

Low-Skill and High-Skill Automation

Daron Acemoglu
Massachusetts Institute of Technology

Pascual Restrepo
Boston University

We present a task-based model in which high- and low-skill workers compete against machines in the production of tasks. Low-skill (high-skill) automation corresponds to tasks performed by low-skill (high-skill) labor being taken over by capital. Automation displaces the type of labor it directly affects, depressing its wage. Through ripple effects, automation also affects the real wage of other workers. Counteracting these forces, automation creates a positive productivity effect, pushing up the price of all factors. Because capital adjusts to keep the interest rate constant, the productivity effect dominates in the long run. Finally, low-skill (high-skill) automation increases (reduces) wage inequality.

I. Introduction

Much has been written on the automation of routine and manual tasks, where machines, computers, and robots replace white-collar and blue-collar workers, typically in middle- and low-wage occupations (e.g., Autor, Levy, and Murnane 2003; Goos and Manning 2007; Michaels, Natraj, and Van Reenen 2014; Acemoglu and Restrepo 2017). In this traditional view, high-skill workers are shielded from automation because they specialize in more complex tasks requiring human judgment, problem solving, analytical skills, or various soft skills. However, recent advances in artificial intelligence cast doubt on this narrative. The automation of the complex tasks in which high-skill workers specialize—what we refer to as “high-skill automation”—is on its way to becoming a potent force in the US labor market. The new generation of artificial-intelligence technology, in conjunction with advances in big data and machine learning, already has the potential to perform many tasks in which human judgment was pre-

This paper is prepared for the special issue of the *Journal of Human Capital* in honor of Gary Becker. We thank two anonymous referees and the editor, Isaac Ehrlich, for useful comments. Financial support from the Sloan Foundation, the Smith Richardson Foundation, and Google are gratefully acknowledged.

[*Journal of Human Capital*, 2018, vol. 12, no. 2]
© 2018 by The University of Chicago. All rights reserved. 1932-8575/2018/1202-0003\$10.00

viously thought to be indispensable. Occupations facing (partial) automation from advances in artificial intelligence include accounting, mortgage origination, management consulting, financial planning, paralegal services, and various medical specialties, including radiology, general practice, or even surgery. A recent McKinsey study, for instance, concludes that “a significant percentage of the activities performed by even those in the highest-paid occupations (for example, financial planners, physicians, and senior executives) can be automated by adapting current technology” (Chui, Mayika, and Miremedi 2016).

In another of its reports, McKinsey declares the end of managers’ comparative advantage and gives the example of a Hong Kong venture-capital firm that has appointed a decision-making algorithm to its board of directors. It points to “the most impressive examples of machine learning substituting for human pattern recognition—such as the IBM supercomputer Watson’s potential to predict oncological outcomes more accurately than physicians by reviewing, storing, and learning from reams of medical-journal articles” (Dewhurst and Willmott 2014). Silicon Valley entrepreneur and author Martin Ford similarly asserts, “It’s not just about lower-skilled jobs either. People with college degrees, even professional degrees, people like lawyers are doing things that ultimately are predictable. A lot of those jobs are going to be susceptible over time” (quoted in McNeal 2015).

Despite this rapid and potentially transformative rise of high-skill automation, there is relatively little work studying its labor market implications. This paper is a first attempt to develop a simple framework incorporating both the more traditional automation of routine and manual jobs—what we refer to as “low-skill automation”—and high-skill automation.

We extend the task-based models originally developed in Acemoglu and Autor (2011) and Acemoglu and Restrepo (2016), which in turn build on Zeira (1998) and Acemoglu and Zilibotti (2001). In our model, a continuum of tasks can be performed by low-skill labor, high-skill labor, or capital. Crucially, the range of tasks that can be performed by capital expands as a result of two types of automation technologies. Low-skill automation expands the range of tasks that capital can perform at the low end of the complexity distribution of tasks.¹ The second, corresponding to high-skill automation, is the new element in our model and is based on the assumption that new developments in artificial intelligence allow capital to compete against high-skill labor in complex tasks.

¹ “Low-skill automation” here refers to the more traditional automation of routine and manual jobs, even though some routine tasks often involve nontrivial skill requirements and some basic tasks have not been much affected by automation at all (e.g., personal services). Likewise, “low-skill labor” refers here to blue-collar and white-collar workers who tend to specialize in the routine and manual tasks that have been more prone to automation in the past 30 years (e.g., clerks, bookkeepers, accountants, welders, assemblers). Thus, we abstract from the role of personal-service jobs performed by low-skill workers that have not been much affected by automation at all (Autor and Dorn 2013).

This framework departs from existing models not only in allowing for two types of automation but also in considering an environment in which there is no simple “comparative advantage” (or single crossing) across factors and tasks. In Acemoglu and Autor (2011) and Acemoglu and Restrepo (2016), as well as in models in the assignment literature, such as Sattinger (1975), Teulings (1995), and Costinot and Vogel (2010), there is a simple comparative-advantage ranking, where some workers are proportionately more productive relative to others in more complex tasks. To study high-skill automation, one needs to generalize this structure and allow for a richer pattern of comparative advantage, where capital has a comparative advantage not only at routine and manual tasks with low complexity but also at complex tasks that would be produced by high-skill labor otherwise.² The development of a tractable framework with a richer comparative-advantage structure for capital is one of the main contributions of our paper.³

We start with a static economy with a given supply of capital as well as inelastically supplied low-skill and high-skill labor. We first establish the existence of an equilibrium in this economy and characterize the potential assignments of tasks to factors. The most novel pattern—and the one that is a direct consequence of the richer structure of comparative advantage that we introduce—is one in which capital performs both the least complex tasks (where it directly competes with low-skill labor) and a disjoint range of more complex tasks (where it directly competes with high-skill labor).

We then characterize the implications of low-skill automation, which corresponds to an expansion in the set of tasks that can be performed by capital at the bottom of the distribution, and high-skill automation, which corresponds to an expansion in the set of tasks that can be performed by capital toward the higher end of the distribution. We show that both types of automation create two distinct impacts: a displacement effect and a productivity effect. The displacement effect, by taking away tasks from the directly affected factor, harms the labor market fortunes of that factor, while the productivity effect tends to increase the wages of all factors. We then demonstrate that the total impact of either type of automation on the wages of low-skill and high-skill labor is given by the sum of its displacement and productivity effects. When the displacement effect dominates, factors affected by automation experience a decline in

² The possibility that automation takes place across a disjoint set of tasks is important to model the possibility that fairly complex functions involved in financial planning, accounting, management, or medical occupations can be automated while other tasks of middle complexity (including various functions in manufacturing, construction, and personal communication) remain nonautomated.

³ A recent paper by Feng and Graetz (2016) also makes a related contribution. They argue that human labor has a comparative advantage not only in nonroutine tasks but also in intuitive tasks with few training requirements—a phenomenon known as Moravec’s paradox. Relatedly, Hémous and Olsen (2016) model the interplay between automation and horizontal innovations in the context of endogenous growth.

their wages. Most interestingly, the displacement caused by automation also creates ripple effects. High-skill automation displaces high-skill labor, which may then compete with low-skill labor in other tasks and displace this latter group. Because of these ripple effects, automation could depress the wages not just of the affected factor but of both factors. For instance, high-skill automation can reduce the real wages of both low-skill and high-skill labor. Nevertheless, the displacement effect on the directly affected group is always greater, and thus low-skill automation increases the inequality between high-skill and low-skill labor, while high-skill automation has the opposite effect.⁴

After this analysis, we turn to the “long-run” implications of automation and allow for capital accumulation to restore the price of capital to its long-run level.⁵ In the long run, the productivity effect becomes stronger. This is for the intuitive reason that automation, by increasing the demand for capital, increases the price of capital in the short run, which dampens the potential productivity gains that can be obtained by substituting the cheaper capital for the more expensive labor in the automated tasks. In the long run, the price of capital remains constant, and thus there will be greater productivity gains. It is for this reason that, in Acemoglu and Restrepo (2016), automation was found to always increase wages in the long run. Here, with two types of labor and two types of automation, we find that automation increases the wage bill in the long run but might still have a negative impact on the wages of the type of labor that it directly displaces.

In addition to the theoretical literature on task-based models and assignment models, which we have already discussed, our paper is related to the empirical literature on the effects of automation and robotics on the labor market. Autor, Levy, and Murnane (2003) documented the decline of employment in jobs comprising routine tasks and argue that these shifts reflect the computerization of such tasks. Michaels, Natraj, and Van Reenen (2014) show that the replacement of routine tasks by information and communication technologies caused a decline in employment opportunities for middle-skill workers.⁶ In Acemoglu and Restrepo (2017), we document that, from 1990 to 2007, US commuting zones that harbored industries more exposed to the use of industrial robots experi-

⁴ This result echoes the work of Ehrlich and Kim (2015), who explore how migrants compete not only against low-skill natives in some segments of the market but also against high-skill natives in others. In their setting, immigrants displace workers in some industries and, depending on which workers they directly substitute for, increase or reduce inequality. As in our context, skilled immigrants also create a productivity effect.

⁵ To economize on space, we do this without explicitly allowing for dynamics, though doing this is straightforward, as in Acemoglu and Restrepo (2016). Note, however, that in contrast to that paper, we do not endogenize technological change or the speed of automation (or the creation of new tasks).

