

Artificial Intelligence, Automation and Work*

Daron Acemoglu
MIT

Pascual Restrepo
Boston University

January 4, 2018

Abstract

We summarize a framework for the study of the implications of automation and AI on the demand for labor, wages, and employment. Our task-based framework emphasizes the displacement effect that automation creates as machines and AI replace labor in tasks that it used to perform. This displacement effect tends to reduce the demand for labor and wages. But it is counteracted by a productivity effect, resulting from the cost savings generated by automation, which increase the demand for labor in non-automated tasks. The productivity effect is complemented by additional capital accumulation and the deepening of automation (improvements of existing machinery), both of which further increase the demand for labor. These countervailing effects are incomplete. Even when they are strong, automation increases output per worker more than wages and reduce the share of labor in national income. The more powerful countervailing force against automation is the creation of new labor-intensive tasks, which reinstates labor in new activities and tends to increase the labor share to counterbalance the impact of automation. Our framework also highlights the constraints and imperfections that slow down the adjustment of the economy and the labor market to automation and weaken the resulting productivity gains from this transformation: a mismatch between the skill requirements of new technologies, and the possibility that automation is being introduced at an excessive rate, possibly at the expense of other productivity-enhancing technologies.

Keywords: AI, automation, displacement effect, labor demand, inequality, productivity, tasks, technology, wages.

JEL classification: J23, J24.

*Prepared for *Economics of Artificial Intelligence*, edited by Ajay Agarwal, Avi Goldfarb and Joshua Gans. We are grateful to David Autor for useful comments. We gratefully acknowledge financial support from Toulouse Network on Information Technology, Google, Microsoft, IBM and the Sloan Foundation.

1 INTRODUCTION

The last two decades have witnessed major advances in artificial intelligence (AI) and robotics. Future progress is expected to be even more spectacular and many commentators predict that these technologies will transform work around the world (Brynjolfsson and McAfee, 2012; Ford, 2016; Boston Consulting Group, 2015; McKinsey, 2017). Recent surveys find high levels of anxiety about automation and other technological trends, underscoring the widespread concerns about their effects (Pew Research Center, 2017).

These expectations and concerns notwithstanding, we are far from a satisfactory understanding of how automation in general, and AI and robotics in particular, impact the labor market and productivity. Even worse, much of the debate in both the popular press and academic circles centers around a false dichotomy. On the one side are the alarmist arguments that the oncoming advances in AI and robotics will spell the end of work by humans, while many economists on the other side claim that because technological breakthroughs in the past have eventually increased the demand for labor and wages, there is no reason to be concerned that this time will be any different.

In this essay, we build on Acemoglu and Restrepo (2016), as well as Zeira (1998) and Acemoglu and Autor (2011) to develop a framework for thinking about automation and its impact on tasks, productivity, and work.

At the heart of our framework is the idea that automation and thus AI and robotics replace workers in tasks that they previously performed, and via this channel, create a powerful *displacement effect*. In contrast to prevailing presumptions in much of macroeconomics and labor economics, which maintain that productivity-enhancing technologies always increase overall labor demand, the displacement effect can reduce the demand for labor, wages and employment. Moreover, the displacement effect implies that increases in output per worker arising from automation will not result in a proportional expansion of the demand for labor. The displacement effect causes a decoupling of wages and output per worker, and a decline in the share of labor in national income.

We then highlight several countervailing forces, which push against the displacement effect and may imply that automation, AI, and robotics could increase labor demand. First, the substitution of cheaper machines for human labor creates a *productivity effect*: as the cost of producing automated tasks declines, the economy will expand and increase the demand for labor in non-automated tasks. The productivity effect could manifest itself as an increase in the demand for labor in the same sectors undergoing automation or as an increase in the demand for labor in non-automating sectors. Second, *capital accumulation* triggered by increased automation (which raises the demand for capital) will also raise the demand for labor. Third, automation does not just operate at the extensive margin—replacing tasks previously performed by labor—but at the intensive

margin as well, increasing the productivity of machines in tasks that have already been automated. This phenomenon, which we refer to as *deepening of automation*, tends to create a productivity effect but no displacement, and thus increases labor demand.

Though these countervailing effects are important, they are generally insufficient to engender a “balanced growth path,” meaning that even if these effects were powerful, ongoing automation would still reduce the share of labor in national income (and possibly employment which tends to be linked to the labor share). We argue that there is a more powerful countervailing force that increases the demand for labor as well as the share of labor in national income: the *creation of new tasks*, functions and activities in which labor has a comparative advantage relative to machines. The creation of new tasks generates a *reinstatement effect* directly counterbalancing the displacement effect.

Indeed, throughout history, we have not just witnessed pervasive automation, but a continuous process of new tasks creating new employment opportunities for labor. As tasks in textiles, metals, agriculture and other industries were being automated in the 19th and 20th centuries, a new range of tasks in factory work, engineering, repair, back-office, management and finance generated demand for displaced workers. The creation of new tasks is not an autonomous process advancing at a predetermined rate, but one whose speed and nature are shaped by the decisions of firms, workers and other actors in society, and which might be fueled by new automation technologies. First, this is because automation, by displacing workers, may create a greater pool of labor that could be employed in new tasks. Second, the currently most discussed automation technology, AI itself, can serve as a platform to create new tasks in many service industries.

Our framework also highlights that even with these countervailing forces, the adjustment of an economy to the rapid rollout of automation technologies could be slow and painful. There are some obvious reasons for this related to the general slow adjustment of the labor market to shocks, for example, because of the costly process of workers being reallocated to new sectors and tasks. Such reallocation will involve both a slow process of searching for the right matches between workers and jobs, and also the need for retraining, at least for some of the workers.

A more critical, and in this context more novel, factor is a potential *mismatch between technology and skills*—between the requirements of new technologies and tasks and the skills of the workforce. We show that such a mismatch slows down the adjustment of labor demand, contributes to inequality, and also reduces the productivity gains from both automation and the introduction of new tasks (because it makes the complementary skills necessary for the operation of new tasks and technologies more scarce).

Yet another major factor to be taken into account is the possibility of *excessive automation*. We highlight that a variety of factors (ranging from a bias in favor of capital in the tax code to labor market imperfections create a wedge between the wage and the

opportunity cost of labor) and will push towards socially excessive automation, which not only generates a direct inefficiency but also acts as a drag on productivity growth. Excessive automation could potentially explain why, despite the enthusiastic adoption of new robotics and AI technologies, productivity growth has been disappointing over the last several decades.

Our framework underscores as well that the singular focus of the research and the corporate community on automation, at the expense of other types of technologies including the creation of new tasks, could be another factor leading to a productivity slowdown, because it forgoes potentially valuable productivity growth opportunities in other domains.

In the next section, we provide an overview of our approach without presenting a formal analysis. Section 3 introduces our formal framework, though to increase readability, our presentation is still fairly non-technical (and formal details and derivations are relegated to the Appendix). Section 4 contains our main results, highlighting both the displacement effect and the countervailing forces in our framework. Section 5 discusses the mismatch between skills and technologies, potential causes for slow productivity growth and excessive automation, and other constraints on labor market adjustment to automation technologies. Section 6 concludes, and the Appendix contains derivations and proofs omitted from the text.

2 AUTOMATION, WORK, AND WAGES: AN OVERVIEW

At the heart of our framework is the observation that robotics and current practice in AI are continuing what other automation technologies have done in the past: using machines and computers to substitute for human labor in a widening range of tasks and industrial processes.

Production in most industries requires the simultaneous completion of a range of tasks. For example, textile production requires production of fiber, production of yarn from fiber (e.g., by spinning), production of the relevant fabric from the yarn (e.g., by weaving or knitting), pre-treatment (e.g., cleaning of the fabric, scouring, mercerizing and bleaching), dyeing and printing, finishing, as well as various auxiliary tasks including design, planning, marketing, transport, and retail.¹ Each one of these tasks can be performed by a combination of human labor and machines. At the dawn of the British Industrial Revolution, most of these tasks were heavily labor-intensive (some of them were merely performed). Many of the early innovations of that era were aimed at automating spinning and weaving by substituting mechanized processes for the labor of skilled artisans (Mantoux, 1928).²

¹See <http://textileguide.chemsec.org/find/get-familiar-with-your-textile-production-processes/>

²It was this displacement effect that motivated Luddites to smash textile machines and agricultural

The mechanization of US agriculture offers another example of machines replacing workers in tasks they previously performed (Rasmussen, 1982). In the first half of the 19th century, the cotton gin automated the labor-intensive process of separating the lint from the cotton seeds. In the second half of the 19th century, horse-powered reapers, harvesters, and plows replaced manual labor working with more rudimentary tools such as hoes, sickles and scythes, and this process was continued with tractors in the 20th century. Horse-powered threshing machines and fanning mills replaced workers employed in threshing and winnowing, two of the most labor-intensive tasks left in agriculture at the time. In the 20th century, combine harvesters and a variety of other mechanical harvesters improved upon the horse-powered machinery, and allowed farmers to mechanically harvest several different crops.

Yet another example of automation comes from the development of the factory system in manufacturing and its subsequent evolution. Beginning in the second half of the 18th century, the factory system introduced the use of machine tools, such as lathes and milling machines, replacing the more labor-intensive production techniques relying on skilled artisans (Mokyr, 1990). Steam power and later electricity greatly increased the opportunities for the substitution of capital for human labor. Another important turning point in the process of factory automation was the introduction of machines controlled via punch cards and then numerically-controlled machines in the 1940s. Because numerically-controlled machines were more precise, faster, and easier to operate than manual technologies, they enabled significant cost savings, while also reducing the role of craft workers in manufacturing production. This process culminated in the widespread use of CNC (computer numerical control) machinery, which replaced the numerically-controlled vintages (Groover, 1983). A major new development was the introduction of industrial robots in the late 1980s, which automated many of the the remaining labor-intensive tasks in manufacturing, including machining, welding, painting, palletizing, assembly, material handling, and quality control (Ayres and Miller, 1983; Groover et al. 1986; Graetz and Michaels, 2015; Acemoglu and Restrepo, 2017).

Examples of automation are not confined to industry and agriculture. Computer software has already automated a number of tasks performed by white-collar workers in retail, wholesale, and business services. Software and AI-powered technologies can now retrieve information, coordinate logistics, handle inventories, prepare taxes, provide financial services, translate complex documents, write business reports, prepare legal briefs, and diagnose diseases. They are set to become much better at these tasks during the next

workers during the Captain Swing riots to destroy threshing machines. Though these workers often appear in history books as misguided, there was nothing misguided about their economic fears. They were quite right that they were going to be displaced. Of course, had they been successful, they might have prevented the Industrial Revolution from gaining momentum with potentially disastrous consequences for technological development and our subsequent prosperity.

several years (e.g., Brynjolfsson and McAfee, 2012; Ford, 2016).

