowment ℓ , we denote the set of efficiency-maximizing technological parameters $A^*(\ell)$ and the associated level of output $y^*(\ell)$ so that

$$A^*(\ell) = \arg \max_A F(\ell; A) \quad \text{and} \quad y^*(\ell) = F(\ell; A^*) \quad (2)$$

For brevity of notation we will omit the argument ℓ on A^* and y^* unless required for clarity. If the technology parameters are specified such that $F(\ell; A)$

is continuously differentiable and concave in A and the maximization problem in (2) has an interior optimum, then production efficiency is described by

$$F_A(\ell; A) = 0$$

Proposition 1 (First-best allocation). *For given welfare weights and factor endowments, the planner chooses the technology parameters in the economy to achieve production efficiency. She chooses the consumption allocations such that they exhaust production and satisfy the optimality conditions*

$$\theta^i u^i(c^i) = \lambda \quad \forall i$$

Proof. The first part follows because if production efficiency was not satisfied, it would be easy to increase welfare by moving to a more efficient technology choice. The second part follows from taking the optimality conditions of the Lagrangian of the planner's maximization problem. \square

The planner simply distributes resources among consumers so that their weighted marginal utilities of consumption are equated – and equal the shadow price on the economy's resource constraint.

The proposition reflects that production efficiency can be pursued independently of distributive concerns – the planner simply maximizes output and then transfers 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 a good guide for public policy. 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 $w = (w_1, \dots, w_L)$ to earn a total factor income of $w \cdot \ell^i$, which she consumes. The problem of an individual consumer is thus

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

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 of the economy and picks the technology parameters A so as to maximize total profits

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

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

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

$$F_\ell(\ell; A) = w \tag{3}$$

$$F_A(\ell; A) = 0 \tag{4}$$

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. \square

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 decen-

tralized firm will also pursue production efficiency – just like the planner in the first best.

We denote the factor shares s_ℓ earned by the different factors ℓ in the economy by

$$s_\ell(\ell; A) = \frac{F_\ell(\ell; A) \circ \ell}{F(\ell; A)}$$

where the operator \circ represents the element-by-element (Hadamard) product of the two factors.

2.4 Constrained Planner

Let us now analyze a constrained planner with weights $\{\theta^i\}$ on individual utilities who is unable to perform outright transfers between the agents of the economy – the only way to affect the income distribution is via competitive factor returns, which depend on the choice of 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.) The consumption of agent i is

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

The constrained planner substitutes the implementability constraint (5) into her objective function and solves

$$\max_A W = \sum_i \theta^i u^i (F_\ell(\ell; A) \cdot \ell^i) \tag{6}$$

For the following proposition, we assume that the planner’s optimization problem is concave in A and has an interior solution.

¹The 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) any associated revenue with lump sum transfers to the same set of agents from whom it was obtained.

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

$$\sum_i \theta^i u^i(c^i) F_{\ell A}(\ell; A) \cdot \ell^i = 0 \quad (7)$$

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

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

2.5 Decomposing the Effects of Technological Change

A constrained planner's choice of technology generically deviates from the benchmark of production efficiency that prevails in both the first best and the decentralized equilibrium. Let us now characterize the trade-off between efficiency and redistribution a bit further.

A useful conceptual benchmark is a technological choice that leaves total output unaffected:

Definition 2 (Efficiency-Neutral Technological Change). For given factor inputs ℓ , the technology parameter A in the production function $\bar{F}(\ell; A)$ represents an efficiency-neutral technology choice if $\bar{F}(\ell; A) = y^* \forall A$.