⁶ Other empirical studies on the impact of the automation and computerization of routine tasks include Goos and Manning (2007), Acemoglu and Autor (2011), Autor and Dorn (2013), Foote and Ryan (2014), Goos, Manning, and Salomons (2014), Jaimovic and Siu (2014), Autor, Dorn, and Hanson (2015), and Gregory, Salomons, and Zierahn (2016).

enced a significant decline in employment and real wages. The negative effects concentrate on blue-collar workers in the lower end of the skill distribution. Relatedly, using a panel of industries in 17 countries from 1993 to 2007, Graetz and Michaels (2015) show that investments in industrial robots were associated with faster productivity growth and higher wages but also created some negative effects on employment for low-skill and middle-skill workers. Overall, the evidence on the impact of the automation of routine tasks and the use of industrial robots is in line with the theoretical implications of our model regarding “low-skill” automation.

The rest of the paper is organized as follows. Section II introduces our model. Section III characterizes the short-run equilibrium (where the supply of capital is taken as given) and highlights the different types of configurations that can arise. Of those, we focus on a situation in which capital competes directly both against low-skill and high-skill labor. Section IV characterizes the impact of automation on factor prices and inequality. In Section V, we study the long-run equilibrium of this model. The main difference in this case is that the productivity effect is amplified by the induced accumulation of capital following automation. As a result, automation cannot reduce the wages of both types of labor in the long run, though it can still depress the wage of the directly affected factor. Section VI returns to the other types of equilibria of the model and shows that they do not permit the simultaneous impact of automation on both low-skill and high-skill labor. Section VII concludes, while the appendix contains the proofs omitted from the text.

II. A Model of Low-Skill and High-Skill Automation

We consider a static economy with a unique final good Y , produced by combining a continuum 1 of tasks $y(i)$ with an elasticity of substitution $\sigma \in (0, \infty)$:

$$Y = \left(\int_0^1 y(i)^{(\sigma-1)/\sigma} di \right)^{\sigma/(\sigma-1)}. \quad (1)$$

The final good is produced competitively. Consumer utility is defined over the unique final good, and we normalize its price to 1.

Final-good producers can produce each task with machines (capital) or labor, and there are two types of labor, high- and low-skill. All tasks can be produced by both types of labor, though they have different productivities in each task. In particular, one unit of high-skill labor can produce $\gamma_H(i)$ units of task i , and one unit of low-skill labor can produce $\gamma_L(i)$ units of task i . Throughout, we assume that these productivities satisfy the following (strict) comparative-advantage structure.

ASSUMPTION 1 (Comparative-advantage assumption). The functions $\gamma_H(i)$, $\gamma_L(i)$, and $\gamma_H(i)/\gamma_L(i)$ are continuous and strictly increasing.

Combined with the pattern of productivity of machines across tasks specified in the next paragraph, the feature that γ_H and γ_L are increasing

enables us to determine the allocation of tasks between capital and labor in a tractable manner.⁷ That their ratio is strictly increasing implies that high-skill labor has (strict) comparative advantage relative to low-skill labor in higher-indexed tasks, which is a feature shared with Sattinger (1975), Teulings (1995), Costinot and Vogel (2010), Acemoglu and Autor (2011), and Acemoglu and Restrepo (2016), among others. Continuity is imposed for simplicity.

In contrast to these papers, however, we depart from the “super-modular” comparative-advantage structure across all factors. Namely, in these papers the productivities of any two factors across tasks satisfy an increasing-differences (or single-crossing) assumption.⁸ Yet such a structure implies that capital could not effectively compete against both types of labor, and this would not allow an interesting analysis of simultaneously ongoing low-skill and high-skill automation. We therefore abandon the supermodular comparative-advantage structure across all factors and tasks by assuming that there exists $J \in (0, 1)$ such that, when automated, tasks $i < J$ can be produced with capital with productivity 1, while tasks $i \geq J$ can be produced with capital with productivity $\gamma_K \geq 1$. (Looking from the viewpoint of capital, we sometimes refer to tasks $i < J$ as “simple” tasks and to $i \geq J$ as “complex” tasks.) We think of the tasks $i < J$ as routine tasks that have been automated in the past 30 years through the use of information-processing technologies or industrial robots. Tasks $i \geq J$, on the other hand, are complex tasks that are in the early stages of automation via artificial intelligence, big data, and a new phase of robotics. When $\gamma_K > 1$, capital will be able to compete simultaneously against high-skill labor in some complex tasks and against low-skill labor in simpler tasks.

Not all tasks can be automated, however. As in Acemoglu and Restrepo (2016), we distinguish between technologically automated tasks, which can be automated if profitable, and tasks automated in equilibrium. We assume that there exists a pair of thresholds $I_L \in (0, J)$ and $i \in [J, I_H]$ such that the tasks $i \in [0, I_L]$ and $i \in [J, I_H]$ are “technologically automated.” They will be automated in equilibrium if it is profitable for them to be produced with capital at the prevailing factor prices. Regardless of factor prices, the tasks in (I_L, J) and $(I_H, 1]$ must be produced with labor.

⁷ Nothing fundamental changes if we make these schedules decreasing and also assume that the productivity of machines is decreasing even more steeply. The structure with the productivity of machines taking the form of a step function and γ_H and γ_L increasing greatly simplifies the exposition.

⁸ This is true of Acemoglu and Restrepo’s (2016) model with two types of labor and capital; of Sattinger’s (1975), Teulings’s (1995), and Costinot and Vogel’s (2010) assignment models with a continuum of tasks and skills; and of Acemoglu and Autor’s (2011) baseline model with three types of labor. The latter paper then introduces automation of “middling” tasks, but in doing so assumes that there are no other tasks in which capital can be used and that it is sufficiently cheap to take over all the tasks that are technologically automated.

We summarize the above discussion by writing the technologically feasible combinations of factors to produce different tasks, given by

$$y(i) = \begin{cases} \gamma_H(i)h(i) + \gamma_L(i)l(i) + k(i) & \text{if } i \in [0, I_L], \\ \gamma_H(i)h(i) + \gamma_L(i)l(i) & \text{if } i \in (I_L, J), \\ \gamma_H(i)h(i) + \gamma_L(i)l(i) + \gamma_K k(i) & \text{if } i \in [J, I_H], \\ \gamma_H(i)h(i) + \gamma_L(i)l(i) & \text{if } i \in (I_H, 1]. \end{cases} \quad (2)$$

Here, $h(i)$, $l(i)$, and $k(i)$ denote the total quantities of high-skill labor, low-skill labor, and capital utilized in the production of task i , respectively.

We start by assuming that all factors are supplied inelastically, and we denote the supply of high-skill labor by H , that of low-skill labor by L , and that of capital by K .

III. Equilibrium

A “short-run equilibrium” is defined by factor prices—wages and a capital rental rate—of high-skill labor, low-skill labor, and capital— W_H , W_L , and R , respectively—such that final-good producers minimize costs and the three factor markets clear. Since final-good producers are competitive and have access to a constant returns to scale production function, cost minimization is equivalent to profit maximization.

We now characterize the equilibrium allocation of tasks to factors in this economy. Throughout, to simplify our notation, we assume that when indifferent between using capital or labor, a firm produces with capital. Likewise, when indifferent between using high- or low-skill labor, a firm produces with high-skill labor.⁹

PROPOSITION 1 (Equilibrium existence). Suppose that assumption 1 holds. For any $H, L, K > 0$ there is a unique equilibrium.¹⁰ The equilibrium is characterized by thresholds $I_L^* \in [0, I_L]$, $I_H^* \in [J, I_H]$, and $M \in (I_L^*, 1)$ such that

- capital produces the tasks in $[0, I_L^*] \cup [J, I_H^*]$;¹¹
- high-skill labor produces the tasks in $[M, 1]$ that are not produced with capital; and
- low-skill labor produces the tasks in $[0, M)$ that are not produced with capital.

⁹ This choice does not affect the results because firms are indifferent between producing with different factors in a set of tasks of measure zero.

¹⁰ Uniqueness here is under the tie-breaking assumption specified before the proposition. Without this assumption, we can instead establish “essential uniqueness,” meaning that the equilibrium allocation will be uniquely determined except at a finite number of threshold tasks at which firms are indifferent between using different factors.

¹¹ Here, we adopt the convention that, when $I_L^* = 0$, capital produces only the tasks in $[J, I_H^*]$. Likewise, when $I_H^* = J$, capital produces only the tasks in $[0, I_L^*]$. In equilibrium, only one of these sets can be empty.

Moreover, the threshold M is given by

$$\frac{W_L}{\gamma_L(M)} = \frac{W_H}{\gamma_H(M)}. \tag{3}$$

For the proof, see the appendix.

The main idea of this proposition is that, to minimize the cost of production, tasks will be allocated to factors depending on their comparative advantage. Our structure of comparative advantage implies that capital produces at most two disjoint sets of tasks $[0, I_L^*] \cup [J, I_H^*]$ (recall that one of these sets could be empty) and that there is a threshold M given by equation (3) defining which of the remaining tasks are allocated to low-skill and high-skill labor.

This result can be illustrated diagrammatically. Figure 1 plots the resulting allocations of tasks to factors when capital performs two disjoint sets of tasks. In the figure, $W_L/\gamma_L(i)$ and $W_H/\gamma_H(i)$ are the effective costs of producing task i with low- and high-skill labor, respectively. Likewise, R for $i \leq I_L$ and R/γ_K for $J \leq i \leq I_H$ are the effective costs of producing these different ranges of tasks with capital. In equilibrium, tasks will




Figure 1.—Cost-minimizing allocation of factors to tasks when capital produces two disjoint (and nonempty) sets of tasks. The top panel depicts an equilibrium in which $M \in (J, I_H^*)$. The middle panel depicts an equilibrium in which $M \in (I_L^*, J)$. The bottom panel shows an equilibrium in which $M \in (I_H^*, 1)$.

be allocated to factors that have the lowest effective cost of producing them.