As these examples illustrate, automation involves the substitution of machines for labor and leads to the displacement of workers from the tasks that are being automated. This displacement effect is not present—or present only incidentally—in most approaches to production functions and labor demand used in macroeconomics and labor economics. The canonical approach posits that production in the aggregate (or in a sector for that matter) can be represented by a function of the form $F(AL, BK)$, where L denotes labor and K is capital. Technology is assumed to take a “factor-augmenting” form, meaning that it multiplies these two factors of production as the parameters A and B do in the production function we wrote down.

It might appear natural to model automation as an increase in B , that is, as capital-augmenting technological change. However, this type of technological change does not cause any displacement and always increases labor demand and wages (see Acemoglu and Restrepo, 2016). Moreover, as our examples above illustrate, automation is not mainly about the development of more productive vintages of existing machines, but involves the introduction of new machinery to perform tasks that were previously the domain of human labor.

Labor-augmenting technological change, corresponding to an increase in A , does create a type of displacement if the elasticity of substitution between capital and labor is small. But in general, this type of technological change also expands labor demand, especially if capital adjusts over the long run (see Acemoglu and Restrepo, 2016). Moreover, our examples make it clear that automation does not directly augment labor; on the contrary, it transforms the production process in a way that allows more tasks to be performed by machines.

Tasks, Technologies and Displacement

We propose, instead, a task-based approach, where the central unit of production is a task as in the textile example discussed above.³ Some tasks have to be produced by labor, while other tasks can be produced either by labor or by capital. Also, labor and capital have *comparative advantages* in different tasks, meaning that the relative productivity of labor varies across tasks.

Our framework conceptualizes *automation* (or automation at the extensive margin) as an expansion in the set of tasks that can be produced with capital. If capital is sufficiently cheap or sufficiently productive at the margin, then automation will lead to the

³See Autor, Leavy and Murnane (2003) and Acemoglu and Autor (2011). Differently from these papers which develop a task-based approach focusing on inequality implications of technological change, we are concerned here with automation and the process of capital replacing tasks previously performed by labor, and their implications for wages and employment.

substitution of capital for labor in these tasks. This substitution results in a displacement of workers from the tasks that are being automated, creating the aforementioned displacement effect.

The displacement effect could cause a decline in the demand for labor and the equilibrium wage rate. The possibility that technological improvements that increase productivity can actually reduce the wage of *all* workers is an important point to emphasize because it is often downplayed or ignored.

With an elastic labor supply (or quasi-labor supply reflecting some labor market imperfections), a reduction in the demand for labor also leads to lower employment. In contrast to the standard approach based on factor-augmenting technological changes, a task-based approach immediately opens the way to productivity-enhancing technological developments that simultaneously reduce wages and employment.

Countervailing Effects

The presence of the displacement effect does not mean that automation will always reduce labor demand. In fact, throughout history, there are several periods where automation was accompanied by an expansion of labor demand and even higher wages. There is a number of reasons why automation will also create a positive impact on labor demand.

1. *The productivity effect:* By reducing the cost of producing a subset of tasks, automation raises the demand for labor in non-automated tasks (Autor, 2015; Acemoglu and Restrepo, 2016). In particular, automation leads to the substitution of capital for labor because at the margin, capital performs certain tasks more cheaply than labor used to. This reduces the prices of the goods and services whose production processes are being automated, making households effectively richer, and increasing the demand for all goods and services.

The productivity effect could manifest itself in two complementary ways. First, labor demand might expand in the same sectors that are undergoing automation.⁴ A telling example of this process comes from the effects of the introduction of automated teller machines (ATMs) on the employment of bank tellers. Bessen (2016) documents that concurrent with the rapid spread of ATMs, a clear example of automating technology which enabled these new machines to perform tasks that were previously performed more expensively by labor—there was an expansion in the employment of bank tellers. Bessen suggests that this is because ATMs reduced the costs of banking and encouraged banks to open more branches, raising the

⁴This requires that the demand for the products of these sectors is elastic. Acemoglu and Restrepo (2017) refer to this channel as the price-productivity effect, because it works by reducing the relative price of products that are being automated and restructuring production towards these sectors.

demand for bank tellers who then specialized in a range of tasks that ATMs did not automate.

Another interesting example of this process is provided by the dynamics of labor demand in spinning and weaving during the British Industrial Revolution as recounted by Mantoux (1928). Automation in weaving (most notably, John Kay’s fly shuttle) made this task cheaper and increased the price of yarn and the demand for the complementary task of spinning. Later automation in spinning reversed this trend and increased the demand for weavers. In the words of John Wyatt, one of the inventors of the spinning machine, installing spinning machines would cause clothiers to “then want more hands in every other branch of the trade, viz. weavers, shearmen, scourers, combers, etc...” (quoted in Mantoux, 1928). This is also probably the reason why the introduction of Eli Whitney’s cotton gin in 1793, which automated the labor-intensive process of separating the cotton lint from the seeds, appears to have led to greater demand for slave labor in southern plantations (Rasmussen, 1982).

The productivity effect also leads to higher real incomes and thus to greater demand for all products, including those not experiencing (much) automation. The greater demand for labor from other industries might then counteract the negative displacement effect of automation. The clearest historical example of this comes from the adjustment of the US and many European economies to the mechanization of agriculture. By reducing food prices, mechanization enriched consumers who then demanded more non-agricultural goods (Herrendorf, Rogerson and Valentinyi, 2013), and created employment opportunities for many of the workers dislocated by the mechanization process in the first place.⁵

This discussion also implies that, in contrast to the popular emphasis on the negative labor market consequences of “brilliant” and highly productive new technologies set to replace labor (e.g., Brynjolfsson and McAfee, 2012; Ford, 2016), the real danger for labor may come not from highly productive but from “so-so” automation technologies that are just productive enough to be adopted and cause displacement, but not sufficiently productive to bring about powerful productivity effects.

2. *Capital accumulation:* As our framework in the next section clarifies, automation corresponds to an increase in the capital intensity of production. The high demand for capital triggers further accumulation of capital (e.g., by increasing the rental rate of capital). Capital accumulation then raises the demand for labor. This may

⁵Acemoglu and Restrepo (2017) refer to it as a “scale effect” because in their setting it acted in a homothetic manner, scaling up demand from all sectors, though in general it could take a non-homothetic form.

have been an important channel of adjustment of the British economy during the Industrial Revolution and of the American economy in the first half of the 20th century in the face of mechanization of agriculture, for in both cases there was rapid capital accumulation (Allen, 2009; Olmstead and Rhode, 2001).

As we discuss in the next section, under some (albeit restrictive) assumptions often adopted in neoclassical models of economic growth, capital accumulation can be sufficiently powerful that automation will always increase wages in the long run (see Acemoglu and Restrepo, 2016), though the more robust prediction is that it will act as a countervailing effect.

3. *Deepening of automation:* The displacement effect is created by automation at the extensive margin—meaning the expansion of the set of tasks that can be produced by capital. But what happens if technological improvements increase the productivity of capital in tasks that have already been automated? This will clearly not create additional displacement, because labor was already replaced by capital in those tasks. But it will generate the same productivity effects we have already pointed out. These productivity effects then raise labor demand. We refer to this facet of advances in automation technology as the deepening of automation (or as automation at the intensive margin, because it is intensifying the productive use of machines).

A clear illustration of the role of deepening automation comes from the introduction of new vintages of machinery replacing older vintages used in already automated tasks. For instance, in U.S. agriculture, the replacement of horse-powered reapers and harvesters by diesel tractors increased productivity, presumably with limited additional substitution of workers in agricultural tasks.⁶ In line with our account of the potential role of deepening automation, agricultural productivity and wages increased rapidly starting in the 1930s, a period that coincided with the replacement of horses by tractors (Olmstead and Rhode, 2001; Manuelli and Seshadri, 2014).

Another example comes from the vast improvements in the efficiency of numerically-controlled machines used for metal cutting and processing (such as mills and lathes) as the early vintages controlled by punched cards were replaced by computerized models during the 1970s. The new computerized machines were used in the same tasks as the previous vintages, and so the additional displacement effects were probably minor. At the same time, the transition to CNC (computer numerical control) machines increased the productivity of machinists, operators, and other workers in

⁶Nevertheless, the move from horse power to tractors contributed to a decline in agricultural employment via a different channel: tractors increased agricultural productivity, and because of inelastic demand, expenditure on agricultural products declined (Rasmussen, 1982).

the industry (Groover, 1983).

The three countervailing forces we have listed here are central for understanding why the implications of automation are much richer than the direct displacement effects might at first suggest, and why automation need not be an unadulterated negative force against the labor market fortunes of workers. Nevertheless, there is one aspect of the displacement effect that is unlikely to be undone by any of these four countervailing forces: as we show in the next section, automation necessarily makes the production process more capital intensive and tends to increase productivity more than the wage, as a consequence reducing the share of labor in national income. Intuitively, this is because it entails the substitution of capital for tasks previously performed by labor, thus squeezing labor into a narrower set of tasks.

If, as we have suggested, automation has been ongoing for centuries, with or without powerful countervailing forces of the form listed here, we should have seen a fairly “non-balanced” growth process with the share of labor in national income declining steadily since the beginning of the Industrial Revolution. That clearly hasn’t been the case (see, e.g., Kuznets, 1966; Acemoglu, 2009). This suggests that there has been another powerful forces making production more labor-intensive to balance the effects of automation. This is what we suggest in the next subsection.

New Tasks

As already discussed in the Introduction, periods of intensive automation have often coincided with the emergence of new jobs, activities, industries and tasks. In 19th-century Britain, for example, there was a rapid expansion of various new industries and jobs ranging from engineers, machinists, repairmen, conductors, back-office workers and managers involved with the introduction and operation of new technologies (e.g., Landes, 1969, Chandler, 1977, and Mokyr, 1990). In early 20th-century America, the mechanization of agriculture coincided with a large increase in employment in new industry and factory jobs (Kuznets, 1966), among others in the burgeoning industries of farm equipment (Olmstead and Rhode, 2001) and cotton milling (Rasmussen, 1982). This is not just a historical phenomenon. As documented in Acemoglu and Restrepo (2016), from 1980 to 2010, the introduction and expansion of new tasks and job titles explains about half of employment growth.

Our task-based framework highlights that the creation of new, labor-intensive tasks (tasks in which labor has a comparative advantage relative to capital) may be the most powerful force balancing the growth process in the face of rapid automation. Without the demand for workers from new factory jobs, engineering, supervisory tasks, accounting and managerial occupations in the second half of the 19th and much of the 20th centuries, it

would have been impossible to employ millions of workers exiting the agricultural sector and traditional labor-intensive tasks.

In the same way that automation has a displacement effect, we can think of the creation of new tasks as engendering a *reinstatement effect*. In this way, the creation of new tasks has the opposite effect of automation. It always generates additional labor demand, which notably increases the share of labor in national income. Consequently, one powerful way in which technological progress could be associated with a balanced growth path is via the balancing of the impacts of automation by the creation of new tasks.