We use the convention of denoting production functions that admit an efficiency-neutral technology choice by bars, e.g. $\bar{F}(\ell; A)$. When A is an efficiency-neutral technology choice, then the technology parameter A does not affect the overall level of output, but it may well impact the *marginal products* and thus the competitive factor rents F_ℓ and factor shares s_ℓ . This property would allow a planner who faces an efficiency-neutral technology choice to pursue her distributive objectives without incurring any efficiency costs. It also enables us to perform the following decomposition:

Lemma 1 (Decomposition of (Marginal) Technological Change). *For given factor inputs ℓ , the effects of a marginal technological change dA on factor returns F_ℓ can be decomposed into an efficiency-neutral redistribution between factors that satisfies $\bar{F}_{\ell A} \cdot \ell = 0$ and a proportional scale parameter on all factor returns so that*

$$F_{\ell A} = \underbrace{(\quad)}_{\text{redistribution}} \bar{F}_{\ell A} + F_\ell \cdot \underbrace{\frac{F_A}{F}}_{\text{scale par.}}$$

Proof. Define $\bar{F}_{\ell A} = F_{\ell A} - F_\ell \cdot F_A/F$ and observe that

$$\bar{F}_{\ell A} \cdot \ell = F_{\ell A} \cdot \ell - F_\ell \cdot \ell \frac{F_A}{F} = F_A - F_A = 0$$

Note that the last step applies Euler's theorem to each of the two terms of the sum, i.e. $F_A = F_{\ell A} \cdot \ell$ and $F = F_\ell \cdot \ell$. \square

We can employ this decomposition to re-formulate the technology choice (7) of a constrained social planner in terms of the traditional equity-efficiency trade-off:

$$\underbrace{E_i [u^{i'}(c^i) \bar{F}_{\ell A} \cdot \ell^i]}_{\text{marg. redistributive effect}} = \underbrace{F_A E_i [u^{i'}(c^i) c^i / F]}_{\text{marg. efficiency effect}}$$

The left-hand side of this expression distills the redistributive effects of technology choice – reflected in the efficiency-neutral terms $\bar{F}_{\ell A}$. The right-hand side captures only the efficiency effects of the technology choice – reflected in the overall change in output F_A converted into units of weighted average marginal utility.

Lemma 2 (Decomposition of (Discrete) Technological Change). *For given factor inputs ℓ , the effects of a technological change ΔA on factor returns F_ℓ can be decomposed into an efficiency-neutral redistribution between factors that satisfies $\Delta \bar{F}_\ell \cdot \ell = 0$ and a proportional scale parameter on all factor returns*

so that

$$\Delta F_\ell = \underbrace{\Delta \bar{F}_\ell}_{\text{redistribution}} + F_\ell \cdot \underbrace{\frac{\Delta F}{F}}_{\text{scale par.}}$$

Proof. Define $\Delta \bar{F}_\ell = \Delta F_\ell - F_\ell \cdot \Delta F/F$ and observe (using Euler's theorem) that

$$\Delta \bar{F}_\ell \cdot \ell = \Delta F_\ell \cdot \ell - F_\ell \cdot \ell \frac{\Delta F}{F} = \Delta F - \Delta F = 0$$

□

2.6 Implementation of Constrained Optimum

Let us now consider how to implement the constrained optimum in a decentralized setting. Assume that the representative 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 profits are

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

and the firm's optimality condition on A becomes

$$F_A(\ell; A) = \tau \tag{8}$$

Compared to optimality condition (4), the tax implies that the firm deviates from production efficiency because of the tax.

To see how to implement the constrained optimal allocation via taxes/subsidies, we identify the tax τ necessary so that expression (8) replicates the constrained planner's optimality condition (7). We find

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

$$\tau = -F_{\ell A} \cdot E_i \{ \ell^i [u^i(c^i) - E_i u^i(c^i)] \} \tag{9}$$

Proof. We use Euler’s theorem to rewrite expression (8) as

$$F_A(\ell; A) = F_{\ell A}(\ell; A) \cdot \ell = \tau$$

We then subtract equation (7) from the resulting expression to obtain

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

Rearranging this expression results in the tax formula (9). □

Intuitively, the tax rate takes into account how much the technological parameter A benefits or hurts each factor h , captured by the cross-derivative $F_{\ell A}$, how much of each factor 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, on average, it benefits factors that are owned by agents who have comparatively high marginal utility.

2.7 Discussion

Our results above offer a sharp analytic description of how to steer technological progress when distribution is a concern. Although we acknowledge the practical difficulties in following this approach, we view our results as a useful guidepost for what direction of technological change is desirable in at least four different settings.