In the top and middle panels, we present the cases in which $M \in (J, I_H^*)$ and $M \in (I_L^*, J)$, respectively. In these two cases, wages and the interest rate are such that low-skill labor specializes in low-indexed tasks and high-skill labor in high-indexed tasks. Capital performs some of the least complex tasks in $[0, I_L^*]$ because for $i \leq J$, its comparative advantage relative to low-skill labor is in lower-complexity tasks. Crucially, capital also performs some complex tasks in $[J, I_H^*]$. The difference between these two cases is merely in whether high-skill labor produces only tasks above those allocated to capital or whether it straddles the set of complex tasks allocated to capital. The bottom panel presents the case in which $M \in (I_H^*, 1)$. Here, capital also produces two disjoint sets of tasks, but it is in direct competition with low-skill labor in both.

In all of the above cases, the thresholds I_L^* and I_H^* , which we introduced in proposition 1, capture the possibility that not all technologically automated tasks will be produced with capital in equilibrium. As noted above, whether this is the case or not depends on factor prices. For instance, we could have that $I_L^* < I_L$ if the price of capital is sufficiently high and, consequently, firms would rather produce task I_L with low-skill labor even if it is possible to do so with capital. If this is the case, a further increase in I_L , corresponding to an expansion of the set of tasks that are technologically automated, will have no impact on the equilibrium allocation (and thus on prices). The same is true for an increase in I_H when $I_H^* < I_H$.

Throughout the paper, we center our analysis around the cases in which capital performs two (nonempty) disjoint sets of tasks and $M \in (I_L^*, I_H^*)$, shown in the top two panels of figure 1. These equilibria capture the more interesting situation in which one form of automation directly competes against low-skill workers and another form of automation directly competes against high-skill workers. We turn to the remaining types of equilibrium where automation competes directly against only a single type of labor in Section VI. Moreover, because our objective is to understand how changes in automation affect wages and inequality, we focus on the case where $I_H^* = I_H$ and $I_L^* = I_L$. Proposition A1, in the appendix, shows that there exists a threshold $\underline{\rho}$ and a threshold $\underline{K}(H, L)$ that is non-decreasing in H and L , such that, for $H/L > \underline{\rho}$ and $K > \underline{K}(H, L)$, the equilibrium features $M \in (I_L, I_H)$ and $I_L^* = I_L$ and $I_H^* = I_H$. Thus, until Section VI, we impose the following assumption on factor supplies:

ASSUMPTION 2. The supplies of labor and capital H, L, K satisfy $H/L > \underline{\rho}$ and $K > \underline{K}(H, L)$.

The condition $H/L > \underline{\rho}$ ensures that

$$\frac{W_L}{\gamma_L(I_H)} > \frac{W_H}{\gamma_H(I_H)}, \quad (4)$$

and so $M < I_H$ —high-skill labor is abundant and will face the competition of automation in the production of tasks near I_H . In addition, for a given

H and L , the condition $K > \underline{K}(H, L)$ ensures that capital is abundant and cheap relative to both types of labor, and so it is cheaper to produce tasks I_H and I_L , respectively, with capital:

$$\frac{W_H}{\gamma_H(I_H)} > \frac{R}{\gamma_K}, \tag{5}$$

$$\frac{W_L}{\gamma_L(I_L)} > R.$$

Let $\min\{J, M\}$ denote the minimum threshold where either there is a switch from simple to complex tasks or the effective costs of production by low-skill and high-skill labor are equated. Under assumption 2, capital performs the tasks in $[0, I_L] \cup [J, I_H]$, low-skill labor performs the tasks in $(I_L, \min\{J, M\})$, and high-skill labor performs the tasks in $[\min\{J, M\}, J) \cup (I_H, 1]$. (Note that when $M \geq J$, the set $[\min\{J, M\}, J)$ is empty.)

Given the allocation of tasks to factors derived above, we can determine the equilibrium prices of a task as the minimum effective cost of producing it:

$$p(i) = \begin{cases} R & \text{if } i \in [0, I_L], \\ \frac{W_L}{\gamma_L(i)} & \text{if } i \in (I_L, \min\{J, M\}), \\ \frac{W_H}{\gamma_H(i)} & \text{if } i \in [\min\{J, M\}, J), \\ \frac{R}{\gamma_K} & \text{if } i \in [J, I_H], \\ \frac{W_H}{\gamma_H(i)} & \text{if } i \in (I_H, 1]. \end{cases} \tag{6}$$

With these task prices $p(i)$, the equilibrium quantity of task i can be determined from the cost-minimization problem of final-good producers as

$$y(i) = Yp(i)^{-\sigma}. \tag{7}$$

Equations (6) and (7) combined imply that the demand for capital in each simple automated task is $YR^{-\sigma}$, the demand for capital in each complex automated task is $Y\gamma_K^{\sigma-1}R^{-\sigma}$, the demand for low-skill labor in each task performed by this factor is $Y\gamma_L(i)^{\sigma-1}W_L^{-\sigma}$, and the demand for high-skill labor in each task performed by this factor is $Y\gamma_H(i)^{\sigma-1}W_H^{-\sigma}$. Integrating these demands over the range of tasks assigned to the relevant factor, we find that factor-market-clearing conditions take the form

$$\begin{aligned} Y\Gamma_H W_H^{-\sigma} &= H, \\ Y\Gamma_L W_L^{-\sigma} &= L, \\ Y\Gamma_K R^{-\sigma} &= K, \end{aligned}$$

where, to simplify notation, we have defined the “effective shares” of high-skill labor, low-skill labor, and capital as

$$\begin{aligned} \Gamma_H &= \int_{\min\{J, M\}}^J \gamma_H(i)^{\sigma-1} di + \int_{I_H}^1 \gamma_H(i)^{\sigma-1} di, \\ \Gamma_L &= \int_{I_L}^{\min\{J, M\}} \gamma_L(i)^{\sigma-1} di, \\ \Gamma_K &= I_L + (I_H - J)\gamma_K^{\sigma-1}. \end{aligned} \tag{8}$$

Why we refer to these objects as effective shares will be clarified below by equation (9).

The following proposition provides explicit expressions for equilibrium factor prices as functions of the thresholds I_H , I_L , J , and M (where the last one is the only endogenous threshold determined in equilibrium).

PROPOSITION 2 (Equilibrium characterization). Suppose that assumptions 1 and 2 hold. Then, equilibrium output and factor prices as functions of the thresholds can be expressed as

$$Y = (\Gamma_H^{1/\sigma} H^{(\sigma-1)/\sigma} + \Gamma_L^{1/\sigma} L^{(\sigma-1)/\sigma} + \Gamma_K^{1/\sigma} K^{(\sigma-1)/\sigma})^{\sigma/(\sigma-1)}, \tag{9}$$

and

$$\begin{aligned} W_H &= Y^{1/\sigma} \Gamma_H^{1/\sigma} H^{-1/\sigma}, \\ W_L &= Y^{1/\sigma} \Gamma_L^{1/\sigma} L^{-1/\sigma}, \\ R &= Y^{1/\sigma} \Gamma_K^{1/\sigma} K^{-1/\sigma}, \end{aligned} \tag{10}$$

where Γ_H , Γ_L , and Γ_K are given by equation (8). Moreover, factor prices satisfy the ideal-price condition

$$\Gamma_H W_H^{1-\sigma} + \Gamma_L W_L^{1-\sigma} + \Gamma_K R^{1-\sigma} = 1, \tag{11}$$

and the endogenous threshold M is given implicitly by the unique solution in the interval (I_L, I_H) to equation (3):

$$\left(\frac{\Gamma_H L}{\Gamma_L H} \right)^{1/\sigma} = \frac{\gamma_H(M)}{\gamma_L(M)}. \tag{12}$$

For the proof, see the appendix.

This proposition clarifies why we refer to the terms Γ_H , Γ_L , and Γ_K as effective shares—they correspond to (endogenous versions of) the distribution parameters in the derived constant elasticity of substitution (CES)

aggregate production function in equation (9). Note also that the ideal-price condition follows as an additional equilibrium condition, since we chose the final good as numeraire.

The unique equilibrium value for M in proposition 2 is implicitly defined by the solution to equation (12). As shown in figure 2, the fact that this equation has a unique solution follows by observing that the right-hand side is a strictly increasing function of M , while the left-hand side is a nonincreasing function of M , which becomes constant for $M \geq J$. The condition $H/L > \underline{p}$ —which we assume to hold throughout—ensures that these two curves intersect for $M \in (I_L, I_H)$. The figure also presents the allocation of tasks to factors, depending on whether $M \leq J$.

The effective shares in the CES aggregator in equation (9) depend on the technology parameters I_H and I_L . Thus, proposition 2 also shows that the task framework provides a richer view of technology, where we are not limited to the usual factor-augmenting technologies but we could also think of changes in effective shares as being driven by technology. This general conception of technology generates many of the new possibilities that we explore in the next section, such as the possibility that automation may reduce all workers' wages in the short run.