The creation of new tasks need not be an exogenous, autonomous process, entirely unrelated to automation, AI, and robotics. This is at least for two reasons.

1. As emphasized in Acemoglu and Restrepo (2016), rapid automation may endogenously generate incentives for firms to introduce new labor-intensive tasks. Automation running ahead of the creation of new tasks reduces the labor share and possibly wages, making further automation less profitable, and new tasks generating employment opportunities for labor more profitable for firms. Acemoglu and Restrepo (2016) show that this equilibrating force could be powerful enough to make the growth process balanced.
2. Some automation technology platforms, especially AI, may facilitate the creation of new tasks. A recent report by Accenture identified entirely new categories of jobs that are emerging in firms using AI as part of their production process (Accenture PLC, 2017). These jobs include “trainers” (to train the AI systems), “explainers” (to communicate and explain the output of AI systems to customers), and “sustainers” (to monitor the performance of AI systems, including their adherence to prevailing ethical standards).

The applications of AI to education, health care, and design may also result in employment opportunities for new workers. Take education. Existing evidence suggests that many students, not least those with certain learning disabilities, will benefit from individualized education programs and personalized instruction (Kolb, 1984). With current technology, it is prohibitively costly to provide such services to more than a small fraction of students. Applications of AI may enable the educational system to become more customized, and in the process create more jobs for education professionals to monitor, design and implement individualized education programs. Similar prospects exist in health care and elderly care services.

Revisiting the False Dichotomy

The conceptual framework outlined above, which will be further elaborated in the next section, clarifies why the current debate is centered on a false dichotomy between disastrous and totally benign effects of automation.

Our task-based framework underscores that automation will always create a displacement effect. Unless neutralized by the countervailing forces, this displacement effect could reduce labor demand, wages, and also in general, employment. At the very least, this displacement effect implies that a falling share of the productivity gains will accrue to labor. The possibility that the displacement effect keeps the demand for labor from expanding at par with output pushes against the benign accounts emphasizing that technology always increases the demand for labor and benefits workers.

Our framework does not support the alarmist perspectives stressing the disastrous effects of automation for labor either. Rather, it highlights several countervailing forces which soften the impact of automation on labor. More importantly, as we have argued in the previous subsection, the creation of new labor-intensive tasks has been a critical part of the adjustment process in the face of rapid automation. The creation of new tasks is not just *mana* from heaven. There are good reasons why market incentives will endogenously lead to the creation of new tasks that gain strength when automation itself becomes more intensive. Also, some of the most defining automation technologies of our age, such as AI, may create a platform for the creation of new sets of tasks and jobs.

At the root of some of the alarmism is the belief that AI will have very different consequences for labor than previous waves of technological change. Our framework highlights that the past is also replete with automation technologies displacing workers, but this need not have disastrous effects for labor. Nor is it technologically likely that AI will replace labor in all or almost all of the tasks it is currently specializing in. This limited remit of AI can be best understood by contrasting the current nature and ambitions of AI with those of its first coming under the auspices of “cybernetics”. The intellectual luminaries of cybernetics, such as Norbert Wiener, envisaged the production of *Human-Level Artificial Intelligence*—computer systems capable of thinking in a way that could not be distinguished from human intelligence—replicating all human thought processes and faculties (Nilsson, 2009). In 1965, Herbert Simon predicted that “machines will be capable, within twenty years, of doing any work a man can do” (Simon, 1965, p. 96). Marvin Minsky agreed, declaring in 1967 that “Within a generation, I am convinced, few compartments of intellect will remain outside the machine’s realm” (Minsky, 1967, p. 2).

Current practice in the field of AI, especially in its most popular and promising forms based on deep learning and various other “big data” methods applied to unstructured data, eschews these initial ambitions and aims at developing *applied Artificial Intelligence*—

commercial systems specializing in clearly delineated tasks related to prediction, decision making, logistics, and pattern recognition (Nilsson, 2009). Though many occupations involve such tasks—and so AI is likely to have a displacement effect in these tasks—there are still many human skills that we still cannot automate, including complex reasoning, judgment, analogy-based learning, abstract problem-solving, and a mixture of physical activity, empathy and communication. This reading of the current practice of AI suggests that the potential for AI and related technological advances to automate a vast set of tasks is limited, even without taking into account the creation of new tasks that will be brought into action because of the economic forces we have already emphasized.

Flies in the Ointment

Our framework so far has emphasized two key ideas. First, automation does create a potential negative impact on labor through the displacement effect and also by reducing the share of labor in national income. But second, it can be counterbalanced by the creation of new tasks (as well as the productivity effect, capital accumulation and the deepening of automation, which tend to increase the demand for labor, even though they do not generally restore the share of labor in national income to its pre-automation levels).

The picture we have painted does underplay some of the challenges of adjustment, however. The economic adjustment following rapid automation can be more painful than the process we have outlined for a number of reasons.

Most straightforwardly, automation changes the nature of existing jobs, and the reallocation of workers from existing jobs and tasks to new ones is a complex and often slow process. It takes time for workers to find new jobs and tasks in which they can be productive, and periods during which workers are laid off from their existing jobs can create a depressed local or national labor market, further increasing the costs of adjustment. These effects are visible in recent studies that have focused on the adjustment of local US labor markets to negative demand shocks, such as Autor, Dorn and Hanson (2015), who study the slow and highly incomplete adjustment of local labor markets in response to the surge in Chinese exports, Mian and Sufi (2015), who investigate the implications of the collapse in housing prices on consumption and local employment, and perhaps more closely related to our focus, Acemoglu and Restrepo (2017), who find employment and wage declines in areas most exposed to one specific type of automation, the introduction of industrial robots in manufacturing.

The historical record also underscores the painful nature of the adjustment. The rapid introduction of new technologies during the British Industrial Revolution ultimately led to rising labor demand and wages, but this was only after a protracted period of stagnant wages, expanding poverty, and harsh living conditions. During an eighty year period

extending from the beginning of the Industrial Revolution to the middle of the 19th century, wages stagnated and the labor share fell, even as technological advances and productivity growth were ongoing in the British economy, a phenomenon which Allen (2009) dubs the “Engel’s pause” (previously referred to as the “living standards paradox”, see Mokyr, 1990).

There should thus be no presumption that adjustment to the changed labor market brought about by rapid automation will be a seamless, costless and rapid process.

Mismatch between Skills and Technologies

It is perhaps telling that wages started growing in the 19th-century British economy only after mass schooling and other investments in human capital expanded the skills of the workforce. Similarly, the adjustment to the large supply of labor freed from agriculture in the early 20th-century America may have been greatly aided by the “high school movement” which increased the human capital of the new generation of American workers (Goldin and Katz, 2010). The forces at work here are likely to be more general than these examples. New tasks tend to require new skills. But to the extent that the workforce does not possess those skills, the adjustment process will be hampered. Even more ominously, if the educational system is not up to providing those skills (and if we are not even aware of the types of new skills that will be required so as to enable investments in them), the adjustment will be greatly impeded. Even the most optimistic observers ought to be concerned about the ability of the current US educational system to identify and provide such skills.

At stake here is not only the speed of adjustment, but potential gains from new technologies. If certain skills are complementary to new technologies, their absence will imply that the productivity of these new technologies will be lower than otherwise. Thus the mismatch between skills and technologies not only slows down the adjustment of employment and wages, but holds back potential productivity gains. This is particularly true for the creation of new tasks. The fact that while there is heightened concerns about job losses from automation, many employers are unable to find workers with the right skills for their jobs underscores the importance of these considerations (Deloitte and the Manufacturing Institute, 2011).

Missing Productivity and Excessive Automation

The issues raised in the previous subsection are important not least because a deep puzzle in any discussion of the impact of new technologies is missing productivity growth—the fact that while so many sophisticated technologies are being adopted, productivity growth

has been slow. As pointed out by Gordon (2016), US productivity growth since 1974 (with the exception of the period from 1995-2004) compares dismally to its post-war performance. While the annual rate of labor productivity growth of the US economy averaged 2.7 percent between 1947 and 1973, it only averaged 1.5 percent between 1974 and 1994. Average productivity growth rebounded to 2.8 percent between 1995 and 2004, and then fell again to only 1.3 percent between 2005 and 2015 (Syverson, 2017). How can we make sense of this?

One line of attack argues that there is plenty of productivity growth, but it is being mismeasured. But, as pointed out by Syverson (2017), the pervasive nature of this slow down, and the fact that it is even more severe in industries that have made greater investments in information technology (Acemoglu et al., 2014) make the productivity mismeasurement hypothesis unlikely to account for all of the slowdown.

Our conceptual framework suggests some possible explanations. They center around the possibility of “excessive automation,” meaning faster automation than socially desirable (Acemoglu and Restrepo, 2016, 2018b). Excessive automation not only creates direct inefficiencies, but may also hold productivity growth down by wastefully using resources and displacing labor.

There are two broad reasons for excessive automation, both of which we believe to be important. The first is related to the biases in the US tax code, which subsidizes capital relative to labor. This subsidy takes the form of several different provisions, including additional taxes and costs employers have to pay for labor, subsidies in the form of tax credits and accelerated depreciation for capital outlays, and additional tax credit for interest rate deductions in case of debt-financed investments (AEI, 2008; Tuzel and Zhang, 2017). All of these distortions imply that, at the margin, when a utilitarian social planner would be indifferent between capital and labor, the market would have an incentive to use machines, giving an additional boost to automation. This inefficiency could translate into slow productivity growth because the substitution of labor for machines worsens the misallocation of capital and labor.

Even absent such a fiscal bias, there are natural reasons for excessive automation. Labor market imperfections and frictions also tend to imply that the equilibrium wage is above the social opportunity cost of labor. Thus a social planner would use a lower shadow wage in deciding whether to automate a task than the market, creating another force towards excessive automation. The implications of this type of excessive automation would again include slower productivity growth than otherwise.

Finally, it is possible that automation has continued at its historical pace, or may have even accelerated recently, but the dismal productivity growth performance we are witnessing is driven by a slowdown in the creation of new tasks or investment in other productivity-enhancing technologies (see Acemoglu and Restrepo, 2016). A deceleration

in the creation of new tasks and technologies other than automation would also explain why the period of slow productivity growth coincided with poor labor market outcomes, including stagnant median wages and a decline in the labor share.

There are natural reasons why too much emphasis on automation may come at the cost of investments in other technologies, including the creation of new tasks. For instance, in a setting where technologies are developed endogenously using a common set of resources (e.g., scientists), there is a natural trade-off between faster automation and investments in other types of technologies (Acemoglu and Restrepo, 2016). Though it is at the moment impossible to know whether the redirection of research resources away from the creation of new tasks and towards automation has played an important role in the productivity slowdown, the almost singular focus in the corporate sector and research community on AI, applications of deep learning, and other big data methods to automate various tasks makes it at least plausible that there may be too much attention devoted to automation at the expense of other technological breakthroughs.