First, many innovators and entrepreneurs in the technology sector are eager to maximize the positive impact of their developments on society. At present, there is a great deal of focus on how AI developers can avoid discrimination, biases, etc. – even if it comes at the expense of somewhat reducing their profit margins (see e.g. Dubber et al., 2020). 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 on workers. 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. Moreover, it may also counteract the tendency of management to automate workers because machines are seen as easier to manage and maintain, even if such a move comes at the expense of production efficiency.²

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 all workers. A natural public policy is to evaluate the likely labor market effects of innovations when determining what type of research the government should pursue or fund.

Fourth, the tax formulas that we derived above would be the most direct instruments to guide technological progress in a desirable direction. However, more generally, our tax system plays an outsized role in affecting the direction of technological progress – whether intentionally or unintentionally: at present, labor is the most highly-taxed factor in our economic system, creating strong incentives for labor-saving innovation (see e.g. Acemoglu et al., 2020). 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 or to even subsidize human labor.

²In a first-best world, such over-automation would be penalized by a reduction in managerial compensation. However, in a world with agency frictions, managers may use their discretion to advantage their well-being at the expense of their workers and shareholder.

3 Applications

3.1 Factor-Augmenting Progress

Our first application considers factor-augmenting technological progress. In its pure form, factor-augmenting progress implies that the same amount of output can be produced using less input of the augmented factor. Examples of labor-augmenting innovation include intelligent assistants that enable a given worker to perform her duties more efficiently, or more efficient techniques of managing workers so that a given amount of labor can effectively provide more labor services. An example of capital-augmenting innovation is the progress in semiconductor technology that is captured by Moore’s Law, whereby a given quantity of silicon chips can perform ever more computation.

Formally, consider a constant-returns production function $F(\ell; A)$ and assume that the technology parameter A determines how much the chosen technology augments the different factors, i.e. for each factor j , the parameter A implies a level of factor-augmenting technology $a_j(A)$ so that $a_j(A)\ell_j$ effective units enter the production function. We collect the factor augmentation functions $a_j(A)$ in a vector function $a(A) = (a_j(A))_{j=1}^J$ and denote the vector of effective units of the different factors that enter the production function by $a(A) \circ \ell$, where \circ is the Hadamard (element-by-element) product. The production function can then be denoted as $F(\ell; A) = F(a(A) \circ \ell)$.

Two-Factor CES Production Function with Factor-Augmenting Technology A typical CES production technology with two factors, say $h = K, L$, and factor-augmenting technology is

$$y = F(a(A) \circ \ell) = [(a_K(A)\ell_K)^\rho + (a_L(A)\ell_L)^\rho]^{\frac{1}{\rho}} \quad (10)$$

where $\rho \in (-\infty, 1) \setminus \{0\}$ implies an elasticity of substitution between the two factors of $\frac{1}{1-\rho}$. For $\rho < 0$ (or equivalently $\frac{1}{1-\rho} < 1$), the two factors are gross complements in production – as is assumed by the majority of the literature on capital-labor substitutability. Conversely for $\rho > 0$ (equivalently $\frac{1}{1-\rho} > 1$), the

two factors are gross substitutes in production.³ The technology parameter A captures how much the technology augments capital $a_K(A)$ and labor $a_L(A)$. W.l.o.g. let us define the technology parameter A directly as the augmentation of one of the factors, say capital so $a_K(A) = A$, and assume $a'_L(A) < 0$ so that the choice of A reflects a trade-off between augmenting capital versus augmenting labor.

For concreteness, consider an economy in which there are only two types of individuals, capitalists K and workers L , who are endowed with one unit of capital and labor, respectively. In a slight abuse of notation, we label them by $i = K, L$ and denote their endowments $\ell^K = (1, 0)'$ and $\ell^L = (0, 1)'$ so the economy's total factor endowment is $\ell = (1, 1)$.

The competitive factor rents of $h = K, L$ are

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

Varying the parameter A traces out the economy's innovation possibilities frontier in the space (w_L, w_K) . We assume $a_L(A)$ is specified such that the frontier is convex. An example is given by the solid line in Figure 1. The ratio of wages to capital rents is

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

If the two factors are gross complements ($\rho < 0$), an increase in A raises wages relative to the returns on capital and corresponds to a movement to the right and downwards on the factor price frontier. If the two factors are gross substitutes ($\rho > 0$), the opposite results apply and an increase in A reduces wages relative to the returns on capital, corresponding to a movement up and to the left along the innovation possibilities frontier.