IV. The Effect of Automation on Factor Prices

In this section we explore the effects of low-skill and high-skill automation. Our analysis is simplified by a straightforward consequence of equation (10): the impact of either type of automation on factor prices (W_H , W_L , and R) can be decomposed into a displacement effect and a productivity effect. To see this, we totally differentiate equation (10) to obtain

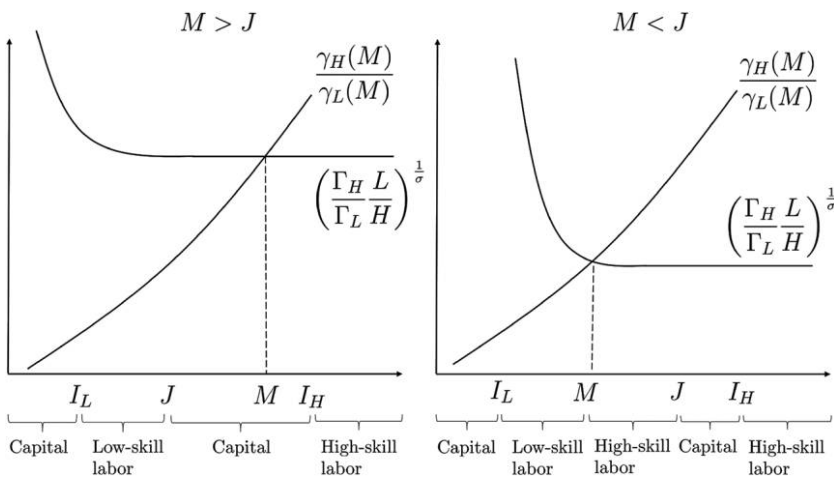


Figure 2.—Equilibrium value for M when assumption 2 holds. The left-hand panel depicts the case in which $M \in (I_L, J)$, while the right-hand panel is for the case in which $M \in (J, I_H)$.

$$\begin{aligned} \sigma \frac{dW_H}{W_H} &= \frac{d\Gamma_H}{\Gamma_H} + \frac{dY}{Y}, \\ \sigma \frac{dW_L}{W_L} &= \frac{d\Gamma_L}{\Gamma_L} + \frac{dY}{Y}, \\ \sigma \frac{dR}{R} &= \frac{d\Gamma_K}{\Gamma_K} + \frac{dY}{Y}. \end{aligned}$$

Here $d\Gamma_H/\Gamma_H$, $d\Gamma_L/\Gamma_L$, and $d\Gamma_K/\Gamma_K$ designate the displacement effects, while dY/Y designates the productivity effect. These expressions imply that the impact of technological change in general, and of the two types of automation in particular, works by changing the effective shares and the overall level of production in the economy—through the terms $d\Gamma_H/\Gamma_H$, $d\Gamma_L/\Gamma_L$, $d\Gamma_K/\Gamma_K$, and dY/Y .

Intuitively, the displacement effect matters because as tasks are reallocated away from a factor, there is a powerful downward pressure on the price of that factor; the reason is that such displacement pushes more of that factor to work in the remaining tasks, running into a downward-sloping demand for these tasks. The productivity effect arises from the fact that automation involves substituting cheaper capital for labor (and we know that capital has to be cheaper, since otherwise it would not have been profitable for firms to use capital instead of labor). Such substitution increases productivity and output in the economy. Because tasks are q -complements in the production of the final good, the increase in output raises the demand for all tasks and hence the price of all factors.

In the next two propositions, we characterize how the two types of automation shape first the displacement effects and then the productivity effects.

PROPOSITION 3 (Displacement effects of technology). Suppose that assumptions 1 and 2 hold. Let $\varepsilon = (\gamma'_H(M)/\gamma_H(M)) - (\gamma'_L(M)/\gamma_L(M)) \geq 0$ be the quasi elasticity of the comparative-advantage schedule. Automation has the following effects on output:¹²

1. An increase in I_L by $dI_L > 0$ —corresponding to low-skill automation—has the following impacts on effective shares: $d\Gamma_K/dI_L = 1$,

$$\frac{d\Gamma_L}{dI_L} = \begin{cases} -\gamma_L(I_L)^{\sigma-1} < 0 & \text{if } M \geq J, \\ -\gamma_L(I_L)^{\sigma-1} \frac{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H)}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} < 0 & \text{if } M < J, \end{cases}$$

and

¹² To economize on notation, we do not explicitly cover the case in which $M = J$, because the left and the right derivatives are different at this point. It can be shown that when $dI_L > 0$, $d\Gamma_L/dI_L$ and $d\Gamma_H/dI_L$ are identical in this case to the expressions for $M > J$ and that when $dI_L < 0$, they are identical to the expressions for $M < J$. Conversely, when $dI_H > 0$, $d\Gamma_H/dI_H$ and $d\Gamma_L/dI_H$ are given by the expressions for $M < J$, and when $dI_H < 0$, they are given by the expressions for $M > J$.

$$\frac{d\Gamma_H}{dI_L} = \begin{cases} 0 & \text{if } M \geq J, \\ -\gamma_L(I_L)^{\sigma-1} \frac{\gamma_H(M)^{\sigma-1}/\Gamma_L}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} < 0 & \text{if } M < J. \end{cases}$$

2. An increase in I_H by $dI_H > 0$ —corresponding to high-skill automation—has the following impact on effective shares: $d\Gamma_k/dI_H = \gamma_k^{\sigma-1}$,

$$\frac{d\Gamma_H}{dI_H} = \begin{cases} -\gamma_H(I_H)^{\sigma-1} < 0 & \text{if } M > J, \\ -\gamma_H(I_H)^{\sigma-1} \frac{\sigma\varepsilon + (\gamma_L(M)^{\sigma-1}/\Gamma_L)}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} < 0 & \text{if } M \leq J, \end{cases}$$

and

$$\frac{d\Gamma_L}{dI_H} = \begin{cases} 0 & \text{if } M > J, \\ -\gamma_H(I_H)^{\sigma-1} \frac{\gamma_L(M)^{\sigma-1}/\Gamma_H}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} < 0 & \text{if } M \leq J. \end{cases}$$

Proof. The proof follows by differentiating equation (8) and then substituting the derivatives involving M using the implicit-function theorem applied to equation (12). The full proof is presented in the appendix.

The main takeaway from this proposition is that both types of automation displace labor and reduce the set of tasks performed by workers. Namely, low-skill automation reduces the share of tasks performed by low-skill labor, and high-skill automation reduces the share of tasks performed by high-skill labor.

Importantly, when $M < J$, both types of automation create ripple effects, also reducing the effective shares of the other type of labor.¹³ For example, when $M < J$, low-skill automation displaces low-skill labor from tasks it previously performed, and these workers then compete for and take over some of the tasks previously performed by high-skill labor. Likewise, when $M < J$, high-skill automation reduces not only the effective share of high-skill labor but also that of low-skill labor. These ripple effects do not arise when $M > J$, because the two types of labor do not compete directly (the sets of tasks they produce are always buffered by tasks produced by capital). The ripple effects also disappear when $\varepsilon \rightarrow \infty$ —so that around the threshold task M , there is a very strong comparative advantage of high-skill labor in more complex tasks and of low-skill labor in simpler tasks. Intuitively, in this case, though the two types of labor do compete for the production of tasks around M , they are such poor substitutes that the ripple effects evaporate. Conversely, when $\varepsilon \rightarrow 0$, the comparative advantage of one type of labor relative to the other around the threshold task M is very small, and the ripple effects are maximized.

¹³ In fact, as indicated in n. 12, an increase in I_H also creates ripple effects when $M = J$.

PROPOSITION 4 (Productivity effect of technology). Suppose that assumptions 1 and 2 hold.

1. An increase in I_L by $dI_L > 0$ —corresponding to low-skill automation—increases aggregate output by

$$\frac{1}{Y} \frac{dY}{dI_L} = \frac{1}{\sigma - 1} \left[R^{1-\sigma} - \left(\frac{W_L}{\gamma_L(I_L)} \right)^{1-\sigma} \right] > 0.$$

2. An increase in I_H by $dI_H > 0$ —corresponding to high-skill automation—increases aggregate output by

$$\frac{1}{Y} \frac{dY}{dI_H} = \frac{1}{\sigma - 1} \left[\left(\frac{R}{\gamma_K} \right)^{1-\sigma} - \left(\frac{W_H}{\gamma_H(I_H)} \right)^{1-\sigma} \right] > 0.$$

Proof. The proof follows by differentiating equation (9). The full argument is presented in the appendix.

This proposition thus shows that there are productivity gains from both types of automation, helping to contribute to higher wages for both types of labor (or higher prices for all factors). Notably, this is true regardless of whether $M \leq J$.

Another noteworthy result in proposition 4 is a quantification of the extent of productivity effects. In particular, the greater is the gap $W_H/\gamma_H(I_H) - R/\gamma_K$ or the gap $W_L/\gamma_L(I_L) - R/\gamma_K$, the greater are the cost savings by substituting capital for the more expensive labor factor and the greater is the productivity effect (assumption 2 guarantees that both of these gaps are positive). This observation also implies that as $W_L/\gamma_L(I_L) \downarrow R$, productivity gains—and thus the productivity effect—from low-skill automation disappear; likewise, as $W_H/\gamma_H(I_H) \downarrow R/\gamma_K$, the productivity effect from high-skill automation disappears.

As observed above, the impact of automation on wages can be directly obtained by combining the displacement and productivity effects. In general, since these two effects go in opposite directions, we cannot unambiguously determine the impact of automation on all factor prices. Nevertheless, it is possible to characterize when one effect will dominate. Though there are different ways of doing this, here we emphasize the role of the gap between the effective cost of production by capital and labor inputs. Since the price of capital (the rental rate) will be higher when capital is more scarce, this leads to a comparison in terms of the level of capital stock in the economy, as shown in the next proposition.

PROPOSITION 5 (Factor prices and automation). Suppose that assumptions 1 and 2 hold. Then for a fixed H, L there exist thresholds $\underline{K}(H, L) < \underline{K}_L < \bar{K}_L$ and $\underline{K}(H, L) < \underline{K}_H < \bar{K}_H$ such that¹⁴

¹⁴ We do not give the comparative statics in the cases in which the capital stock, K , is exactly equal to the thresholds, to shorten the proposition. As is evident from the rest of the proposition, in these cases, it will have no effect on the price of one of the factors.