3 A MODEL OF AUTOMATION, TASKS, AND THE DEMAND FOR LABOR

In the previous section, we provided an intuitive discussion of how automation in general, and robotics and AI in particular, is expected to impact productivity and the demand for labor. In this section, we outline a formal framework which underlines these conclusions. Our presentation will be somewhat informal and without any derivations, which are all collected in the Appendix.

A Task-Based Framework

We start with a simplified version of the task-based framework introduced in Acemoglu and Restrepo (2016). Aggregate output is produced by combining the services of a unit measure of tasks $x \in [N - 1, N]$ according to the following Cobb-Douglas (unit elastic) aggregator

$$\ln Y = \int_{N-1}^N \ln y(x) dx, \quad (1)$$

where Y denotes aggregate output and $y(x)$ is the output of task x . The fact that tasks run between $N - 1$ and N enables us to consider changes in the range of tasks, for example because of the introduction of new tasks, without altering the total measure of tasks in the economy.

Each task can be produced by human labor, $\ell(x)$, or by machines, $m(x)$, depending

on whether it has been (technologically) automated or not. In particular, tasks $x \in [0, I]$ are technologically automated, so can be produced by either labor or machines, while the rest are not technologically automated, so must be produced with labor:

$$y(x) = \begin{cases} \gamma_L(x)\ell(x) + \gamma_M(x)m(x) & \text{if } x \in [0, I] \\ \gamma_L(x)\ell(x) & \text{if } x \in (I, N]. \end{cases} \quad (2)$$

Here, $\gamma_L(x)$ is the productivity of labor in task x and is assumed to be increasing, while $\gamma_M(x)$ is the productivity of machines in automated tasks. We assume that $\gamma_L(x)/\gamma_M(x)$ is increasing in x , and thus labor has a *comparative advantage* in higher-indexed tasks.⁷

The threshold I denotes the frontier of automation possibilities: it describes the range of tasks that can be automated using current available technologies in AI, industrial robots, various computer-assisted technologies, and other forms of “smart machines”.

We also simplify the discussion by assuming that both the supply of labor, L , and the supply of machines, K , are fixed and inelastic. The fact that the supply of labor is inelastic implies that changes in labor demand impact the share of labor in national income and the wage, but not the level of employment. We outline below how this framework can be easily generalized to accommodate changes in employment and unemployment.

Types of Technological Change

Our framework incorporates four different types of technological advances. All advances increase productivity but as we will see with a very different impact on the demand for labor and wages.

1. *Labor-augmenting technological advances:* Standard approaches in macroeconomics and labor economics typically focus on labor-augmenting technological advances. Such technological changes correspond to increases (or perhaps an equi-proportionate increase) in the function $\gamma_L(x)$. Our analysis will show that they are in fact quite special, and the implications of automation and AI are generally very different from those of labor-augmenting advances.

⁷Our theoretical framework builds on Zeira (1998) who develops a model where firms produce intermediates using labor-intensive or capital-intensive technologies. Zeira focuses on how wages affect the adoption of capital-intensive production methods and how this margin amplifies productivity differences across countries and over time. In contrast, we focus on the implications of automation—modeled here as an increase in the set of tasks that can be produced by machines, represented by I —for the demand for labor, wages, and employment, and we also study the implications of the introduction of new tasks. In Acemoglu and Restrepo (2016), we generalize Zeira’s framework in a number of other dimensions and also endogenize the development of automation technologies and new tasks.

2. *Automation (at the extensive margin)*: We consider automation to be an expansion of the set of tasks that are technologically automated as represented by the parameter I .
3. *Deepening of automation (or automation at the intensive margin)*: Another dimension of advances in AI and robotics technology will tend to increase the productivity of machines in tasks that are already automated, for example, by replacing existing machines with newer, more productive vintages. In terms of our model, this corresponds to an increase in the $\gamma_M(x)$ function for tasks $x < I$. We will see that this type of deepening of automation has very different implications for labor demand than automation (at the extensive margin).
4. *Creation of new tasks*: As emphasized in Acemoglu and Restrepo (2016), another important aspect of technological change is the creation of new tasks and activities in which labor has a comparative advantage. In our model this can be captured in the simplest possible way by an increase in N .

Equilibrium

Throughout, we denote the equilibrium wage rate by W and the equilibrium cost of machines (or the rental rate) by R . An equilibrium requires firms to choose the cost-minimizing way of producing each task, and labor and capital markets to clear.

To simplify the discussion, we impose the following assumption

$$\frac{\gamma_L(N)}{\gamma_M(N-1)} > \frac{W}{R} > \frac{\gamma_L(I)}{\gamma_M(I)}. \quad (\text{A1})$$

The second inequality implies that all tasks in $[N-1, I]$ will be produced by machines. The first inequality implies that the introduction of new tasks—an increase in N —will increase aggregate output. This assumption is imposed on the wage to rental rate ratio, which is an endogenous object; the Appendix provides a condition on the stock of capital and labor which is equivalent to this assumption (see Assumption A2).

We also show in the Appendix that aggregate output (GDP) in the equilibrium takes the form

$$Y = B \left(\frac{K}{I - N + 1} \right)^{I - N + 1} \left(\frac{L}{N - I} \right)^{N - I}, \quad (3)$$

where

$$B = \exp \left(\int_{N-1}^I \ln \gamma_M(x) dx + \int_I^N \ln \gamma_L(x) dx \right). \quad (4)$$

Aggregate output is given by a Cobb-Douglas aggregate of the capital stock and employment. This resulting aggregate production function in (3) is itself derived from the

allocation of the two factors of production to tasks. More importantly, the exponents of capital and labor in this production function depend on the extent of automation, I , and the creation of new tasks, as captured by N .

Central to our focus is not only the impact of new technologies on productivity but also on the demand for labor. The Appendix shows that the demand for labor can be expressed as

$$W = (N - I) \frac{Y}{L}. \quad (5)$$

This equation is intuitive in view of the Cobb-Douglas production function in (3), since it shows that the wage (the marginal product of labor) is equal to the average product of labor—which we will also refer to as “productivity”—times the exponent of labor in the aggregate production function. Equation (5) can be inverted to obtain a downward-sloping labor demand curve as a function of the wage.

Equation (5) implies that the share of labor in national income is given by

$$s_L = \frac{WL}{Y} = N - I. \quad (6)$$

4 TECHNOLOGY AND LABOR DEMAND

The Displacement Effect

Our first result shows that automation (at the extensive margin) indeed creates a *displacement effect*, reducing labor demand as emphasized in Section 2, but also that it is counteracted by a *productivity effect*, pushing towards greater labor demand.

Specifically, from equation (5) we directly obtain

$$\frac{d \ln W}{dI} = \underbrace{\frac{d \ln(N - I)}{dI}}_{\substack{\text{Displacement} \\ \text{effect} < 0}} + \underbrace{\frac{d \ln(Y/L)}{dI}}_{\substack{\text{Productivity} \\ \text{effect} > 0}}. \quad (7)$$

Without the productivity effect, automation would always reduce labor demand, because it is directly replacing labor in tasks that were previously performed by workers. Indeed, if the productivity effect is limited, automation will reduce labor demand and wages.

Counteracting the Displacement Effect I: The Productivity Effects

The productivity effect, on the other hand, captures the important idea that technologies that increase productivity will tend to raise labor demand and wages. As highlighted in the previous section, there are two complementary manifestations of the productivity effect. The first works by increasing the demand for labor in non-automated tasks in the industries where automation is ongoing. The second works by raising the demand for labor in other industries that are not automating as much. The productivity effect shown in equation (7) combines these two mechanisms.

One important implication of the decomposition in equation (7) is that, in contrast to some popular discussions, the new AI and robotics technologies that are more likely to reduce the demand for labor are not those that are brilliant and highly productive, but those that are “so-so”—just productive enough to be adopted but not much more productive or cost-saving than the production processes that they are replacing. Interestingly, and related to our discussion on missing productivity, if new automation technologies are so-so, they would not bring major improvements in productivity either.

To elaborate further on this point and to understand the productivity implications of automation technologies better, let us also express the productivity effect in terms of the physical productivities of labor and machines and factor prices as follows

$$\frac{d \ln(Y/L)}{dI} = \ln\left(\frac{W}{\gamma_L(I)}\right) - \ln\left(\frac{R}{\gamma_M(I)}\right) > 0.$$

The fact that this expression is positive, and that new automation technologies will be adopted, follows from Assumption A1. Using this expression, the overall impact on labor demand can be alternatively written as

$$\frac{d \ln W}{dI} = - \underbrace{\frac{1}{N-I}}_{\text{Displacement effect} < 0} + \underbrace{\ln\left(\frac{W}{\gamma_L(I)}\right) - \ln\left(\frac{R}{\gamma_M(I)}\right)}_{\text{Productivity effect} > 0}. \quad (8)$$

This expression clarifies that the displacement effect of automation will dominate the productivity effect and thus reduce labor demand (and wages) when $\gamma_M(I)/R \approx \gamma_L(I)/W$, which is exactly the case where new technologies are so-so—only marginally better than labor at newly-automated tasks. In contrast, when $\gamma_M(I)/R \gg \gamma_L(I)/W$, automation will increase productivity sufficiently to raise the demand for labor and wages.

Turning next to the implications of automation for the labor share, equation (6) immediately implies

$$\frac{ds_L}{dI} = -1 < 0, \quad (9)$$

so that regardless of the magnitude of the productivity effect, automation always reduces the share of labor in national income. This negative impact on the labor share is a direct consequence of the fact that automation always increases productivity more than the wage, $d \ln(Y/L)/dI > d \ln W/dI$ (itself directly following from equation (7), which shows that the impact on wages is given by the impact on productivity minus the displacement effect).

Below, we investigate the implications of deepening automation and the creation of new tasks for labor demand and wages. At this point, we simply note that the implications of standard labor-augmenting technological change, which corresponds to a (marginal) shift-up of the $\gamma_L(x)$ schedule, are very different. This type of technological change leaves the form of the wage equation, (5), unchanged, and increases average output per worker, Y/L , and the equilibrium wage, W , proportionately, and thus does not impact the share of labor in national income.⁸

Counteracting the Displacement Effect II: Capital Accumulation

We have so far emphasized the displacement effect created by new automation technologies. We have also seen that the productivity effect—resulting from the fact that automation technologies reduced costs—counteracts the displacement effects to some degree. In this and the next subsection, we discuss two additional countervailing forces.

The first discussed in this section is capital accumulation. The analysis so far assumed that the economy has a fixed supply of capital that could be devoted to new machines (automation technologies). As a result, a further increase in automation (at the extensive margin) increases the demand for capital and thus the equilibrium rental rate, R . This may be understood as the short-run effect of automation.