Let us also characterize the labor-augmentation coefficient $a_L(A)$ that implies that the technology choice A is *efficiency-neutral* for this production

³We do not consider the Cobb-Douglas case $\rho = 0$ since labor and capital are always affected proportionately by factor-augmenting progress in that case and no interesting distributive effects arise.




Figure 1: Innovation possibilities frontier and welfare isoquants

function. Output (10) remains constant at a given level \bar{y} as we vary $a_K = A \in [0, \bar{y}]$ if and only if we set labor augmentation $a_L(A)$ to satisfy $\bar{y}^\rho = A^\rho + a_L(A)^\rho$ or equivalently

$$a_L(A) = [\bar{y}^\rho - A^\rho]^{\frac{1}{\rho}} \in [0, \bar{y}]$$

This satisfies $a'_L(A) < 0$. The resulting innovation possibilities frontier exhibits $dw_K/dw_L = -1 \forall A$ and is indicated by the dashed line in Figure 1.

A constrained planner who optimizes the social welfare function (6) solves

$$\max_A \theta^K u^K(w_K(\ell; A)) + \theta^L u^L(w_L(\ell; A)) \quad (11)$$

Graphically, the planner chooses the welfare isoquant (convex curve) that forms a tangent on the economy's innovation possibilities frontier, guaranteeing the highest level of welfare possible. This level is characterized by the following version of condition (7),

$$\frac{\theta^K u^{K'}(c^K)}{\theta^L u^{L'}(c^L)} = - \frac{w_{LA}(\ell; A)}{w_{KA}(\ell; A)}$$

Proposition 4. *If factors are gross complements ($\rho < 0$), then the planner's optimal choice of A is a strictly increasing function of the planner's relative*

weight on workers versus capitalists θ^L/θ^K . If factors are gross substitutes, the opposite results apply.

Proof. See discussion above. □

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 done by augmenting capital relatively more than labor; if they are gross complements, it requires augmenting labor relative to capital. The planner's willingness to deviate from the first-best solution depends both on the efficiency cost (captured by the distance between the innovation possibilities frontier and the dashed line) and the curvature of her welfare isoquants, i.e. her desire to redistribute.

An increase in the planner's welfare relative 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 innovation possibilities frontier, such that the economy ends up with higher consumption for workers and lower consumption for capitalists.

Examples of factor-augmenting technologies One example of a labor-augmenting technology is intelligent assistants, which are frequently cited by AI developers as holding promise for improving the productivity of workers. These are AI-powered devices that assist workers and increase their productivity by complementing their cognitive capabilities. A specific example of such assistants are Augmented-Reality devices that help to upskill lesser-skilled workers by providing them with 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, e.g. by analyzing the emotional content of voices. Even navigation systems can be interpreted as intelligent assistants that augment

human drivers and allow them to navigate more efficiently and to navigate routes in areas that they are not familiar with.⁴

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. An important example are ride-sharing platforms, which match demand and supply for drivers.⁵ 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 (while also devaluing their human capital). 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.

Whether labor-augmenting technological progress ultimately benefits workers depends on the elasticity of substitution/demand for labor, as highlighted in our proposition. If the elasticity is less than unity, then productivity increases raise the effective supply of labor by more than they raise demand for it, resulting in lower returns to labor. This is what seems to have happened e.g. in the ride-sharing market. If the elasticity is above unity, returns to labor rise.

3.2 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

⁴We 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.

⁵There are justified concerns about 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.

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

$$F(K, L; A) = \left(\frac{K}{A}\right)^A \left(\frac{L}{1-A}\right)^{1-A} \quad (12)$$

Lemma 3. *Production efficiency implies that a fraction $A = \alpha$ of tasks is automated. This is what is replicated by the laissez faire equilibrium.*

Proof. Maximizing the log production function (12) delivers the optimality condition $\frac{A}{\alpha} = \frac{1-A}{1-\alpha}$ which is satisfied for $A = \alpha$. Proposition 2 implies that the same holds in the laissez faire equilibrium. \square

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 (11).

