

SHOULD WE FEAR THE ROBOT REVOLUTION? (THE CORRECT ANSWER IS YES)*

Andrew Berg
aberg@imf.org

Ed Buffie
buffie@indiana.edu

Felipe Zanna
fzanna@imf.org

This version September 29, 2017

Abstract

We may be on the cusp of a "second industrial revolution" based on advances in artificial intelligence and robotics. We present a model with the minimum necessary features to analyze the implications for inequality and output. Two assumptions are key: "robot" capital is distinct from traditional capital in its degree of substitutability with human labor; and only capitalists and skilled workers save. We analyze a range of variants that reflect widely different views of how automation may transform the labor market. Our main results are surprisingly robust: automation is good for growth and bad for equality; in the benchmark model real wages fall in the short run and eventually rise, but "eventually" can easily take generations.

JEL Codes: E23, E25, O30, O40

Keywords: Technological change, robots, growth, income distribution.

Preliminary. This paper has been prepared for the Carnegie-Rochester NYU Conference Series

*We would like to thank Robert Chirinko, Anton Korinek, Andy Levin, Chris Papageorgiou, and many IMF colleagues. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the authors and are published to elicit comments and to further debate.

1 Introduction

The factory of the future will have only two employees, a man and a dog. The man will be there to feed the dog. The dog will be there to keep the man from touching the equipment. (Warren Bennis, management consultant)

There is a growing perception that advances in artificial intelligence (AI) and robotics will radically transform the workplace in upcoming decades (Brynjolfsson and McAfee, 2014; Ford, 2015). Robots load, unload, retrieve and send out products with minimal human supervision at Symbolic LLC’s automated distribution centers. AI programs have started working as paralegals, accountants, and teaching assistants, and self-driving vehicles may soon eliminate millions of jobs held by truck, bus, and taxicab drivers. Uber aims to be driverless by 2030. Robots staff more assembly lines each year, kiosks are replacing cashiers at fast-food restaurants,¹ and Watson recently co-authored a song. The list goes on, with each week bringing a new or imminent application of smart machines. According to estimates by Frey and Osborne (2017), Chui et al. (2015), and the World Bank (2016), anticipated advances in automation threaten 45-57 percent of all jobs in the United States. The White House’s Council of Economic Advisors projects that automation will affect 83 percent of jobs paying \$20 an hour or less.²

The sense that we are on the cusp of a robot revolution has sparked a lively debate among economists, journalists, and technophiles about the likely impact of automation on growth and the distribution of income. Broadly speaking, there are two camps with starkly different views of what the future holds. Technology pessimists fear that we are headed toward an economic dystopia of extreme inequality and class conflict: "Without ownership stakes, workers will become serfs working on behalf of robots’ overlords [in] a new form of economic feudalism" (Freeman, 2015). Summers (2016) shares Freeman’s vision (if not his colorful language), predicting that the prime-age employment rate for American males will drop below 25 percent by mid-century in the absence of an aggressive policy response.

Technology optimists do not deny that automation will prove disruptive in the short run. They point out, however, that historically periods of rapid technological change have created more jobs than they have destroyed and have raised wages and per capita income in rough proportion. The AI revolution may be different, but there are good reasons to believe that a resilient, adaptable economy will again vanquish the specter of technological unemployment: income growth raises the demand for labor in sectors that produce non-automatable goods and for workers that perform manual-intensive tasks; higher productivity stimulates investment throughout the economy in cooperating capital inputs; and while automation renders some jobs obsolete, it complements many others, especially jobs that place a premium on creativity, flexibility, and abstract reasoning. Criticizing technology pessimists for missing the big picture, Autor (2014) argues that "journalists and even expert commentators tend to overstate the extent of machine substitution for human labor and ignore the strong complementarities between automation and labor that increase productivity, raise earnings, and augment the demand for labor . . . Focusing only on what is lost misses a central economic mechanism by which automation affects the demand for labor: raising the value of the tasks that workers supply uniquely."

¹McDonald’s has announced that it will introduce digital self-serve ordering stations at all 14,000 of its American restaurants.

²Grace et al. (2017) find that AI researchers on average expect to see AI outperforming humans at translating languages by 2024, driving a truck by 2027, and working in retail by 2031, for example, with considerable variation around these estimates.