Instead, we may envisage the “medium-run” effect as the impact of these technologies after the supply of machines used in newly automated tasks expands as well. Because machines and labor are q -complements, an increase in the capital stock, with the level of employment held constant at L , increases the real wage and reduces the rental rate. Equation (8) shows that this change in factor prices makes the productivity effect more powerful and the impact on the wage more likely to be positive.

⁸A small shift-up of $\gamma_L(x)$ does not violate Assumption A1, because at the margin it was strictly cost-saving to use machines. A larger labor-augmenting technological change may result in a violation of Assumption A1. At this point, only tasks below an endogenous threshold $\tilde{I} < I$ would be automated, and labor augmenting technologies could also reduce $\tilde{I}$, increasing the labor share in national income.

In the limit, if capital accumulation fixes the rental rate at a constant level (which will be the case, for example, when we have a representative household with exponential discounting and time-separable preferences), the productivity effect will always dominate the displacement effect.⁹

Crucially, however, equation (6) still applies, and thus automation continues to reduce the labor share, even after the adjustment of the capital stock.

Counteracting the Displacement Effect III: Deepening of Automation

Another potentially powerful force counteracting the displacement effect from automation at the extensive margin comes from the deepening of automation (or automation at the intensive margin), for example, because of improvements in the performance of already existing automation technologies or the replacement of such technologies with newer, more productive vintages. This increase in the productivity of machines in tasks that are already automated corresponds in our model to an increase in the function $\gamma_M(x)$ in tasks below I .

To explore the implications of this type of change in the simplest possible way, let us suppose that $\gamma_M(x) = \gamma_M$ in all automated tasks, and consider an increase in the productivity of machines by $d \ln \gamma_M > 0$, with no change in the extensive margin of automation, I . The implications of this change in the productivity of machines on equilibrium wages and productivity can be obtained as:

$$d \ln W = d \ln Y / L = (I - N + 1) d \ln \gamma_M > 0.$$

Hence, deepening of automation will tend to increase labor demand and wages, further counteracting the displacement effect. Note, however, that as with capital accumulation, in our model this has no impact on the share of labor in national income, as can be seen from the fact that wages and productivity increase by exactly the same amount.

⁹Assuming that production exhibits constant returns to scale, the productivity gains from any technology accrue to both capital and labor. In particular, for any constant returns to scale production function, we have $d \ln Y|_{K,L} = s_L d \ln W + (1 - s_L) d \ln R$, where $d \ln Y|_{K,L} > 0$ denotes the productivity gains brought by technology holding the use of capital and labor constant, and s_L is the labor share. If the rental rate is constant in the long run, then $d \ln R = 0$ and all productivity gains accrue to the relatively inelastic factor, labor.

New Tasks and the Comparative Advantage of Labor

Much more powerful than the countervailing effects of capital accumulation and the deepening of automation is the creation of new tasks in which labor has a comparative advantage. These tasks include both new, more complex versions of existing tasks and the creation of new activities, which are made possible by advances in technology. In terms of our framework, they correspond to increases in N .

An increase in N —the creation of new tasks—raises productivity by

$$\frac{d \ln Y/L}{dN} = \ln \left(\frac{R}{\gamma_M(N-1)} \right) - \ln \left(\frac{W}{\gamma_L(N)} \right) > 0,$$

which is positive from Assumption A1.

More importantly for our focus here, the creation of new tasks also increases labor demand and equilibrium wages by creating a *reinstatement effect* counteracting the displacement effect. In particular,

$$\frac{d \ln W}{dN} = \underbrace{\ln \left(\frac{R}{\gamma_M(N-1)} \right) - \ln \left(\frac{W}{\gamma_L(N)} \right)}_{\substack{\text{Productivity} \\ \text{effect} > 0}} + \underbrace{\frac{1}{N-I}}_{\substack{\text{Reinstatement} \\ \text{effect} > 0}}. \quad (10)$$

In contrast to capital accumulation and the deepening of automation, which increase the demand for labor but do not affect the labor share, equation (6) implies that new tasks increase the labor share, i.e.,

$$\frac{ds_L}{dN} = 1.$$

The centrality of new tasks can be understood when viewed from a complementary historical angle. Automation is not a recent phenomenon. As we already discussed in Section 2, the history of technology of the last two centuries is full of examples of automation, ranging from weaving and spinning machines to the mechanization of agriculture as discussed in the previous section. Even with capital accumulation and the deepening of automation, if there were no other counteracting force, we would see the share of labor in national income declining steadily. Our conceptual framework highlights a major force preventing such a decline—the creation of new tasks in which labor has a comparative advantage.

This can be seen by putting together equations (7) and (10), which yields

$$d \ln W = \left[\ln \left(\frac{R}{\gamma_M(N-1)} \right) - \ln \left(\frac{W}{\gamma_L(N)} \right) \right] dN + \left[\ln \left(\frac{W}{\gamma_L(I)} \right) - \ln \left(\frac{R}{\gamma_M(I)} \right) \right] dI \quad (11)$$

$$+ \frac{1}{N-I} (dN - dI),$$

and also from (6),

$$ds_L = dN - dI.$$

Consequently, for the labor share to remain stable and for wages to increase in tandem with productivity, we need I —capturing the extensive margin of automation—to grow by the same amount as the range of new tasks, N . When that happens, equilibrium wages grow proportionately with productivity, and the labor share, s_L , remains constant, as can be seen from the fact that the first line of (11) is in this case equal to the increase in productivity or GDP per worker. Indeed, rewriting (11) imposing $dN = dI$, we have

$$d \ln W = \left[\ln \left(\frac{\gamma_L(N)}{\gamma_M(N-1)} \right) - \ln \left(\frac{\gamma_L(I)}{\gamma_M(I)} \right) \right] dI > 0,$$

which is strictly positive because of Assumption A1.

A False Dichotomy: Recap

With our conceptual framework exposited in a more systematic manner, we can now briefly revisit the false dichotomy highlighted in the Introduction. Our analysis, in particular equation (7), highlights that there is always a negative displacement effect on labor resulting from automation. Equation (11) reiterates that there is no presumption that this displacement effect could not reduce overall demand for labor.

However, several countervailing effects imply that a negative impact from automation on labor demand is not a forgone conclusion. Most importantly, the productivity effect could outweigh the displacement effect, leading to an expansion in labor demand and equilibrium wages from automation. The presence of the productivity effect as the key counterweight to the displacement created by automation highlights an important conceptual issue, however. In contrast to the emphasis in the popular discussions it is not the brilliant, super-productive automation technologies that threaten labor, but the “so-so” ones which create the displacement effect as they replace labor in tasks that it previously performed, but do not engender the countervailing productivity effect.

The productivity effect is supplemented by the capital accumulation that automation sets in motion and the deepening of automation, which increases the productivity of

machines in tasks that have already been automated. But even with these countervailing effects, equation (9), shows that automation will always reduce the share of labor in national income. All the same, this does not signal the demise of labor either, because the creation of new tasks in which labor has a comparative advantage could counterbalance it, which is our interpretation of why the demand for labor has kept up with productivity growth in the past despite several rapid waves of automation.

Our framework suggests that the biggest shortcoming of the alarmist and the optimist views is their failure to recognize that the future of labor depends on the balance between automation and the creation of new tasks. Automation will often lead to a healthy growth of labor demand and wages if it is accompanied with a commensurate increase in the set of tasks in which labor has a comparative advantage—a feature that alarmists seem to ignore. Even though there are good economic reasons for why the economy will create new tasks, this is neither a foregone conclusion nor something we can always count on without the requisite investments and adjustments—as the optimists seem to assume. AI and robotics could be permanently altering this balance, causing automation to pace ahead of the creation of new tasks with negative consequences for labor, at the very least in regards to the share of labor in national income.

Generalizations

Many of the features adopted in the previous subsection are expositional simplifications. In particular, the aggregate production function (1) can be taken to be any constant elasticity of substitution aggregate. One implication of this would be that aggregate output in equation (3) would be a constant elasticity aggregate itself. This does not affect any of our main conclusions, including the negative impact of automation on the labor share (see Acemoglu and Restrepo, 2016).¹⁰

We also do not need Assumption A1 for any of the results. If the second inequality in this assumption does not hold, changes in automation technology have no impact on the equilibrium because it is not cost-effective to adopt all available automation technologies (for this reason, in the general case, Acemoglu and Restrepo, 2016, distinguish technologically-automated tasks from equilibrium automation). Given our focus here, there is no loss of generality in making this assumption.

A final feature that is worth commenting on is the fact that in the aggregate production function, (1), the limits of integration are $N - 1$ and N , ensuring that the total measure of tasks is one. This is useful for several reasons. First, when the introduction of new tasks expands the total measure of tasks, it becomes more challenging to obtain

¹⁰Recent work by Aghion, Jones and Jones (2017) points out, however, that if the elasticity of substitution between tasks is less than one and there is an exogenous and high saving rate, the labor share might asymptote to a positive value even with continuously ongoing automation.

a balanced growth path (see Acemoglu and Restrepo, 2016). Second, in this case some minor modifications are necessary so that an expansion in the total measure of tasks leads to productivity improvements. In particular, consider the general case where the elasticity of substitution between tasks is not necessarily equal to one. If it is greater than one, an increase in N leads to higher productivity, but not necessarily when it is less than or equal to one. In this latter case, we then need to introduce direct productivity gains from task variety. For example, in the present case where the elasticity of substitution between tasks is equal to one, we could modify (1) to $\ln Y = \frac{1}{N} \sum_0^N \ln[N^{1+\alpha}y(i)]$, where $\alpha \geq 0$ represents these productivity gains from task variety and ensures that the qualitative results explicit here continue to apply.

Employment and Unemployment

An additional generalization concerns the endogenous adjustment of employment in the face of new automation technologies. We have so far taken labor to be supplied inelastically for simplicity. There are two ways in which the level of employment responds to the arrival of new technologies. The first is via a standard labor supply margin. Acemoglu and Restrepo (2016) show that the endogenous adjustment of labor supply, including income effects and the substitution of consumption and leisure, links the level of employment to the share of labor in national income.

The second possibility is through labor market frictions, for example as in Acemoglu and Restrepo (2018b). Under appropriate assumptions, the endogenous level of employment in this case is also a function of the share of labor in national income. Though both models with and without labor market frictions endogenize employment as a function of the labor share, their normative implications are potentially different as we discuss below.

For now, however, the more important implications of such extensions is to link the level of employment (or unemployment) to labor demand. Automation, when it reduces labor demand, will also reduce the level of employment (or increase the level of unemployment). Moreover, because the supply of labor depends on the labor share, in our framework automation results in a reduction in employment (or an increase in unemployment). As such, our analysis so far also sheds light on (and clarifies the conditions for) the claims that new automation technologies will reduce employment. It also highlights, however, that the fact that automation has been ongoing does not condemn the economy to a declining path of employment. If automation is met by equivalent changes in the creation of new tasks, the share of labor in national income can remain stable and ensure a stable level of employment (or unemployment) in the economy.