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. \square

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.

Examples of Task Automation Task automation occurs when a machine acts as a perfect substitute for a task in a productive process that was previously performed by labor. A tangible example is an assembly line that consists of a series of steps performed by humans, and a machine is introduced to perform one of them. Choosing the level of automation A in our analytic framework then corresponds to deciding how many of the steps are automated and how many are performed by humans. In the given example, tasks are perfect complements in the sense that each task along the assembly line is required in fixed proportion to produce the output.

More generally, tasks may also be combined in a more elastic fashion. When machines substitute for tasks performed with labor and simultaneously increase the productivity at which the task is performed, then this can also be thought of as task-augmenting progress akin to the factor augmentation in the previous section. If tasks are gross complements (with elasticity less than one), then augmenting the automated task will benefit other tasks performed by labor. For example, if doctors produce health services both by diagnosing and by providing advice in a complementary fashion, then automating diagnosis and making it more efficient may actually increase their returns from providing advice. If the elasticity is below one, the opposite result applies. In the production function (12) above, tasks are combined in Cobb-Douglas fashion, i.e. with unitary elasticity, so no such effects occur.

Acemoglu and Restrepo (2018) observe that the fraction A of automated tasks can be affected in two ways: by changing the degree of automation of existing tasks or by creating new tasks. Accordingly, choosing a lower degree of automation than the decentralized equilibrium can be achieved not only by holding back the automation of existing tasks performed by human labor but

also by inventing new human-focused tasks that are performed by labor.

4 Multiple Goods

In a multi-sector economy, one of the important questions is at what sectors innovative efforts should be targeted. This section expands our baseline model to include a role for multiple goods and sectors of production. First, we focus on how to steer progress across multiple different sectors when consumers have homothetic preferences so that the relative demand for goods is independent of income distribution and distributive consequences of technological change arise solely from changes in factor prices. Then we analyze how to steer progress when consumers have different consumption baskets so that changing relative goods prices redistributes real income.

4.1 Homothetic Preferences

Multi-sector economy A crucial factor in this question is consumers' elasticity of substitution between the goods produced by the different sectors of the economy. (These findings are analogous to our analysis of the elasticity of substitution in production in section 3.1.) If the elasticity is less than unity, then innovation that makes the production of a good more efficient will in fact reduce overall demand for the factors producing the good. Conversely, if the elasticity is greater than one, then greater efficiency in producing a good will raise demand for the good and ultimately factor demand in the sector.

We assume an economy in which there is a continuum of consumers who each supply differentiated labor to intermediate goods sectors that are combined into a final good according to a CES production function

$$F(\{y(j)\}) = \left[\int_0^1 y(j)^\nu dj \right]^{\frac{1}{\nu}}$$

[to be completed]

4.2 Different Consumption Baskets

When consumer preferences are heterogeneous, the distribution of income interacts with relative goods prices: a planner can actively use innovations to change relative goods prices and redistribute real income among the agents of the economy; moreover, changes in incomes affect relative demand for the different goods of the economy, which in turn impacts factor earnings differentially.

[to be completed]

5 Steering Progress under Imperfect Competition

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

5.1 Specialization and Labor's Market Power

The following application captures firms' tradeoff of how specialized of a production process they choose versus how much market power their hires will enjoy. In general, highly specialized production processes may yield significant productivity gains but also imply that the firms rely on specialized and/or highly skilled labor, which enjoys greater market power than undifferentiated unskilled labor.

Consider an economy with a single final good and a unit mass $i \in [0, 1]$ of agents who are consumer-workers. Each agent i derives CES utility from consumption $u(c) = c^{1-\sigma}/(1-\sigma)$ and elastically supplies specialized labor of type i subject to a disutility $d(\ell) = \ell^{1+\psi}/(1+\psi)$ with Frisch elasticity ψ . We assume that $\sigma < 1$ so that the substitution effects from wage changes dominate any income effects.