1. When $M < J$, low-skill automation (an increase in I_L) has the following effects on wages:
 - if $K \in (\underline{K}(H, L), \underline{K}_L)$, it reduces both W_H and W_L ;
 - if $K \in (\underline{K}_L, \bar{K}_L)$, it reduces W_L and increases W_H ;
 - if $K > \bar{K}_L$, it increases both W_H and W_L .

Also, again when $M < J$, high-skill automation (an increase in I_H) has the following effects on wages:

- if $K \in (\underline{K}(H, L), \underline{K}_H)$, it reduces both W_H and W_L ;
 - if $K \in (\underline{K}_H, \bar{K}_H)$, it reduces W_H and increases W_L ;
 - if $K > \bar{K}_H$, it increases both W_H and W_L .
2. If, on the other hand, $M > J$, we have that
 - if $K > \bar{K}_L$, low-skill automation increases both W_H and W_L , and if $K < \bar{K}_L$, it reduces W_L and increases W_H ;
 - similarly, if $K > \bar{K}_H$, high-skill automation increases both W_H and W_L , and if $K < \bar{K}_H$, it reduces W_H and increases W_L .
 3. Both types of automation always increase the rental rate of capital, R .

For the proof, see the appendix.

This proposition is one of the main results of the paper. First, it shows that, when the price of capital (the rental rate) is high relative to wages, automation directed to a particular type of labor reduces the wage rate of that type of labor—so low-skill automation reduces low-skill wages and high-skill automation reduces high-skill wages. This result is reversed, however, when the productivity effect is sufficiently powerful, which, as shown in proposition 4, happens when capital is sufficiently abundant and the price of capital (the rental rate) is low. Second, this proposition also demonstrates the implications of the ripple effect, which was noted in our discussion of proposition 3. When there is a ripple effect ($M < J$) and when the productivity effect is not too powerful, low-skill automation also reduces high-skill wages and high-skill automation also reduces low-skill wages. This result, which to the best of our knowledge is unique to the framework with the two types of automation developed here, is important in highlighting how very specific types of automation technologies can depress wages throughout the wage distribution.

Nevertheless, the effects of the two types of automation technologies on inequality, which in our model is given by the ratio of high-skill to low-skill wages and is proportional to $\omega = (W_H/W_L)^\sigma$, always goes in the intuitive direction, as shown in the next proposition for factor prices.

PROPOSITION 6 (Automation and inequality). Suppose that assumptions 1 and 2 hold.

1. Low-skill automation increases wage inequality, that is,

$$\frac{1}{\omega} \frac{d\omega}{dI_L} = \begin{cases} \frac{\gamma_L(I_L)^{\sigma-1}}{\Gamma_L} > 0 & \text{if } M \geq J, \\ \frac{\gamma_L(I_L)^{\sigma-1}}{\Gamma_L} \frac{\sigma\varepsilon}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} > 0 & \text{if } M < J. \end{cases}$$

2. High-skill automation reduces wage inequality, that is,

$$\frac{1}{\omega} \frac{d\omega}{dI_H} = \begin{cases} -\frac{\gamma_H(I_H)^{\sigma-1}}{\Gamma_H} < 0 & \text{if } M > J, \\ -\frac{\gamma_H(I_H)^{\sigma-1}}{\Gamma_H} \frac{\sigma\varepsilon}{\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)} < 0 & \text{if } M \leq J. \end{cases}$$

Proof. The proof follows from differentiating the expression $\omega = (\Gamma_H/\Gamma_L)(L/H)$. QED

One noteworthy feature is that the effect of either type of automation on wage inequality is lower when $M < J$ (for $\sigma\varepsilon/[\sigma\varepsilon + (\gamma_H(M)^{\sigma-1}/\Gamma_H) + (\gamma_L(M)^{\sigma-1}/\Gamma_L)] < 1$), because in this case there are ripple effects on the wages of the factor that are not directly affected by automation, and this limits the impact on inequality. In fact, when $M < J$, we can also see that as $\varepsilon \rightarrow 0$, the ripple effects become so powerful (because there is little comparative advantage protecting the factor not directly affected by automation) that the impact of both types of automation on inequality vanishes.

Propositions 3–6 provide a new perspective on what to expect from technological developments that automate tasks performed by high-skill workers. High-skill automation increases productivity and reduces inequality. But importantly, it might indirectly hurt low-skill workers by pushing more skilled ones to compete against them in simpler tasks. These predictions resemble the reversal in the demand for skills and cognitive tasks documented by Beaudry, Green, and Sand (2016). The reversal started in 2000 but has accelerated in recent years. Interestingly, the reversal has been accompanied by a movement of high-skill workers down the skill ladder, displacing low-skill workers in less skill-intensive jobs (see also Modestino, Shoag, and Ballance 2016). This cascading pattern is reminiscent of the ripple effects that arise in our model when $M < J$.

The empirical evidence on the impact of the use of industrial robots in manufacturing and the automation of routine tasks, described in Section I, is broadly in line with the implications of our model regarding low-skill automation. For instance, in Acemoglu and Restrepo (2017), we document that industrial robots tend to reduce employment and wages in the most exposed local labor markets, but these effects are significantly more pronounced for low-skill workers in blue-collar jobs. Nevertheless, the evidence also supports the ripple effects that arise in our model when $M < J$, as we find negative effects on workers higher up the skill ladder as well.

V. Long-Run Equilibrium

We have so far assumed that capital is in fixed supply. It is straightforward to embed this model in a dynamic setting in which a representative household maximizes its intertemporal utility (e.g., as in Acemoglu and Restrepo 2016). An immediate implication of that setup is that a long-run, steady-state equilibrium requires that the interest rate, and thus the rental rate of capital, is constant. Here, to economize on space, we study the long-run equilibrium by simply imposing that the rental rate of capital is constant at some level R_ℓ and that the stock of capital adjusts to maintain this level.

More formally, a long-run equilibrium is given by high-skill and low-skill wages, W_H and W_L , respectively, and a level of capital stock, K , such that the final-good producers maximize profits, the three factor markets clear, and $R = R_\ell$. Because we now have an elastic supply of capital, the equivalent of assumption 2 now becomes

ASSUMPTION 2'. The supplies of labor, H, L , and the rental rate of capital, R_ℓ , satisfy $H/L > \underline{\rho}$ and $R_\ell < \bar{R}(H, L)$.

Here, the threshold $\bar{R}(H, L)$ is nonincreasing in H, L and ensures that the endogenous supply of capital satisfies $K > \underline{K}(H, L)$.

The analysis in this case is very similar to that for our short-run equilibrium, with the only difference being that the productivity effect is strengthened because the rental rate is constant at R_ℓ and the stock of capital changes in response to changes in technology and other parameters. Consequently, the productivity effect can now be expressed as

$$\begin{aligned} \sigma \frac{dW_H}{W_H} &= \frac{d\Gamma_H}{\Gamma_H} + \frac{dY}{Y} \Big|_{K\text{fixed}} + s_K \frac{dK}{K}, \\ \sigma \frac{dW_L}{W_L} &= \frac{d\Gamma_L}{\Gamma_L} + \frac{dY}{Y} \Big|_{K\text{fixed}} + s_K \frac{dK}{K}, \end{aligned} \tag{13}$$

where $s_K \in (0, 1)$ is the share of capital in national income, and $(dY/Y)|_{K\text{fixed}}$ denotes the fixed-capital productivity effect characterized in the previous section. The term $s_K(dK/K)$ corresponds to the additional productivity gains due to the induced change in the capital stock.

PROPOSITION 7 (The productivity effect in the long run). Suppose that assumptions 1 and 2' hold. In the long-run equilibrium,

1. Low-skill automation increases the capital stock by

$$\frac{1}{K} \frac{dK}{dI_L} = \frac{1}{1 - s_K} \left(\frac{1}{\Gamma_K} \frac{d\Gamma_K}{dI_L} + \frac{1}{Y} \frac{dY}{dI_L} \Big|_{K\text{fixed}} \right) > 0.$$

2. High-skill automation increases the capital stock by

$$\frac{1}{K} \frac{dK}{dI_H} = \frac{1}{1 - s_K} \left(\frac{1}{\Gamma_K} \frac{d\Gamma_K}{dI_H} + \frac{1}{Y} \frac{dY}{dI_H} \Big|_{K\text{fixed}} \right) > 0.$$

Proof. The results follow by implicitly differentiating the capital-market-clearing condition $KR_k^a = \Gamma_K Y$. QED

Recalling that the increases in the capital stock multiplied by the share of capital in national income, s_k , give the additional productivity effect in this case, we directly obtain the main result of this section in the next proposition.

PROPOSITION 8 (Wages in the long run). Suppose that assumptions 1 and 2' hold. In the long-run equilibrium, both types of automation raise the total wage bill, $W = W_L L + W_H H$. In particular,

$$\frac{dW}{dI_L} = \frac{dY}{dI_L} \Big|_{K \text{ fixed}} > 0$$

and

$$\frac{dW}{dI_H} = \frac{dY}{dI_H} \Big|_{K \text{ fixed}} > 0.$$

Thus, high-skill automation always increases low-skill wages, and low-skill automation always increases high-skill wages. Moreover, when $M < J$ and ε are sufficiently high, both wages increase from either type of automation.

Proof. The appendix presents the derivation for the formulas for dW/dI_L and dW/dI_H .

Because in the long run both types of automation increase the wage bill $W_L L + W_H H$, automation cannot simultaneously reduce both low-skill and high-skill wages. Thus, high-skill (low-skill) automation always increases low-skill (high-skill) wages.