5 CONSTRAINTS AND INEFFICIENCIES

Even in the presence of the countervailing forces limiting the displacement effect from automation, there are potential inefficiencies and constraints limiting the smooth adjustment of the labor market and hindering the productivity gains from new technologies.

Here we focus on how the mismatch between skills and technologies not only increases inequality, but also hinders the productivity gains from automation and new tasks. We then explore the possibility that, concurrent with rapid automation, we are experiencing a slowdown in the creation of new tasks, which could result in slow productivity growth. Finally, we examine how a range of factors lead to excessive automation, which not only creates inefficiency but also hinders productivity.

Mismatch of Technologies and Skills

The emphasis on the creation of new tasks counter-balancing the potential negative effects of automation on the labor share and the demand for labor ignores an important caveat and constraint—the potential mismatch between the requirements of new technologies (tasks) and the skills of the workforce. To the extent that new tasks require skilled employees or even new skills to be acquired, the adjustment may be much slower—or even severely hampered—than our analysis so far may suggest.

To illustrate these ideas in the simplest possible fashion, we follow Acemoglu and Restrepo (2016) and assume that there are two types of workers, low-skill with supply L and high-skill with supply H , both of them supplied elastically. We also assume that low-skill workers can only perform tasks below a threshold $S \in (I, N)$, while high-skill workers can perform all tasks. For simplicity, we assume that the productivity of both low-skill and high-skill workers in the tasks that they can perform is still given by $\gamma_L(x)$.¹¹ As a result, while low-skill workers will earn a wage W_L , high-skill workers will earn a wage $W_H \geq W_L$.

In this simple extension of the framework presented so far, the threshold S can be considered as an inverse measure of the mismatch between new technologies and skills. A greater value of S implies that there are plenty of additional tasks for low-skill workers, while a low value of S implies the presence of only a few tasks left that low-skill workers can perform.

¹¹We can also introduce differential comparative advantages and also an absolute productivity advantage for high-skill workers, though we choose not to do so to increase transparency (see Acemoglu and Restrepo, 2016). The more restrictive assumption here is that automation happens at the bottom of the range of tasks. In general, automation could take place in the middle range, and its impact would depend on whether automated tasks are competing predominantly against low-skill or high-skill workers (see Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018a).

Assuming that in equilibrium $W_H > W_L$,¹² which implies that low-skill workers will perform all tasks in the range (I, S) , equilibrium wages satisfy

$$W_H = \frac{Y}{H}(N - S) \quad \text{and} \quad W_L = \frac{Y}{L}(S - I).$$

Thus, the impact of automation on inequality—defined here as the wage premium between high and low-skill workers—is given by

$$\frac{d \ln W_H / W_L}{dI} = \frac{1}{S - I} > 0.$$

This equation shows that automation increases inequality. This is not surprising, since the tasks that are automated are precisely those performed by low-skill workers. But in addition, it also demonstrates that the impact of automation on inequality becomes Worse when there is a severe skill mismatch—the threshold S is close to I . In this case, displaced workers will be squeezed into a very small range of tasks, and hence, each of these tasks will receive a large number of workers and will experience a substantial drop in its price, which translates into a sharp decline in the wage of low-skill workers. In contrast, when S is large, displaced workers can spread across a larger set of tasks without depressing their wage as much.

A severe mismatch in the skills required in automated tasks and other jobs also affects the productivity gains from automation. In particular, we have

$$\frac{d \ln(Y/L)}{dI} = \ln \left(\frac{W_L}{\gamma_L(I)} \right) - \ln \left(\frac{R}{\gamma_M(I)} \right) > 0.$$

This equation shows that the productivity gains from automation depend positively on W_L/R : it is precisely when displaced workers have a high opportunity cost that automation raises productivity. Using the fact that $R = \frac{Y}{K}(I - N + 1)$, we obtain

$$\frac{W_L}{R} = \frac{S - I}{I - N + 1} \frac{K}{L}.$$

A worse mismatch (a lower S) reduces the opportunity cost of displaced workers further, and via this channel, it makes automation less profitable. This is because a severe mismatch impedes reallocation, reducing the productivity gains of freeing workers from automated tasks.

Equally important are the implications of a skill mismatch for the productivity gains

¹² This is equivalent to $\frac{N-S}{S-I} > \frac{H}{L}$, so that high-skill workers are scarce relative to the range of tasks that only they can produce.

from new tasks. Namely,

$$\frac{d\ln(Y/L)}{dN} = \ln\left(\frac{R}{\gamma_M(N-I)}\right) - \ln\left(\frac{W_H}{\gamma_H(N)}\right) > 0,$$

which depends negatively on W_H/R : it is precisely when high-skill workers have a relatively high wage that the gains from new tasks will be limited. With similar arguments to before, we also have

$$\frac{W_H}{R} = \frac{N-S}{I-N+1} \frac{K}{L},$$

which implies that in the presence of a worse mismatch (a lower S), the productivity gains from new tasks will be limited. This is because new tasks require high-skill workers who are scarce and expensive when S is low.

An important implication of this analysis is that to limit increasing inequality and to best deploy new tasks and harness the benefits of automation, society may need to simultaneously increase the supply of skills. A balanced growth process requires not only automation and the creation of new tasks to go hand-in-hand, but also the supply of high-skill workers to grow in tandem with these technological trends.

Automation at the Expense of New Tasks

As discussed in Section 2, a puzzling aspect of recent macroeconomic developments has been the lack of robust productivity growth despite the bewildering array of new technologies. While some have argued that this is because of mismeasurement of true productivity, our conceptual framework provides three novel (and at least to us, more compelling) reasons for slow productivity growth. The first was discussed in the previous subsection.

The second one, discussed in this subsection, is that concurrent with the rapid introduction of new automation technologies, we may be experiencing a slow down in the creation of new tasks and investments in other technologies that benefit labor.

This explanation comes in two flavors. First, we may be running out of good ideas to create new jobs, sectors, and products capable of expanding the demand for labor (e.g., Gordon, 2016; Bloom et al. 2017), even if automation continues at a healthy or accelerating pace. Alternatively, the rapid introduction of new automation technologies may redirect resources that were devoted to other technological advances, in particular, the creation of new tasks (see Acemoglu and Restrepo, 2016). To the extent that the recent enthusiasm—or even “frenzy” —about deep learning and some aspects of AI can be viewed as such a redirection, our framework pinpoints a potential powerful mechanism for slower productivity growth in the face of rapid automation.

Both explanations hinge on the redirection of research activity from the creation of

new tasks to automation—in the first case exogenously and in the second for endogenous reasons. Recall from our analysis so far that the productivity gains from new tasks in our baseline framework are given by

$$\frac{d \ln(Y/L)}{dN} = \ln \left(\frac{R}{\gamma_M(N-1)} \right) - \ln \left(\frac{W}{\gamma_L(N)} \right) > 0,$$

while productivity gains from automation are

$$\frac{d \ln(Y/L)}{dI} = \ln \left(\frac{W}{\gamma_L(I)} \right) - \ln \left(\frac{R}{\gamma_M(I)} \right) > 0.$$

If the former expression is greater than the latter, then the redirection of research effort from the creation of new tasks towards automation, or a lower research efficiency in creating new tasks, will lead to a slowdown of productivity growth, even if advances in automation are accelerating and being adopted enthusiastically. This conclusion is strengthened if additional effort devoted to automation at the expense of the creation of new tasks is runs into diminishing returns.

Excessive Automation

In this subsection, we highlight the third reason for why there may be missing productivity growth: socially excessive automation (see Acemoglu and Restrepo, 2016, 2018b).

To illustrate why our framework can generate excessive automation, we modify the assumption that the supply of capital, K , is given, and instead suppose that machines used in automation are produced—as intermediate goods—using the final good at a fixed cost R . Moreover, suppose that because of subsidies to capital, accelerated depreciation allowances, tax credit for debt-financed investment or simply because of the tax cost of employing workers, capital receives a marginal subsidy of $\tau > 0$.

Given this subsidy, the rental rate for machines is $R(1 - \tau)$, and Assumption A1 now becomes

$$\frac{\gamma_L(N)}{\gamma_M(N-1)} > \frac{W}{R(1-\tau)} > \frac{\gamma_L(I)}{\gamma_M(I)}.$$

Let us now compute GDP as value added, subtracting the cost of producing machines. This gives us

$$GDP = Y - RK.$$

Suppose next that there is an increase in automation. Then we have

$$\frac{dGDP}{dI} = \left. \frac{dY}{dI} \right|_K + R(1-\tau) \frac{dK}{dI} - R \frac{dK}{dI},$$

which simplifies to

$$\frac{dGDP}{dI} = \underbrace{\ln\left(\frac{W}{\gamma_L(I)}\right) - \ln\left(\frac{R(1-\tau)}{\gamma_M(I)}\right)}_{\substack{\text{Productivity} \\ \text{effect} > 0}} - \underbrace{R\tau \frac{dK}{dI}}_{\substack{\text{Excessive} \\ \text{automation} < 0}}.$$

The first term is positive and captures the productivity increase generated by automation. However, when $\tau > 0$ —so that the real cost of using capital is distorted—we have an additional negative effect originating from excessive automation.¹³ At the root of this negative effect is the fact that subsidies induce firms to substitute capital for labor even when this is not socially cost-saving (though it is privately beneficial because of the subsidy).

This conclusion is further strengthened when there are also labor market frictions as pointed out in Section 2. To illustrate this point in the simplest possible fashion, let us assume that there is a threshold $J \in (I, N)$ such that, when performing the tasks in $[I, J]$, workers earn rents $\omega > 0$ proportional to their wage in other tasks. In particular, workers are paid a wage W to produce tasks in $[J, N]$, and a wage $W(1 + \omega)$ to produce tasks in (I, J) .¹⁴ Let L_A denote the total amount of labor allocated to the tasks in (I, J) , and note that these are the workers that will be displaced by automation, i.e., by a small increase in I . Given this additional distortion, Assumption A1 now becomes

$$\frac{\gamma_L(N)}{\gamma_M(N-1)} > \frac{W}{R(1-\tau)} > \frac{1}{1+\omega} \frac{\gamma_L(I)}{\gamma_M(I)}.$$

The demand for labor in tasks where workers earn rents is now

$$L_A = \frac{Y}{W(1+\omega)}(J-I).$$

The demand for labor in tasks where workers do not earn rents is

$$L - L_A = \frac{Y}{W}(N-J).$$

¹³We show in the Appendix that $K = \frac{Y}{R}(I - N + 1)$, which implies that K increases in I .