There is a representative firm in the economy that hires labor $h \in [0, 1]$ for

a unit mass of tasks and combines them according to the production function

$$y = A(\eta) \int_0^1 (\ell^h)^{1-\alpha} dh$$

The parameter $\eta \in [0, 1]$ reflects the degree of specialization of labor that the firm chooses for the production process and simultaneously drives how much market power workers enjoy. We assume that $A(\eta)$ is strictly increasing and concave in η , i.e. specialization makes production more efficient but at decreasing speed. Moreover, we assume that the range is $A(\eta) = [\underline{A}, \overline{A}]$, and that the function satisfies the two Inada conditions $\lim_{\eta \rightarrow 0} A'(\eta) = \infty$ and $A'(1) = 0$. However, the downside for the representative firm is that more specialization gives more monopoly power to workers. At $\eta = 0$, productivity is at its lowest level \underline{A} and labor is completely unspecialized, so all types of labor are perfect substitutes and individual workers do not have any market power. Conversely, at $\eta = 1$, productivity is at its highest level \overline{A} , but each type of labor i is specific for a particular task $h = i$ so each agent i enjoys significant monopoly power. Intermediate levels of specialization imply that there is some limited substitutability between different types of labor. For example, at $\eta = 1/2$, each task $h \in [0, 1]$ can be accomplished by precisely two agents $i, j \in [0, 1]$, and the two supply labor in Cournot fashion so, in a symmetric equilibrium, they internalize that each supplies a fraction $\eta = 1/2$ of the labor within each of their sectors of employment.

The optimization problem of consumer-worker i is thus given by

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

where $\ell^{\setminus i}$ denotes the supply of labor by all agents other than agent i in the agent's sectors of employment. The agent's optimality condition is then

$$w(1 - \eta \epsilon_{w, \ell}) = \frac{d'(\ell^i)}{u'(c^i)} = (\ell^i)^\psi (c^i)^\sigma$$

where $\epsilon_{w, \ell} = -\frac{dw}{d\ell} \cdot \frac{\ell^h}{w}$ is the inverse demand elasticity for labor of firms, which

reflects by what percentage wages need to go down for firms to demand one percent more labor. This defines an inverse demand relationship $w(\ell^i; \eta)$ with the derivative w.r.t. specialization

$$\frac{\partial w}{\partial \eta} = \frac{\epsilon_{w,\ell} \cdot (\ell^i)^\psi (c^i)^\sigma}{(1 - \eta \epsilon_{w,\ell})^2} > 0 \quad (13)$$

The representative firm hires labor and picks the technology parameters $A(\eta)$ to maximize total profits

$$\max_{\ell, \eta, A(\eta)} \Pi = A(\eta) \int_0^1 (\ell^h)^{1-\alpha} dh - \int_0^1 w^h(\eta) \ell^h dh$$

Observe that the representative firm is small and has no effect on the overall labor demand faced by each agent i . Therefore it acts competitively in labor markets in the sense that the wage does not depend on the quantity of labor that it hires. However, the firm internalizes that the degree of monopoly power enjoyed by the labor that it is hiring is endogenous and depends on its choice of specialization η , as captured by a wage function $w^h(\eta)$ for each variety h . The firm's optimality condition of labor for a given degree of specialization η is

$$(1 - \alpha) A(\eta) (\ell^h)^{-\alpha} = w^h(\eta)$$

which implies an inverse demand elasticity $\epsilon_{w,\ell} = -\frac{dw}{d\ell} \cdot \frac{\ell}{w} = \alpha$. In a symmetric equilibrium, the optimal choice of specialization can be rewritten as

$$A'(\eta) \ell^{1-\alpha} = w'(\eta) \ell$$

The left-hand side captures the marginal efficiency gain from specialization and is strictly decreasing in η from infinity to zero. The right-hand side reflects the marginal rise in labor costs associated with greater specialization, where $w'(\eta)$ is given by equation (13) and is increasing in η . The condition therefore yields a unique solution for the optimum level of specialization.

Proposition 6 (Steering Progress and Employee Market Power). *The greater*

the weight θ^L placed on workers, the more specialized the production technology that the planner will employ.