1.1 The Theoretical Literature: Searching for the Right Approach

Although economists have extensively analyzed the effects of technological change, the formal theoretical macroeconomic literature that bears on the current debate is quite small. This is important, because the models that tend to shape intuitions fail to capture the the key features of the new technologies. The macroeconomic literature has focused on technological change that increases the supply of some factor in efficiency units. In this formulation, skill-biased technological change exacerbates wage inequality if the elasticity of substitution between low- and high-skill labor exceeds unity.³ Similarly, capital-augmenting technological progress lowers or raises labor’s share in national income depending on whether the elasticity of substitution between capital and labor is above or below unity. But while inequality may worsen, real wages increase across-the-board for all types of labor.⁴ This is a serious limitation. The benign conclusion that everyone gains may correctly characterize the effects of technological in bygone eras, but it is hard to reconcile with the post-1980 stylized facts and the experiences of working-class citizens of many advanced countries. As Acemoglu (2002) observes, “pure technological approaches” fail to explain “how sustained technological change can be associated with *an extended period* of falling wages of low-skill workers and stagnant average wages.” And it fails to address the central idea that the latest technologies differ fundamentally in their ability to allow machines to substitute for human labor in a broad range of tasks.

The shortcomings of existing theory have motivated a new line of research that models robots/automation more realistically as a special type of capital that greatly increases the supply of labor services while reducing the marginal product of human labor that competes for employment in the same production tasks. Sachs and Kotlikoff (2012) frame an overlapping generations (OG) model with a production function in which old skilled workers combine with a CES aggregate of robots and young unskilled workers. When the elasticity of substitution between robots and unskilled labor is sufficiently high, an improvement in robot productivity reduces the demand for unskilled workers. After their wage decreases, unskilled workers — the only group that saves in the economy — cut investment in robots and human capital.⁵ This leads to successive rounds of contraction, with aggregate income and wages declining and each new generation investing less than the generation before. Under weak conditions, welfare of the current young generation and all future generations decreases.⁶

Sachs et al. (2015) analyze OG models that are simpler in some ways and more complex in others. In the one-sector model, firms produce a homogeneous good by operating either a traditional Cobb-Douglas production function with capital and labor or a fully-automated production function linear in robots. The two-sector model adds a non-automatable sector with Cobb-Douglas technology. Capital and robots are putty-putty, not putty-clay, so when robot productivity improves capitalists take apart some factories and convert them overnight into robots. Since capital and labor are complements, the results are qualitatively similar to those in Sachs and Kotlikoff (2012): wage income, saving, and investment decrease and welfare declines for all except the first old generation.⁷ There is an important new twist, however, in the two-sector model. In scenarios where the traditional technology disappears

³For example Acemoglu and Autor’s (2011*a*) ‘canonical model’ is CES two-factor production function of skilled and unskilled labor.

⁴This result assumes a non-nested CES production function containing capital and labor, as in for example Karabarbounis and Neiman (2014) and (implicitly) in Nordhaus (2015).

⁵Under the assumption of log utility, saving is independent of returns to investment in robots and human capital.

⁶The current old generation gains. They receive a pay raise (the skilled wage increases) and earn a higher return on their robots.

⁷In the most relevant case where traditional technology and robots both produce the automatable good, the condition for immiserization in their equation (68) holds for plausible parameter values, including their baseline calibration.

and robots take over the automatable sector, the economy either ascends to a virtuous circle of ongoing endogenous growth or descends into a death spiral of perpetual contraction. Unfortunately, the odds strongly favor the death spiral.⁸

The latest addition to the literature is an ingenious, innovative paper by Acemoglu and Restrepo (2016a). In their model, technological progress proceeds on two fronts: automation and the creation of new more complex tasks that only human labor can perform. A representative Ramsey agent invests in AI capital that serves as a perfect substitute for labor in a subset of potentially automatable tasks. Provided labor supply is increasing in the ratio of the wage to capital income, endogenous directed technological change drives the economy onto a balanced growth path where automation and the creation of complex, new AI-immune tasks advance at the same rate.⁹ Acemoglu and Restrepo show that, *ceteris paribus*, advances in automation then reduce the demand for labor *relative* to capital if the elasticity of substitution between tasks that produce the final good is close to the elasticity of substitution between task-specific intermediate inputs and labor services.¹⁰ The short-/medium-run impact on absolute labor demand is unclear, but in the long run, after the capital stock fully adjusts, the real wage increases and labor's share in national income returns to its original level. When labor is divided into high- and low-skill workers, the same restrictions ensure that the skill premium increases in the short run but not the long run. The results pertain only to relative wages; nothing can be inferred about the short-/medium-term impact on the absolute wage paid to low-skill labor.