¹⁴The assumption that there are rents only in a subset of tasks is adopted for simplicity. The same results apply: (i) when there are two sectors and one of the sectors has higher rents/wages for workers and enables automation; (ii) there is an endogenous margin between employment and nonemployment and labor market imperfections (such as search, bargaining or efficiency wages) create a wedge between wages and outside options. In both cases the automation decisions of firms fail to internalize the gap between the market wage and the opportunity cost of labor, leading to excessive automation (see Acemoglu and Restrepo, 2018b).

Dividing these two expressions, we obtain the equilibrium condition for L_A ,

$$\frac{L_A}{L - L_A} = \frac{1}{1 + \omega} \frac{J - I}{N - J},$$

which implies that the total number of workers earning rents declines with automation.

Moreover, the Appendix shows that (gross) output is now given by

$$Y = B \left(\frac{K}{I - N + 1} \right)^{I - N + 1} \left(\frac{L_A}{J - I} \right)^{J - I} \left(\frac{L - L_A}{N - J} \right)^{N - J}, \quad (12)$$

and GDP is still given by $Y - RK$. Equation (12) highlights that there is now a misallocation of labor across tasks—output can be increased by allocating more workers to tasks (I, J) where their marginal product is greater (because of the rents they are earning).

Equation (12) further implies that the impact of automation on GDP is given by

$$\frac{dGDP}{dI} = \underbrace{\ln \left(\frac{W(1 + \omega)}{\gamma_L(I)} \right) - \ln \left(\frac{R(1 - \tau)}{\gamma_M(I)} \right)}_{\text{Productivity effect} > 0} - \underbrace{R\tau \frac{dK}{dI}}_{\text{Excessive automation} < 0} + \underbrace{W\omega \frac{dL_A}{dI}}_{\text{Excessive displacement of labor} < 0}.$$

The new term $W\omega \frac{dL_A}{dI}$ captures the first-order losses from a decline in employment in tasks (I, J) . These losses arise because by automating jobs where workers earn rents, firms are effectively displacing workers to other tasks in which they have a lower marginal product and earn a strictly lower wage, which increases the extent of misallocation.

The point highlighted here is much more general. Without labor market frictions, automation increases GDP (and net output), so at the very least it is possible to redistribute the gains that it creates to make workers—of different skill levels—better off. Labor market frictions change this picture dramatically. In the presence of such frictions, firms' automation decisions do not internalize the fact that the marginal product of labor is above its opportunity cost, or equivalently, do not recognize that there are first-order losses that workers will suffer as a result of automation. Consequently, equilibrium automation could reduce GDP and welfare, and there may not be a way to make (all) workers better off even with tools for costless redistribution. Under these circumstances, a utilitarian planner would choose a lower level of automation than the equilibrium.¹⁵

¹⁵Naturally, if the planner could remove the rents, or the labor market frictions underpinning them, then the equilibrium would be restored to efficiency. Nevertheless, most sources of rents, including search, bargaining and efficiency wages, would be present in the constrained efficient allocations as well.

6 CONCLUDING REMARKS

Despite the growing concerns and intensifying debate about the implications of automation for the future of work, the economics profession and popular discussions lack a satisfactory conceptual framework. To us this lack of appropriate conceptual approach is also the key reason why much of the debate is characterized by a false dichotomy—between the view that automation will spell the end of work for humans and the argument that technologies will always tend to increase the demand for labor as they have done in the past.

In this paper, we summarized a conceptual framework that can help understand the implications of automation and bridge the opposite sides of this false dichotomy. At the center of our framework is a task-based approach, where automation is conceptualized as replacing labor in tasks that it used to perform. This type of replacement causes a direct displacement effect, reducing labor demand. If this displacement effect is not counterbalanced by other economic forces, it will reduce labor demand, wages and employment. But our framework also emphasizes that there are several countervailing forces. These include the fact that automation will reduce the costs of production and thus create a productivity effect, the induced capital accumulation, and the deepening of automation—technological advances that increase the productivity of machines in tasks that have already been automated.

Our framework also emphasizes that these first-order countervailing forces are generally insufficient to totally balance out the implications of automation. In particular, even if these forces are strong, the displacement effect of automation tends to cause a decline in the share of labor in national income. But we know from the history of technology and industrial development that despite several waves of rapid automation, the growth process has been more or less balanced, with no secular downward trend in the share of labor in national income. We argue this is because another powerful force is balancing the implications of automation: the creation of new tasks in which labor has a comparative advantage, which fosters a countervailing reinstatement effect for labor. These tasks increase the demand for labor and tend to raise the labor share. When they go hand-in-hand with automation, the growth process is balanced, and it need not imply a dismal scenario for labor.

Nevertheless, the adjustment process is likely to be slower and more painful than this account of balance between automation and new tasks at first suggests. This is because the reallocation of labor from its existing jobs and tasks to new ones is a slow process, in part owing to time-consuming search and other labor market imperfections. But even more ominously, new tasks require new skills, and especially when the education sector does not keep up with the demand for new skills, a mismatch between skills and technologies is

bound to complicate the adjustment process. We have also argued why such a mismatch will hinder the productivity gains from new technologies.

Our framework further suggests that there are additional reasons for the productivity slowdown. At the center of these is a tendency for excessive automation because of the tax treatment of capital investments and labor market imperfections. Excessive automation directly reduces productivity, but may have even more powerful indirect effects, because it redirects technological improvements away from productivity-enhancing activities that lead to the creation of new tasks and deepening automation to excessive efforts at the extensive margin of automation, a picture that receives informal support from the current singular focus on AI and deep learning.

We would like to conclude by pointing out a number of additional issues that may be important in understanding the full impact of AI and other automation technologies on future prospects of labor. We believe that these issues can be studied using simple extensions of the framework presented here.

First, we have emphasized the role of the productivity effect in partially counterbalancing the displacement effect created by automation. However, this countervailing effect works by increasing the demand for products. As we have also seen, automation tends to increase inequality. If, as a consequence of this distributional impact, the rise in real incomes resulting from automation ends up in the hands of a narrow segment of the population with much lower marginal propensity to consume than those losing incomes and their jobs, these countervailing forces would be weakened and might operate much more slowly. This imbalance in the distribution of the gains from automation might slow down the creation of new tasks as well.

Second, our analysis highlighted the negative consequences of a shortage of skills for realizing the productivity gains from automation and for inequality. In practice, the problem may be workers acquiring the wrong types of skills rather than a general lack of skills. For example, if AI and other new automation technologies necessitate a mix of numeracy, communication, and problem-solving skills different than those emphasized in current curricula, this would have implications similar to those of a shortage of skills, but it cannot be overcome by just increasing educational spending with current educational practices remaining intact. One important consideration in this respect is that there is little concrete information about what types of skills new technologies will complement, underscoring the importance of further empirical work in this area.

Third, government policies and labor market institutions may impact not just the speed of automation (and thus whether there is excessive automation), but what types of automation technologies will receive more investments. To the extent that some uses of AI may complement labor more or generate opportunities for more rapid creation of new tasks, an understanding of the impact of various policies, including support for academic

and applied research, and social factors on the path of development of AI is critical.

Last but not least, the development and adoption of productivity-enhancing AI technologies cannot be taken for granted. If we do not find a way of creating shared prosperity from the productivity gains generated by AI, there is a danger that the political reaction to these new technologies may slow down or even completely stop their adoption and development. This underscores the importance of studying the distributional implications of AI, the political economy reactions to it, and the design of new and improved institutions for creating more broadly shared gains from these new technologies.

APPENDIX

Derivations for the Basic Model

Suppose that Assumption A1 holds. We first derive the demand for factors:

- Denote by $p(x)$ the price of task x . Assumption A1 implies

$$p(x) = \begin{cases} \frac{R}{\gamma_M(x)} & \text{if } x \in [0, I] \\ \frac{W}{\gamma_L(x)} & \text{if } x \in (I, N]. \end{cases} \quad (13)$$

- In addition, the demand for task x is given by

$$y(x) = \frac{Y}{p(x)}.$$

- Thus, the demand for smart machines in task x is

$$k(x) = \begin{cases} \frac{Y}{R} & \text{if } x \in [0, I] \\ 0 & \text{if } x \in (I, N] \end{cases},$$

and the demand for labor in task x is

$$\ell(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ \frac{Y}{W} & \text{if } x \in (I, N] \end{cases}.$$

- Aggregating the demand for machines from this expression and setting it equal to the supply of capital, K , we have the following market-clearing condition for capital

$$K = \frac{Y}{R}(I - N + 1).$$

Similarly, aggregating the demand for labor and setting it equal to its inelastic supply, L , we obtain the market-clearing condition for labor as

$$L = \frac{Y}{W}(N - I).$$

- Rearranging these two equations, the equilibrium rental rate and wage can be obtained as

$$R = \frac{Y}{K}(I - N + 1) \quad \text{and} \quad W = \frac{Y}{L}(N - I), \quad (14)$$

which are the expressions used in the text.

We next turn to deriving the expression for aggregate output.

- Because we normalized the price of the final good to 1 as numeraire, we have

$$\int_{N-1}^N \ln p(x) dx = 0.$$

- Plugging in the expressions for $p(x)$ from equation (13) yields

$$\int_{N-1}^I [\ln R - \ln \gamma_M(x)] dx + \int_I^N [\ln W - \ln \gamma_L(x)] dx = 0$$

- Substituting the expressions for R and W from (14), we obtain

$$\int_{N-1}^I [\ln Y - \ln(K/(I - N + 1)) - \ln \gamma_M(x)] dx + \int_I^N [\ln Y - \ln(L/(N - I)) - \ln \gamma_L(x)] dx = 0.$$

- This equation can be rearranged as

$$\begin{aligned} \ln Y &= \int_{N-1}^I \left[\ln \left(\frac{K}{I - N + 1} \right) + \ln \gamma_M(x) \right] dx + \int_I^N \left[\ln \left(\frac{L}{N - I} \right) + \ln \gamma_L(x) \right] dx \\ &= \int_{N-1}^I \ln \gamma_M(x) dx + \int_I^N \ln \gamma_L(x) dx \\ &\quad + (I - N + 1) \ln \left(\frac{K}{I - N + 1} \right) + (N - I) \ln \left(\frac{L}{N - I} \right), \end{aligned}$$

which, after taking exponentials on both sides of the equation, yields the expression for aggregate output in equation (1) in the text.

Assumption A1

We now show that Assumption A1 is equivalent to the capital-labor ratio of the economy taking an intermediate value. In particular, there exist two positive thresholds $\underline{\kappa} < \bar{\kappa}$ such that Assumption A1 holds whenever

$$\frac{K}{L} \in (\underline{\kappa}, \bar{\kappa}). \quad (\text{A2})$$

Equation (14) shows that

$$\frac{W}{R} = \frac{K}{L} \frac{N - I}{I - N + 1}.$$

Define

$$\underline{\kappa} = \frac{I - N + 1}{N - I} \frac{\gamma_L(I)}{\gamma_M(I)}, \quad \text{and} \quad \bar{\kappa} = \frac{I - N + 1}{N - I} \frac{\gamma_L(N)}{\gamma_M(N - I)}.$$

Then equation A2 is equivalent to Assumption A1.