Proof. See appendix. □

Discussion A tangible example of this result is that firms have incentives to de-skill jobs so that workers are more replaceable and have less bargaining power. If a given worker is the only one who can do a certain job, she can extract significant surplus; if anyone can do the job, then workers are perfect substitutes and are paid competitive wages. For example, the introduction of highly standardized production processes, say the conveyor belt or work procedures in the fast food industry, can be interpreted along these lines. More generally, this result reflects that there may be a broad set of innovations that do not increase productivity but that make jobs more undifferentiated and unskilled so as to reduce workers' bargaining power.

5.2 Monopsony Power in Factor Markets

Next we consider a setup in which firms have monopsony power in factor markets. We assume that each factor h is supplied by a single type of consumer-worker with CES consumption utility $u(c) = c^{1-\sigma} / (1 - \sigma)$ who elastically supplies type h labor subject to a disutility cost $d^h(\ell) = \ell^{1+\psi} / (1 + \psi)$ with Frisch elasticity ψ . The resulting optimization problem is $\max_{\ell^h} u(w^h \ell^h) - d^h(\ell^h)$, with optimality condition $w^h = d'(\ell^h) / u'(c^h)$. This gives rise to an inverse labor supply function

$$w^h(\ell^h) = (\ell^h)^{\frac{\psi+\sigma}{1-\sigma}}$$

with elasticity $\epsilon_{w,\ell}^h = \frac{\psi+\sigma}{1-\sigma}$. We assume that $\sigma < 1$ so that the substitution effects from wage changes dominate any income effects.

Assume a set of oligopsonistic firms, for which the extent of market power in factor markets is described by a vector $\alpha^m = (\alpha^{m1}, \dots, \alpha^{mH})$, where each α^{mh} captures what fraction of the demand for factor h derives from the firm.

The optimization problem of firm m is

$$\max_{A, \ell^m} F(\ell^m; A) - w(\ell^m + L) \cdot \ell^m$$

where we denote by L the labor demand from all other firms and observe that for each factor h , we find $\alpha^{mh} = \frac{\ell^{mh}}{\ell^{mh} + L^h}$. The firm's optimality condition for labor demand equates marginal product to marginal revenue,

$$F_\ell(\cdot) = w + w'(\cdot) \ell^m = w(1 + \alpha^m \epsilon_{w, \ell})$$

Similarly, the firm's optimal choice of technology is given by

$$F_A(\cdot) = 0$$

Specific examples of the ways in which firms increase their monopsony power in labor markets are (i) to put no-compete clauses in employment contracts, which prevent other employers in the same sector to compete for them and (ii) to provide training to workers in ways that are not easily portable to other firms.

6 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,

$$U^i = u^i(c^i) + d^i \quad \mathbf{with} \quad d^i = v(A) \cdot \ell^i \quad (14)$$

where $v(A)$ is a vector function that reflects how much the different technologies $A = (A^1, \dots, A^k)$ 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 $E_i [U^i]$ as specified above in (14) subject to the resource constraint $\sum_i c^i = F(\ell; A)$, where we assume a single consumption good. The optimality condition for the choice of technology parameter A^k is

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

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 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 (3). 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

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

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*

$$\underbrace{E_i [u^{i'}(c^i) F_{\ell A^k}(\ell; A) \cdot \ell^i]}_{\text{factor compensation}} + \underbrace{E_i [v_{A^k} \cdot \ell^i]}_{\text{non-monetary}} = 0 \quad (15)$$

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

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 (9) can easily be derived,

$$\tau^k = - \sum_h F_{\ell^h A^k}(\ell; A) E_i \{ \ell^{hi} \cdot [u^{i'}(c^i) - E_i u^{i'}(c^i)] \} - v_{A^k} \cdot E_i [\ell^i] \quad \forall k \quad (16)$$

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

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.

6.1 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 (15) 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 (15). 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.

7 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.

More generally, technological progress is by definition always 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 ascertain this, innovators might be able to. And it may also be possible to guide innovation by committing to implement some of the proposed policies with ex post measures that are taken once the impact of an innovation is clear.

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., Manera, A., and Restrepo, P. (2020). Does the u.s. tax code favor automation? *Brookings Papers on Economic Activity*, forthcoming.
- 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.
- Baqae, 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.
- Dubber, M. D., Pasquale, F., and Das, S. (2020). *Oxford Handbook of Ethics of AI*. Oxfo.
- 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)