For the second part of the proposition, observe from proposition 6 that when $M < J$ and $\varepsilon \rightarrow \infty$, inequality does not increase, and thus for ε sufficiently high, neither wage could fall as a result of automation. QED

There are several important implications of this proposition. First, it shows that both wages cannot be reduced by automation. This result is closely related to that in Acemoglu and Restrepo (2016), that automation cannot reduce wages in the long run in a model with only one type of labor. In fact, as we have just seen, its proof first establishes that the total wage bill always goes up with automation. Second, more importantly for our focus and differently from Acemoglu and Restrepo (2016), it establishes that the factor directly affected by automation can lose even in the long run. This is because the productivity effects created by automation are shared by both types of labor, whereas the negative displacement effects are borne only by the directly affected factor.

A third implication of this proposition is also noteworthy. Because automation at first pushes up the share of capital in national income (as it will increase the rental rate) and may also increase wage inequality, a natural policy reaction may be to tax capital (e.g., Piketty 2014). Proposition 8, however, implies that this has the potential to be counterproductive. Automation tends to have a more positive effect on wages in the long run than

in the short run because capital accumulates in response to the increased demand for capital. Taxing and discouraging further accumulation of capital would stop these beneficial effects in their tracks.

VI. Equilibrium in Other Cases

In this section, we briefly describe the impact of automation on the allocation of tasks to factors and factor prices. In the remaining cases, we have one of two possible situations.

First, we could have that capital directly competes against low-skill labor only. This will be the case when capital produces tasks in $[0, I_L^*]$ only and $M > I_L^* > 0$; when capital produces tasks in $[0, I_H^*]$ only and $M > I_H^* > J$; or when capital produces tasks in $[0, I_L^*] \cup [0, I_H^*]$ only, $M > I_H^* > J$, and $I_L^* > 0$. In all these cases, automation (an increase in I_L or I_H) takes away only tasks previously assigned to low-skill labor. The analysis is essentially identical to the one developed in Acemoglu and Restrepo (2016), where we show that automation always increases inequality. In addition, displaced low-skill workers then compete against high-skill workers assigned to more complex tasks; thus, ripple effects are still present. As in our analysis in Section IV, when the productivity effect is small (e.g., because capital is scarce), automation not only reduces low-skill wages but may even depress high-skill wages. However, because automation always increases inequality, the long-run effect of automation on high-skill wages is always positive. To summarize, both types of automation have the exact same impact as low-skill automation in Section IV, but there is no equivalent of high-skill automation and thus no room for the simultaneous analysis of low- and high-skill automation.

Alternatively, we could have that capital directly competes against high-skill labor only. This will be the case when capital produces tasks in $[J, I_H^*]$ only and $M < I_H^*$. Here, high-skill automation has the same impact as in our analysis in Section IV, but there is no equivalent automation of low-skill jobs.

VII. Conclusion

This paper developed a task-based model in which high- and low-skill workers compete against machines in the production of different tasks. In contrast to other models in the literature, we have allowed for both low-skill and high-skill automation. The former corresponds to tasks previously performed by low-skill labor being taken over by machines, while the latter involves what is arguably a new phase of automation, in which machines start competing in tasks in which high-skill workers specialize. Our model incorporates both types of automation by departing from the simple structure of comparative advantage adopted by most other models in this genre—so that capital might end up performing two disjoint sets

of tasks, one competing against low-skill labor and the other one against high-skill labor.

After characterizing the different types of equilibria that can arise in this setup, we focus on the cases in which capital does indeed compete against both types of labor. In these cases, improvements in automation technology directly affect either one or the other type of labor. We show that automation always displaces the type of labor it directly affects, depressing its wage. Counteracting this, it also creates a positive productivity effect, pushing up the price of all factors. The net impact of automation on the directly affected factor depends on the balance between the displacement and productivity effects, which are in turn shaped by the gap between the effective cost of producing marginal tasks by labor and that by capital. The simplest characterization of whether the displacement or the productivity effect dominates, then, depends on how large the capital stock is (and thus how high is the rental rate of capital). Importantly, we also establish that both types of automation may create ripple effects, further displacing the type of labor that is not directly affected. The major implication of the ripple effects for us is that automation may reduce the wage of not just the factor it directly affects but of other imperfectly substitutable factors—put differently, either type of automation can depress the real wage of both high-skill and low-skill labor. Though the effects of automation on wages are potentially ambiguous, we also establish that it always has an unambiguous impact on inequality. In particular, low-skill automation always increases wage inequality, whereas high-skill automation always reduces it.

Finally, we extend our model to allow for the adjustment of the stock of capital to changes in technology. In such a “long-run equilibrium,” automation induces further accumulation of capital and thus amplifies the productivity effect. As a result, we show that automation cannot reduce the real wage of both types of labor but may still depress the wage of the directly affected factor.

We view our paper as a contribution to the analysis of how different types of automation technologies, which are arguably permeating our lives even more deeply, will affect the labor market. Many promising research areas are open for future work. Here we mention a few.

1. We took one simple departure from the common, supermodular structure of task-based or assignment models by allowing capital not to have a strict comparative advantage relative to labor. Much richer forms of comparative advantage can be studied beyond the one we have utilized here. The difficulty is to obtain a tight characterization of both the assignment of tasks to factors and factor prices.
2. Automation technology has been assumed to be entirely exogenous, and we have also taken the set of tasks that labor can perform to be given exogenously. In Acemoglu and Restrepo (2016), both of these assumptions are relaxed, which allows for an analysis of how differ-

ent types of technological shocks trigger further technological responses. Extending that type of analysis to our richer framework is an interesting step for future work.

3. We have followed other task-based models in assuming that tasks can be fully unbundled. One of the important effects of new information and communication technologies is to increase the extent to which tasks can be unbundled, and simultaneously studying this process together with the impact of automation technologies appears as a fruitful area for research as well.
4. Our framework does not make any policy recommendations, because we have not modeled various labor market imperfections that might interact with automation, and as a result, the equilibrium in our model is Pareto optimal. Acemoglu and Restrepo (2016) show that when there is a gap between wages and the opportunity cost of labor, automation may reduce welfare and even aggregate output. A similar approach can be developed in the context of our model here, and the welfare implications of different types of automation would depend on the gap between the wage of affected workers and their opportunity cost of labor. For instance, if this gap is greater for high-skilled workers (e.g., because these workers have greater bargaining power), high-skill automation may have more negative welfare consequences than low-skill automation.
5. Last but not least, our analysis has been purely theoretical. How different types of automation technologies affect wages, unemployment, and inequality is an important area for research, and one we are pursuing in ongoing work (Acemoglu and Restrepo 2017).

Appendix

A1. Proof of Proposition 1

The second welfare theorem applies in this competitive economy. Thus, it is sufficient to focus on Pareto optimal allocations, which are those that maximize total output.

We start by showing that, in every potential Pareto optimal allocation, the levels of production of all tasks, $y = (y(i))_{i=0}^1$, are uniquely determined. In any Pareto optimal allocation, y maximizes aggregate output, Y , subject to the constant returns to scale production function (eq. [2]) and the resource constraints ($\int_0^1 l(i) di = L$, $\int_0^1 h(i) di = H$, and $\int_0^1 k(i) di = K$). Uniqueness of the solution, y^* , follows because the production function is strictly quasi-concave.

To see that the production function is strictly quasi-concave, let us express it as $Y = f \circ g$, with $f(x) = x^{\sigma/(\sigma-1)}$ and

$$g(y) = \int_0^1 y(i)^{(\sigma-1)/\sigma} di.$$

When $\sigma > 1$, the function g is an integral of strictly concave functions and is thus itself strictly concave. Because Y is a strictly increasing transform, f , of a strictly

concave function, it is strictly quasi-concave in y . Likewise, when $\sigma < 1$, the function g is an integral of strictly convex functions and is thus itself strictly convex. Because in this case Y is a strictly decreasing transform of a strictly convex function, it is strictly quasi-concave in y . Finally, when $\sigma = 1$, we have that $f(x) = \exp(x)$ and $g(y) = \int_0^1 \ln(y(i)) di$, so the same argument presented above applies. This establishes the uniqueness of y^* in any Pareto optimal allocation and thus in any equilibrium.

Let $x(i) = (l(i), h(i), k(i))$, and $x = (x(i))_{i=0}^1$. Suppose that there are two different Pareto optimal allocations, x and x' . Because y^* is uniquely determined, we must have

$$\begin{aligned} y^*(i) &= l(i)\gamma_L(i) + h(i)\gamma_H(i) + k(i)\gamma_K(i) \\ &= l'(i)\gamma_L(i) + h'(i)\gamma_H(i) + k'(i)\gamma_K(i). \end{aligned}$$

This equality implies that either $x(i) = x'(i)$ or there are several combinations of $\{l(i), h(i), k(i)\}$ that yield the same output. The latter case implies that the final-good producer of task i must be indifferent between using different factor proportions to produce this task, which is not possible, given the assumption that when indifferent between using capital or labor, a firm produces with capital, and when indifferent between using high- or low-skill labor, a firm produces with high-skill labor. Thus, we must have that $x(i) = x'(i)$ for all $i \in [0, 1]$, and there is a unique Pareto optimal allocation x .

Because $\gamma_H(i)$ and $\gamma_L(i)$ are (strictly) increasing, the set of tasks performed by capital can take only the form $[0, I_L^*] \cup [J, I_H^*]$, with $I_L^* \in [0, I_L]$ and $I_H^* \in [J, I_H]$. As in the main text, we adopt the convention that, when $I_L^* = 0$, capital produces only the tasks in $[J, I_H^*]$. Likewise, when $I_H^* = J$, capital produces only the tasks in $[0, I_L^*]$. Also, the market-clearing condition for capital implies that it must perform a range of tasks with positive measure. Thus, $I_L^* = 0$ and $I_H^* = J$ cannot hold simultaneously.