Summing up, the debate between the pessimists and the optimists is still unsettled. The papers by Sachs and co-authors make a powerful case for technology pessimism. But the case is built on the OG structure in which all saving and investment is done by wage earners. This is likely to be a sticking point for many people — in the U.S., the majority of regular wage earners live check to check. The Acemoglu and Restrepo paper highlights a new mechanism that strengthens the case for technology optimism. It is hard to ascertain, however, what happens to the absolute level of the real wage on the transition path and how robust the results are to the assumption that capital and labor are perfect substitutes in automatable tasks, to the unusual quasi-labor supply function required for convergence to balanced growth, and to the strong restrictions on technology that ensure automation does not change the skill premium or labor's share in national income in the long run.¹¹

⁸The condition for perpetual contraction in equation (73) of Sachs et al. (2015) holds easily in the baseline calibration of the model in their Section 6.

⁹See Hemous and Olsen (2015) for an alternative model of endogenous technological change. In their model, growth is driven by automation and horizontal innovation that increases the variety of intermediate inputs. (Greater variety of intermediate inputs acts like Hicks-neutral technological progress in a Dixit-Stiglitz gains-from-variety production function.) Automation kicks in when high wages for low-skill labor make it profitable for firms to hire high-skill labor to do research on automation technology. A variety of results are possible depending on the exogenous parameters that control the productivity of horizontal innovation and automation.

¹⁰This condition is not needed if either the production function for tasks is Cobb-Douglas or if the share of intermediate inputs in the task production function is close to zero. The alternative restrictions ensure that factor demands do not stray too far from homotheticity.

¹¹Nordhaus (2015) examines macro data from recent decades to see whether the U.S. economy is undergoing a robot revolution. In his formulation, information technology capital is different. However, he frames his analysis using a CES production function of IT capital and labor, implicitly bundling other capital with labor. This two-factor production function leads him to limit the substitutability between capital and labor to fairly low values. This plus the absence of explicit treatment of both IT-intensive and traditional capital leads him to conclude that real wages should rise strongly with technological progress in the IT sector, counter to the evidence he examines.

1.2 This Paper

This paper analyzes the short and long-run effects of robots on output and its distribution in a family of dynamic general equilibrium models designed to include the minimum necessary features.¹² The models depart from the existing literature by making two critical additions to the standard neoclassical framework. First, they incorporate investment in both robots and traditional capital. In standard production functions, an increase in the supply of labor stimulates investment by raising the productivity of capital. Since the same logic applies when robots increase the effective supply of labor services, the two types of capital should be gross complements. The positive impact of robot labor services on traditional capital accumulation is missing, however, in existing models. In Acemoglu and Restrepo (2016*a*), robots are the only type of capital. Both types of capital are present in Sachs et al. (2015), but not as cooperating inputs; robots reside in a separate production function and do not therefore affect the productivity of non-robot capital. At the core of the story here is that robot capital substitutes for human labor as it complements traditional capital, while increases in traditional capital spur demand for robot capital that attenuates the usual effect of capital accumulation on labor demand.

Second, in keeping with the empirical evidence for the U.S. and other developed countries, we allow for two types of agents in the economy: capitalists who save and invest (and in some variants perform skilled labor), and workers who live check to check. Thus we do not assume that everyone saves, as in the representative agent model of Acemoglu and Restrepo (2016*a*), or that only young wage-earners save, as in the OG models of Sachs and co-authors. We can thus look at inequality along two dimensions: the distribution of income between capitalists and workers, and (in variants with skilled workers) the skilled/unskilled wage differential. Insofar as in reality the number of capitalists is relatively small and that their capital income is generally much greater than wage income, the we interpret the capital share as an index of income inequality. When we allow for skilled as well as unskilled workers in the models, skilled workers are also the capitalists, so that both an increase in the wage premium for skilled workers and the capital labor ratio increase inequality.¹³