Derivations in the Presence of Technology-Skill Mismatch

- Denote by $p(x)$ the price of task x . Assumption A1 together with the fact that $W_H > W_L$ (see footnote 12) implies

$$p(x) = \begin{cases} \frac{R}{\gamma_M(x)} & \text{if } x \in [0, I] \\ \frac{W_L}{\gamma_L(x)} & \text{if } x \in (I, S) \\ \frac{W_H}{\gamma_L(x)} & \text{if } x \in [S, N] \end{cases}.$$

- Following the same steps as in our baseline model, we obtain the market-clearing condition for capital,

$$K = \frac{Y}{R}(I - N + 1).$$

- The demand for low-skill labor in task x is given by

$$\ell(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ \frac{Y}{W_L} & \text{if } x \in (I, S) \\ 0 & \text{if } x \in [S, N]. \end{cases}.$$

- Aggregating the demand for low-skill labor and setting it equal to its inelastic supply,

L , we obtain the market-clearing condition for low-skill labor as

$$L = \frac{Y}{W_L}(S - I),$$

which implies the expression for W_L given in the main text.

- The demand for high-skill labor in task x is given by

$$h(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ 0 & \text{if } x \in (I, S) \\ \frac{Y}{W_H} & \text{if } x \in [S, N]. \end{cases}.$$

- Aggregating the demand for high-skill labor and setting it equal to its supply, H , we obtain the market-clearing condition for high-skill labor as

$$H = \frac{Y}{W_H}(N - S),$$

which implies the expression for W_H given in the main text.

Derivations for the Model with Distortions

- Denote by $p(x)$ the price of task x . The variant of Assumption A1 introduced in Section 5 implies

$$p(x) = \begin{cases} \frac{R(1-\tau)}{\gamma_M(x)} & \text{if } x \in [0, I] \\ \frac{W(1+\omega)}{\gamma_L(x)} & \text{if } x \in (I, J) \\ \frac{W}{\gamma_L(x)} & \text{if } x \in [J, N]. \end{cases}$$

- Following the same steps as in the model with no distortions, we obtain the market-clearing condition for capital,

$$K = \frac{Y}{R(1-\tau)}(I - N + 1).$$

- The demand for labor in task x is

$$\ell(x) = \begin{cases} 0 & \text{if } x \in [0, I] \\ \frac{Y}{W(1+\omega)} & \text{if } x \in (I, J) \\ \frac{Y}{W} & \text{if } x \in [J, N] \end{cases}.$$

- The expression for $\ell(x)$ implies that the total amount of labor employed in tasks where labor gets rents is

$$L_A = \frac{Y}{W(1 + \omega)}(J - I).$$

The total amount of labor employed in tasks where labor does not get rents is

$$L - L_A = \frac{Y}{W}(N - J).$$

To derive the expression for (gross) output we proceed as follows:

- Again from our choice of numeraire, we have

$$\int_{N-1}^N \ln p(x) dx = 0.$$

- Plugging in the expressions for $p(x)$ we obtain

$$\begin{aligned} \int_{N-1}^I [\ln R - \ln \gamma_M(x)] dx + \int_I^J [\ln W + \ln(1 + \omega) - \ln \gamma_L(x)] dx \\ + \int_J^N [\ln W - \ln \gamma_L(x)] dx = 0 \end{aligned}$$

- Substituting for factor prices using the expressions for K , L_A and $L - L_A$, we obtain

$$\begin{aligned} \int_{N-1}^I [\ln Y - \ln(K/(I - N + 1)) - \ln \gamma_M(x)] dx \\ + \int_I^J [\ln Y - \ln(L_A/(J - I)) - \ln \gamma_L(x)] dx \\ + \int_I^J [\ln Y - \ln((L - L_A)/(N - J)) - \ln \gamma_L(x)] dx = 0. \end{aligned}$$

- This equation can be rearranged as

$$\begin{aligned} \ln Y &= \int_{N-1}^I \left[\ln \left(\frac{K}{I - N + 1} \right) + \ln \gamma_M(x) \right] dx + \int_I^J \left[\ln \left(\frac{L_A}{J - I} \right) + \ln \gamma_L(x) \right] dx \\ &\quad + \int_J^N \left[\ln \left(\frac{L}{N - J} \right) + \ln \gamma_L(x) \right] dx \\ &= \int_{N-1}^I \ln \gamma_M(x) dx + \int_I^N \ln \gamma_L(x) dx \\ &\quad + (I - N + 1) \ln \left(\frac{K}{I - N + 1} \right) + (J - I) \ln \left(\frac{L_A}{J - I} \right) + (N - J) \ln \left(\frac{L - L_A}{N - J} \right), \end{aligned}$$

which yields equation (12) in the text.

REFERENCES

- Accenture PLC (2017)** “How Companies Are Reimagining Business Processes With IT,” <https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/>
- Acemoglu, Daron and David Autor (2011)** “Skills, tasks and technologies: Implications for employment and earnings,” *Handbook of Labor Economics*, 4: 1043–1171.
- Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price (2014)** “Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing,” *American Economic Review: Papers & Proceedings* 104(5):394–99.
- Acemoglu, Daron and Pascual Restrepo (2016)** “The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment,” forthcoming *American Economic Review*.
- Acemoglu, Daron and Pascual Restrepo (2017)** “Robots and Jobs: Evidence from US Labor Markets” NBER Working Paper No. 23285.
- Acemoglu, Daron and Pascual Restrepo (2018a)** “Low-Skill and High-Skill Automation” forthcoming *Journal of Human Capital*.
- Acemoglu, Daron and Pascual Restrepo (2018b)** “Excessive Automation: Technology Adoption and Worker Displacement in a Frictional World,” Mimeo.
- AEI (2008)** “Taxing Capital,” Report by the American Enterprise Institute, available at <https://www.aei.org>.
- Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones (2017)** “Artificial Intelligence and Economic Growth,” NBER Working Paper No. 23928.
- Allen, Robert C. (2009)** “Engels’ Pause: Technical Change, Capital Accumulation, and Inequality in the British Industrial Revolution,” *Explorations in Economic History*, 46(4): 418–435.
- Autor, David H., Frank Levy and Richard J. Murnane (2003)** “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118(4): 1279–1333.
- Autor, David H., David Dorn, and Gordon H. Hanson (2013)** “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103(6): 2121–68
- Autor, David H. (2015)** “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, 29(3):3–30.
- Ayres, Robert, and Steven M. Miller (1983)** *Robotics: Applications and Social Implications*, Ballinger Publishing Company.
- James, Bessen (2016)** *Learning by Doing: The Real Connection Between Innovation, Wages, and Wealth*. New Haven: Yale University Press.

Boston Consulting Group (2015) “The Robotics Revolution: The Next Great Leap in Manufacturing.”

Bloom, Nicholas, Charles I. Jones, John Van Reenen and Michael Webb (2017) “Are Ideas Getting Harder to Find?” NBER Working Paper No. 23782.

Brynjolfsson, Erik and Andrew McAfee (2014) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton & Company.

Chandler, Alfred D. (1977) *The Visible Hand: The Managerial Revolution in American Business*, Harvard University Press: Cambridge, MA.

Deloitte and The Manufacturing Institute (2011) “Boiling Point? The Skills Gap in U.S. Manufacturing.” Report from <http://www.themanufacturinginstitute.org>.

Ford, Martin (2015) *The Rise of the Robots*, Basic Books, New York.

Graetz, Georg and Guy Michaels (2015) “Robots at Work,” CEP Discussion Paper No 1335.

Goldin, Claudia and Larry Katz (2010) *The Race Between Education and Technology*, Belknap Press for Harvard University Press.

Gordon, Robert J. (2016) *The Rise and Fall of American Growth The U.S. Standard of Living since the Civil War*, Princeton University Press.

Groover, Mikell (1983) *CAD/CAM: Computer-aided Design and Manufacturing*, Englewood Cliffs, NJ: Prentice Hall.

Groover, Mikell, Mitchell Weiss, Roger N. Nagel, and Nicholas G. Odrey (1986) *Industrial Robotics: Technology, Programming and Applications*, McGraw-Hill Inc.

Herrendorf, Berthold, Richard Rogerson and Ákos Valentinyi (2013) “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, 103(7): 2752–2789.

Kolb, David A. (1984) *Experiential Learning: Experience as the Source of Learning and Development*, Englewood Cliffs, NJ: Prentice Hall.

Kuznets, Simon (1966) *Modern Economic Growth*. New Haven: Yale University Press.

Landes, David (1969) *The Unbound Prometheus*, Cambridge University Press: New York.

Manuelli, Rodolfo E. and Ananth Seshadri (2014) “Frictionless Technology Diffusion: The Case of Tractors,”

Mantoux, Paul (1928) *The Industrial Revolution in the Eighteenth Century: An Outline of the Beginnings of the Modern Factory System in England*. New York: Harcourt.

McKinsey (2017) “Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation.” Report from <https://www.mckinsey.com>.

Mian, Atif and Amir Sufi (2014) “What Explains the 2007-2009 Drop in Employment?” *Econometrica* 82(6): 2197–2223.

- Minsky, Marvin (1967)** *Computation: Finite and Infinite Machines*, Englewood Cliffs, N.J.: Prentice-Hall.
- Mokyr, Joel (1990)** *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press: New York.
- Nilsson, Nils J. (2009)** *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge University Press.
- Olmstead, Alan L. and Paul W. Rhode (2001)** “Reshaping the Landscape: The Impact and Diffusion of the Tractor in American Agriculture, 1910-1960,” *The Journal of Economic History* 61(3): 663–98.
- Pew Research Center (2017)** “Automation in Everyday Life.” Online report obtained from <http://www.pewinternet.org/2017/10/04/automation-in-everyday-life/>
- Rasmussen, Wayne D. (1982)** “The Mechanization of Agriculture,” *Scientific American*, 247(3): 76–89.
- Simon, Herbert A. (1965)** *The Shape of Automation for Men and Management*, New York: Harper & Row.
- Syverson, Chad (2017)** “Challenges to Mismeasurement Explanations for the US Productivity Slowdown,” *Journal of Economic Perspectives* 31(2): 165–186.
- Tuzel, Selale and Miao Ben Zhang (2017)** “Economic Stimulus at the Expense of Routine-Task Jobs,” Mimeo, Marshall School of Business, USC.
- Zeira, Joseph (1998)** “Workers, Machines, and Economic Growth,” *Quarterly Journal of Economics*, 113(4): 1091–1117.