Next, let M denote the smallest index i such that task i can be more cheaply produced with high-skill labor than with low-skill labor. It follows that low-skill labor can produce only the tasks in $[0, M]$ that are not produced by capital, while high-skill labor can produce only tasks in $[M, 1]$ that are not produced by capital.

To conclude the proof, note that low-skill labor must perform some positive measure set of tasks (otherwise, its market-clearing condition would be violated). Thus, we must have that $M > I_L^*$, as stated in the proposition. QED

PROPOSITION A1. There exist a threshold $\underline{\rho}$ and a threshold $\underline{K}(H, L)$ nondecreasing in H, L , such that, for $H/L > \underline{\rho}$ and $K > \underline{K}(H, L)$, the unique equilibrium features $I_H^* = I_H, I_L^* = I_L$, and $M \in (I_L, I_H)$.

Proof. Consider an allocation of tasks to factors where capital performs the tasks in $[0, I_L^*] \cup [J, I_H^*]$; low-skill labor performs the tasks in $(I_L, \min\{J, M\})$; and high-skill labor performs the tasks in $[\min\{J, M\}, J] \cup (I_L, 1]$. The factor prices that support this allocation are given by equation (10).

This allocation minimizes the cost of production and thus corresponds to the unique equilibrium of the model, if and only if the conditions(4) and (5) hold.

We now show that these conditions are satisfied when the supply of factors satisfies the restrictions provided in the proposition.

First, we show that in this allocation, the wage-to-rental-rate ratios, W_H/R and W_L/R , are decreasing in the supply of high-skill, H , and low-skill, L , labor, respectively, and increasing in K . We have that

$$\frac{W_H}{R} = \left(\frac{\Gamma_H}{H}\right)^{1/\sigma} \left(\frac{K}{\Gamma_K}\right)^{1/\sigma}. \tag{A1}$$

If $M > J$, so that Γ_H is constant, then W_H/R is decreasing in H . Otherwise, if $M < J$, an increase in H shifts the left-hand side of equation (12) down in the second panel of figure 2, reducing M . From equation (8), the decline in M also increases Γ_H and reduces Γ_L . From equation (12), we also have

$$\frac{\Gamma_H}{H} = \left(\frac{\gamma_H(M)}{\gamma_L(M)}\right)^\sigma \frac{\Gamma_L}{L},$$

which implies that, as Γ_L declines, Γ_H/H also declines in response to an increase in H in the region where $M < J$. Then, from equation (A1), W_H/R is also decreasing in H . An analogous argument establishes that W_L/R is decreasing in L .

Moreover, equation (A1) shows that the wage-to-rental-rate ratios, W_H/R and W_L/R , are also increasing in the stock of capital K .

It follows that there exists a threshold $\underline{K}(H, L)$ such that the inequalities in equation (5) hold for $K > \underline{K}(H, L)$. Moreover, because the wage-to-rental-rate ratios, W_H/R and W_L/R , are decreasing in the supply of high-skill, H , and low-skill, L , labor, respectively, the threshold $\underline{K}(H, L)$ is nondecreasing.

To establish the existence of $\underline{\rho}$, we show that, in this allocation, the wage ratio W_H/W_L is decreasing in H/L . From equation (10), we have

$$\frac{W_H}{W_L} = \left(\frac{\Gamma_H L}{\Gamma_L H}\right)^{1/\sigma} = \frac{\gamma_H(M)}{\gamma_L(M)}.$$

Thus, when $M > J$ and Γ_L and Γ_H are constant, this expression shows that W_H/W_L is decreasing in H/L . When $M < J$, the increase in H/L shifts the left-hand side of equation (12) down in the second panel of figure 2, once again reducing M , establishing the desired result.

Because W_H/W_L is decreasing in H/L , it follows that there exists a threshold $\underline{\rho}$ such that condition (4) holds. This condition ensures that $M < I_H$, completing the proof of the proposition. (Note that proposition 1 guarantees that $M > I_L$.) QED

A2. *Proof of Proposition 2*

Substituting for factor prices from equation (10) into the ideal-price index condition, equation (11), we obtain

$$\begin{aligned} 1 &= \Gamma_H W_H^{1-\sigma} + \Gamma_L W_L^{1-\sigma} + \Gamma_K R^{1-\sigma} \\ &= \Gamma_H \left(\frac{Y\Gamma_H}{H}\right)^{(1-\sigma)/\sigma} + \Gamma_L \left(\frac{Y\Gamma_L}{L}\right)^{(1-\sigma)/\sigma} + \Gamma_K \left(\frac{Y\Gamma_K}{K}\right)^{(1-\sigma)/\sigma} \\ &= Y^{(1-\sigma)/\sigma} \left(\Gamma_H^{1/\sigma} H^{(\sigma-1)/\sigma} + \Gamma_L^{1/\sigma} L^{(\sigma-1)/\sigma} + \Gamma_K^{1/\sigma} K^{(\sigma-1)/\sigma}\right), \end{aligned}$$

which confirms equation (9). As already noted in the proof of proposition A1, equation (12) follows immediately by combining equations (3) and (10). The fact that equation (12) has a unique solution for M follows by observing that, from assumption 1, the right-hand side is a strictly increasing function of M , while the left-hand side is nonincreasing in M .

The condition $\bar{x} > H/L > \underline{x}$ —which is assumed to hold throughout—ensures that the curves intersect at some interior $M \in (I_L, I_H)$. To see this, note that the condition $H/L > \underline{x}$ guarantees that at I_H we have

$$\frac{\gamma_H(M)}{\gamma_L(M)} > \left(\frac{\Gamma_H L}{\Gamma_L H}\right)^{1/\sigma}.$$

Likewise, the condition $\bar{x} > H/L$ guarantees that at I_L we have

$$\frac{\gamma_H(M)}{\gamma_L(M)} < \left(\frac{\Gamma_H L}{\Gamma_L H}\right)^{1/\sigma}.$$

Thus, the left- and right-hand sides of equation (12) cross at a unique point, M , in the interval (I_L, I_H) , completing the proof. QED

A3. Proof of Proposition 3

When $M > J$, the formulae for the displacement effects follow straightforwardly by differentiating the expressions for Γ_H and Γ_L in equation (8). When $M < J$, we also need to take into account the change in M . Differentiating equation (12), we obtain

$$\frac{d\Gamma_H}{\Gamma_H} - \frac{d\Gamma_L}{\Gamma_L} = \sigma \varepsilon dM. \tag{A2}$$

Moreover, when $M < J$, a change of low-skill automation by dI_L changes the effective labor shares by

$$\begin{aligned} d\Gamma_L &= -\gamma(I_L)^{\sigma-1} dI_L + \gamma_L(M)^{\sigma-1} dM, \\ d\Gamma_H &= -\gamma_H(M)^{\sigma-1} dM. \end{aligned} \tag{A3}$$

Solving equations (A2) and (A3), we obtain the formula provided in the proposition for $d\Gamma_L/dI_L$ and $d\Gamma_H/dI_L$ for the case where $M < J$. Likewise, when $M < J$, a change of high-skill automation by dI_H changes labor shares by

$$\begin{aligned} d\Gamma_H &= -\gamma(I_H)^{\sigma-1} dI_H - \gamma_H(M)^{\sigma-1} dM, \\ d\Gamma_L &= \gamma_L(M)^{\sigma-1} dM. \end{aligned} \tag{A4}$$

Solving equations (A2) and (A4), we obtain the formulae in the proposition for $d\Gamma_L/dI_H$ and $d\Gamma_H/dI_H$ for the case where $M < J$. QED

A4. Proof of Proposition 4

We present the proof for a change in low-skill automation, dI_L . The proof for a change in high-skill automation is entirely analogous and is omitted.

Differentiating the expression for total output, we obtain

$$\begin{aligned} dY &= \frac{1}{\sigma-1} Y^{1/\sigma} \left(\Gamma_H^{(1-\sigma)/\sigma} H^{(\sigma-1)/\sigma} \frac{d\Gamma_H}{dI_L} + \Gamma_L^{(1-\sigma)/\sigma} L^{(\sigma-1)/\sigma} \frac{d\Gamma_L}{dI_L} + \Gamma_K^{(1-\sigma)/\sigma} K^{(\sigma-1)/\sigma} \frac{d\Gamma_K}{dI_L} \right) \\ &= \frac{1}{\sigma-1} Y \left[(Y\Gamma_H)^{(1-\sigma)/\sigma} H^{(\sigma-1)/\sigma} \frac{d\Gamma_H}{dI_L} + (Y\Gamma_L)^{(1-\sigma)/\sigma} L^{(\sigma-1)/\sigma} \frac{d\Gamma_L}{dI_L} + (Y\Gamma_K)^{(1-\sigma)/\sigma} K^{(\sigma-1)/\sigma} \frac{d\Gamma_K}{dI_L} \right] \\ &= \frac{1}{\sigma-1} Y \left(W_H^{1-\sigma} \frac{d\Gamma_H}{dI_L} + W_L^{1-\sigma} \frac{d\Gamma_L}{dI_L} + R^{1-\sigma} \frac{d\Gamma_K}{dI_L} \right). \end{aligned}$$

Plugging the expressions derived in proposition 3 for the terms $d\Gamma_H/dL_L$, $d\Gamma_L/dL_L$, and $d\Gamma_K/dL_L$ yields the desired expressions, completing the proof of proposition 4. QED

A5. Proof of Proposition 5

We present the proof for low-skill automation. The argument for high-skill automation is entirely analogous and is omitted.