While we are confident that the right model features complementary capital inputs and households too poor to save, a number of potentially important design decisions are less clearcut. The basic problem is that nobody knows what the world will look like in say 2035. There is considerable disagreement among and between economists and technology experts about whether automation will (i) destroy jobs for just low-skill labor or labor at all skill levels; (ii) penetrate most or only a small subset of sectors; (iii) reduce the demand for labor in all tasks or decrease it in some and increase it in others. The lack of consensus coupled with the uncertainty inherent in predicting a "future that ain't what it used to be" (Yogi Berra) makes it hard to justify choosing one set of answers over any other.

Our solution to this problem is to start with a benchmark model and then examine the implications of plausible variations that reflect widely different views about how automation may transform the labor market. In Model 1, robots compete against all labor in all tasks. Subsequent extensions

¹²By "robots" we mean the combination of computers, artificial intelligence, big data and the digitalization of information, networks, sensors and servos that are emphasized in the literature on the new machine age (Brynjolfsson and McAfee, 2014, e.g.)

¹³In the model, we assume without loss of generality that there is one capitalist and one (unskilled) worker, and we solve for the evolution through time of the capital share, the wage, and the wage premium. Thus, we can describe the evolution of inequality along both dimensions (capital share and wage premium) relative to the initial equilibrium. If we were interested in comparing the evolution in inequality across runs that involved very different dynamics for the two dimensions of inequality, we could assign relative weights to the two types of agents and create a single index of inequality. In practice, the two dimensions move similarly so this does not seem necessary.

of the model assume that robots (i) compete only for some tasks (Model 2); substitute only for unskilled labor while complementing skilled labor (Model 3); and contribute to production only in one sector — elsewhere, production requires only labor and traditional capital (Model 4). We examine the implications of an increase in the level of robot productivity on the level of output and its distribution, both in the long run and during the transition. Our main results, previewed below, are surprisingly robust. Automation is very good for growth and very bad for equality in all variants, including those reputed to be conducive to technological optimism:

- Real per capita income increases 30-240 percent in the long run. The large positive impact on growth does *not* require dramatic advances in robot technology. Small improvements suffice when robots and human labor are very close substitutes. In runs for this scenario, the direct gains from more investment in more productive robots account for only 6-16 percent of the increase in GDP. The remaining 84-94 percent reflects the strong positive feedback effects between robot and non-robot capital accumulation.
- In the benchmark model, the real wage decreases in the short run under weak conditions. In the long run, however, growth in the non-robot capital stock *raises* the demand for labor and the real wage. Both the long-run increase in the real wage and the depth and duration of the low-wage phase are greater the higher the elasticity of substitution between robots and labor. The intertemporal trade-off for labor is thus sharply defined: more short-run pain for a larger long-run gain.
- Following up on the hint in the last statement, the transition path is difficult for labor. When multiple parameters happen to fall within specific narrow ranges, it takes as little as twelve years for positive real wage growth to materialize. In other scenarios, the low-wage phase lasts 20-50+ years. (The "short run" can consume an entire working life.)
- Although the real wage increases in the long run, labor's share in income decreases most when real output increases most. The bigger the increase in the GDP pie, the less equitable the distribution of the pie.
- In the limiting case of perfect substitutability between robots and labor, the long run never comes. There is a dramatic "singularity": the increase in the level of robot productivity sends the economy on a trajectory that converges to endogenous growth, in which the accumulation of robot and traditional capital continues forever, wages fall and stay below initial levels forever, and the labor share of income converges to zero.
- The distributional outcome is much worse when robots substitute only for low-skill labor (Model 3). While skilled labor enjoys continuous large gains, the wage for low-skill labor decreases in the short/medium run under conditions much weaker than in the benchmark model. Nor is there any assurance that growth eventually raises the low-skill wage. Quite the contrary: there is a strong presumption the real wage decreases more in the long run than in the short run. And the magnitude of the worsening in inequality is horrific. In our base case calibration, the skilled wage increases 56-157 percent in the long run while the wage paid to low-skill labor drops 26-56 percent and the group's share in national income decreases from 31 percent to 8-18 percent.
- The most common arguments for technology optimism do not stand up to scrutiny. Neither the assumption that robots complement labor in some production tasks (Model 2) or that a

non-automatable sector co-exists alongside the automation-vulnerable sector (Model 4) delivers optimistic results. Rather, they tend to underscore an underlying trade-off: variations in which inequality worsens by less also tend to deliver less output growth and lower wage growth.

The rest of the paper is organized into five sections. Section 2 lays out the benchmark model in which robots compete with homogeneous human labor in a single production task. Following this, Sections 3–5 develop models with low- and high-skill labor, automatable and non-automatable sectors, and human-specific production tasks. Section 6 revisits the debate between optimists and pessimists, concluding with an appeal for research aimed at finding policy measures that promote a more equitable distribution of the gains from automation-led growth.

2 Model 1: Robots Do Everything

In the benchmark model, robots substitute for all labor in all tasks. This may seem fanciful, but many AI experts and some economists believe that the future marriage of machine learning and big data will enhance robot pattern recognition to the point where most tasks can be automated.¹⁴ It turns out, moreover, that the focus on one type of labor and one output yields a surprisingly robust pattern of results.

Technology and Factor Demands

Competitive firms operate the linearly homogeneous production function

$$Q = n[a^{1/\sigma_1}K^{(\sigma_1-1)/\sigma_1} + (1-a)^{1/\sigma_1}V^{(\sigma_1-1)/\sigma_1}]^{\sigma_1/(\sigma_1-1)},$$

where

$$V = e^{1/\sigma_2}L^{(\sigma_2-1)/\sigma_2} + (1-e)^{1/\sigma_2}(bZ)^{(\sigma_2-1)/\sigma_2}]^{\sigma_2/(\sigma_2-1)},$$

K , Z , and L denote traditional capital, robots, and labor; the parameter b determines the productivity of robots; σ_2 is the elasticity of substitution between robots and labor; and σ_1 is the elasticity of substitution between traditional capital and the composite input V .¹⁵

Multi-level CES production functions are inelegant and ungainly. This is especially true of the three-tiered functions that appear in subsequent sections. To minimize clutter and facilitate the derivation of analytical results, we bypass the production function and work with the firm’s unit cost function:

$$C[r_k, f(w, r_z/b)] = [ar_k^{1-\sigma_1} + (1-a)f^{1-\sigma_1}]^{1/(1-\sigma_1)}/n, \tag{1}$$

where w is the real wage; r_k and r_z are the real capital and robot rentals; and

$$f = [ew^{1-\sigma_2} + (1-e)(r_z/b)^{1-\sigma_2}]^{1/(1-\sigma_2)}$$

is the sub-cost function dual to the composite input V . When deriving analytical results, we can work with $C[r_k, f(w, r_z/b)]$ and invoke well known formulas that link the derivatives of the cost function

¹⁴AI researchers on average believe that there is a 50% chance of AI outperforming humans in all tasks in 45 years, with substantial variation; Asian researchers on average expect the same point to be reached in 30 years, for example (Grace et al., 2017).

¹⁵ n is the TFP parameter and α and ϵ are CES distribution parameters.

to the substitution elasticities and factor cost shares; it is not necessary to manipulate the actual, cumbersome cost function.

Factor prices adjust to equate demand and supply for each input. Using Shepherd's lemma, the market-clearing conditions may be written as

$$K = C_{r_k} Q, \quad (2)$$

$$Z = C_f f_z Q/b, \quad (3)$$

$$L = C_f f_w Q, \quad (4)$$

where $f_z \equiv \partial f / \partial (r_z/b)$. Labor supply is perfectly inelastic. The capital stocks vary with net investment but are predetermined in the short run.

The Zero-Profit Condition

Perfect competition rules out supranormal profits. Hence price always equals unit cost:

$$1 = C(r_k, w, r_z/b). \quad (5)$$

Saving and Investment

There are two types of agents: workers, who consume all of their income each period, and capitalists. The representative capitalist chooses consumption c , investment in robots I_z , and investment in traditional capital I_k to maximize

$$U = \int_0^\infty \frac{c^{1-1/\tau}}{1-1/\tau} e^{-\rho t}, \quad (6)$$

subject to

$$c + I_z + I_k = r_k K + r_z Z - \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 K - \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 Z, \quad (7)$$

$$\dot{K} = I_k - \delta K, \quad (8)$$

$$\dot{Z} = I_z - \delta Z, \quad (9)$$

where ρ , δ , and τ denote the pure time preference rate, the depreciation rate, and the intertemporal elasticity of substitution, respectively. In the budget constraint (7), the terms $v_k(\bullet)^2 K/2$ and $v_z(\bullet)^2 Z/2$ capture adjustment costs incurred in changing the two capital stocks.