Consider the productivity effect $PE_L(K) = (1/Y)(dY/dL_L)$, written as an implicit function of the capital stock, K , and whose expression was derived in proposition 4,

$$PE_L(K) = \frac{1}{\sigma - 1} \left[R^{1-\sigma} - \left(\frac{W_L}{\gamma_L(L_L)} \right)^{1-\sigma} \right].$$

We now show that the productivity effect $PE_L(K)$ increases with the capital stock, K , and that as K becomes smaller, PE_L converges to 0. To prove this, note that an increase in capital reduces the rental rate of capital by

$$\frac{dR}{R} = -\frac{1}{\sigma} (1 - s_K) \frac{dK}{K},$$

while it increases low-skill wages by

$$\frac{dW_L}{W_L} = \frac{1}{\sigma} s_K \frac{dK}{K}.$$

Moreover, as $K \rightarrow 0$, we have that $W_L/R \rightarrow 0$. Thus, as K declines, we necessarily reach a point in which $PE_L(K) = 0$.

To finalize the proof, we start by considering the case $M < J$. Proposition 3 shows that, in this case, both types of labor are affected by a strictly negative displacement effect. Moreover, the proposition shows that

$$\frac{1}{\Gamma_L} \frac{d\Gamma_L}{dL_L} < \frac{1}{\Gamma_H} \frac{d\Gamma_H}{dL_L} < 0.$$

Thus, the negative displacement effect on low-skill labor is stronger than the ripple effect on high-skill labor.

We now make use of these properties of the function $PE_L(K)$ to define the thresholds \bar{K}_L and \underline{K}_L . Because the capital stock does not change the size of the displacement and ripple effects, it follows that there exist two thresholds $\bar{K}_L > \underline{K}_L$ such that

$$PE_L(\underline{K}_L) = -\frac{1}{\Gamma_H} \frac{d\Gamma_H}{dL_L} > 0,$$

$$PE_L(\bar{K}_L) = -\frac{1}{\Gamma_L} \frac{d\Gamma_L}{dL_L} > 0.$$

These thresholds satisfy the properties stated in proposition 5.

Finally, when $M > J$, proposition 3 shows that

$$\frac{1}{\Gamma_H} \frac{d\Gamma_H}{dL_L} = 0.$$

Thus, for high-skill labor we have that $PE_L(K) > -(1/\Gamma_H)(d\Gamma_H/dL) = 0$ for all K and the threshold \underline{K}_L is not defined.

We complete the proof by noting that low-skill automation always increases R because, in this case, the productivity and displacement effects are both strictly positive. QED

A6. Proof of Proposition 6

The proof follows from differentiating the expression $\omega = (\Gamma_H/\Gamma_L)(L/H)$ and then using the formulas derived in proposition 3. QED

A7. Proof of Proposition 7

We present the proof for low-skill automation. The argument for high-skill automation is entirely analogous and is omitted.

The endogenous capital stock satisfies

$$K = \Gamma_K Y R_t^{-\sigma}.$$

Totally differentiating separation, we obtain

$$\begin{aligned} \frac{1}{K} \frac{dK}{dL} &= \frac{1}{\Gamma_K} \frac{d\Gamma_K}{dL} + \frac{1}{Y} \frac{dY}{dL} \Big|_{K\text{fixed}} + \frac{1}{Y} \frac{dY}{dK} \frac{dK}{dL} \\ &= \frac{1}{\Gamma_K} \frac{d\Gamma_K}{dL} + \frac{1}{Y} \frac{dY}{dL} \Big|_{K\text{fixed}} + s_K \frac{1}{K} \frac{dK}{dL} \\ &= \frac{1}{1 - s_K} \left(\frac{1}{\Gamma_K} \frac{d\Gamma_K}{dL} + \frac{1}{Y} \frac{dY}{dL} \Big|_{K\text{fixed}} \right), \end{aligned}$$

which yields the desired result. QED

A8. Proof of Proposition 8

We present the derivation for low-skill automation. The argument for high-skill automation is entirely analogous and is omitted.

Totally differentiating the ideal-price condition, equation (11), we obtain

$$\Gamma_L W_L^{1-\sigma} \frac{1}{W_L} \frac{dW_L}{dL} + \Gamma_H W_H^{1-\sigma} \frac{1}{W_H} \frac{dW_H}{dL} = \frac{1}{\sigma - 1} \left[R^{1-\sigma} - \left(\frac{W_L}{\gamma_L(L)} \right)^{1-\sigma} \right].$$

Let s_L and s_H be the shares of low-skill and high-skill labor in national income, respectively. Note that $\Gamma_L W_L^{1-\sigma} = s_L$ and $\Gamma_H W_H^{1-\sigma} = s_H$. Thus, automation raises the total wage bill by

$$s_L \frac{1}{W_L} \frac{dW_L}{dL} + s_H \frac{1}{W_H} \frac{dW_H}{dL} = \frac{1}{\sigma - 1} \left[R^{1-\sigma} - \left(\frac{W_L}{\gamma_L(L)} \right)^{1-\sigma} \right] > 0.$$

The desired result follows by noting that the left-hand side is equal to $(1/Y)(dW/dL)$ and the right-hand side is equal to $(1/Y)(dY/dL)|_{K\text{fixed}}$.

An alternative derivation uses the accounting identity $R_t K + W = Y$. Differentiating this identity, we obtain

$$R \cdot dK + dW = dY|_{K\text{fixed}} + R \cdot dK,$$

which implies that

$$dW = dY|_{K\text{fixed}},$$

also establishing the same result. QED

References

- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4B, edited by David Card and Orley C. Ashenfelter, 1043–171. Amsterdam: Elsevier Science.
- Acemoglu, Daron, and Pascual Restrepo. 2016. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *A.E.R.*, forthcoming.
- . 2017. "Robots and Jobs: Evidence from US Labor Markets." Working Paper no. 23285 (March), NBER, Cambridge, MA.
- Acemoglu, Daron, and Fabrizio Zilibotti. 2001. "Productivity Differences." *Q.J.E.* 116 (2): 563–606.
- Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *A.E.R.* 103 (5): 1553–97.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2015. "Untangling Trade and Technology: Evidence from Local Labor Markets." *Econ. J.* 125 (584): 621–46.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Q.J.E.* 118 (4): 1279–333.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2016. "The Great Reversal in the Demand for Skill and Cognitive Tasks." *J. Labor Econ.* 34 (1): S199–S247.
- Chui, Michael, James Mayika, and Mehdi Miremadi. 2016. "Four Fundamentals of Workplace Automation." *McKinsey Q.* 2016 (1). <http://www.mckinsey.com/business-functions/business-technology/our-insights/four-fundamentals-of-workplace-automation>.
- Costinot, Arnaud, and Jonathan Vogel. 2010. "Matching and Inequality in the World Economy." *J.P.E.* 118 (4): 747–86.
- Dewhurst, Martin, and Paul Willmott. 2014. "Manager and Machine: The New Leadership Equation." *McKinsey Q.* 2014 (3). <http://www.mckinsey.com/global-themes/leadership/manager-and-machine>.
- Ehrlich, Isaac, and Jinyoung Kim. 2015. "Immigration, Human Capital Formation, and Endogenous Economic Growth." *J. Human Capital* 9(4): 518–63.
- Feng, Andy, and Georg Graetz. 2016. "Rise of the Machines: The Effects of Labor-Saving Innovations on Jobs and Wages." Manuscript, Uppsala Univ.
- Foote, Christopher L., and Richard W. Ryan. 2014. "Labor-Market Polarization Over the Business Cycle." In *NBER Macroeconomics Annual*, no. 29, edited by Jonathan A. Parker and Michael Woodford, 371–413. Chicago: Univ. Chicago Press.
- Goos, Maarten, and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Rev. Econ. and Statis.* 89 (1): 118–33.
- Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *A.E.R.* 104 (8): 2509–26.
- Graetz, Georg, and Guy Michaels. 2015. "Robots at Work." Discussion Paper no. 1335, Centre for Economic Performance, London.

- Gregory, Terry, Anna Salomons, and Ulrich Zierahn. 2016. "Racing With or Against the Machine? Evidence from Europe." Discussion Paper no. 16-053, Centre for European Economic Research (ZEW), Mannheim.
- Hémous, David, and Morten Olsen. 2016. "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality." Manuscript, Dept. Econ., Univ. Zurich.
- Jaimovich, Nir, and Henry E. Siu. 2014. "The Trend is the Cycle: Job Polarization and Jobless Recoveries." Working Paper No. 18334 (March), NBER, Cambridge, MA.
- McNeal, Marguerite. 2015. "Rise of the Machines: The Future Has Lots of Robots, Few Jobs for Humans." *Wired* 2015 (4). <http://www.wired.com/brandlab/2015/04/rise-machines-future-lots-robots-jobs-humans/>.
- Michaels, Guy, Ashwini Natraj, and John Van Reenen. 2014. "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years." *Rev. Econ. and Statis.* 96 (1): 60–77.
- Modestino, Alicia S., Daniel Shoag, and Joshua Ballance. 2016. "Upskilling: Do Employers Demand Greater Skill When Workers are Plentiful?" Manuscript, Northeastern Univ.
- Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Cambridge, MA: Harvard Univ. Press.
- Sattinger, Michael. 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica* 43 (3): 455–68.
- Teulings, Coen N. 1995. "The Wage Distribution in a Model of the Assignment of Skills to Jobs." *J.P.E.* 103 (2): 280–315.
- Zeira, Joseph. 1998. "Workers, Machines, and Economic Growth." *Q.J.E.* 113 (4): 1091–117.