The Maximum Principle furnishes the first-order conditions for an optimum. These may be compressed into two Euler equations for investment

$$\frac{v_k}{K} \dot{I}_k = \left[1 + v_k \left(\frac{I_k}{K} - \delta \right) \right] \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 - r_k, \quad (10)$$

$$\frac{v_z}{Z} \dot{I}_z = \left[1 + v_z \left(\frac{I_z}{Z} - \delta \right) \right] \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 - r_z, \quad (11)$$

which require, as usual, that the real interest rate $(\rho + \dot{c}/c\tau)$ equal the capital and robot rentals, net of depreciation and adjustment costs, at every point in time.

Advances in Robotics

Advances in robotics come in waves. Following an initial jump at $t = 0$, b increases monotonically until it reaches its new steady-state level \bar{b} :

$$\dot{b} = s(\bar{b} - b), \quad s > 0, \quad b_o < b(0) < \bar{b}. \quad (12)$$

2.1 The Short- vs. Long-Run Outcome

Equations (2) - (5) can be solved for w , r_z , r_k , and Q as a function of K , Z , and b . The solution for w is (see Appendix A)

$$\hat{w} = \frac{\sigma_1 - \sigma_2 \theta_K}{\sigma_1 \sigma_2} \alpha_z (\hat{b} + \hat{Z}) + \frac{\theta_K}{\sigma_1} \hat{K}, \quad (13)$$

where a circumflex denotes the percentage change in a variable ($\hat{x} = dx/x$); σ_1 is the elasticity of substitution between capital and labor services (V); σ_2 is the elasticity of substitution between robots and labor; θ_j is the cost share of factor j ($\theta_V = \theta_L + \theta_Z$); and α_i is the cost share of input i ($i = Z, L$) in the production of labor services ($\alpha_L + \alpha_z = 1$).

In the short run, the capital stocks are fixed and

$$\begin{aligned} \hat{w}|_{t=0} &= \frac{\sigma_1 - \sigma_2 \theta_K}{\sigma_1 \sigma_2} \alpha_z \hat{b}, \\ \implies \hat{w}|_{t=0} < 0 &\quad \text{iff} \quad \sigma_2 > \sigma_1 / \theta_K. \end{aligned} \quad (14)$$

The condition in (14) is the same as in DeCanio (2016) and Sachs and Kotlikoff (2012). To understand its logic, suppose that $\sigma_2 = \sigma_1$. We then have a standard non-nested CES production function in which all inputs are gross complements. Hence $\sigma_2 > \sigma_1$ is necessary for automation to reduce w . Note also that as $\theta_K \rightarrow 0$, the production function has only two inputs, implying, again, that labor and robots must be gross complements. Hence the threshold value of σ_2 — the value at which L and Z become gross substitutes — varies inversely with θ_K . Putting this together with $\sigma_2 > \sigma_1$ suggests the condition $\sigma_2 > \sigma_1 / \theta_K$.

Production theory asserts that normally inputs are gross complements (Rader, 1968). The normal rule may not apply, however, to robots and human labor. For $\theta_K = .33$ -.40, the condition in (14) requires σ_2 to be 2.5-3 times larger than σ_1 . Empirical estimates of σ_1 range from .4 to 1.2, so the borderline value of σ_2 could be on the order of 3.6. This is very high for an elasticity of substitution between two primary inputs. But a value of 3.6, or even a much larger number, is consistent with the view that robots are a special type of capital that greatly increases the total supply of labor services, e.g.: "Counting both humans and machines, the world's labor force will be able to do more work than ever before. But this abundance of labor — both those made of cells and those made of bits — could create a glut of labor. The machines may render humans as redundant as so many vintage washing machines" (The Atlantic, 2016).

When robot technology improves and wages decline, capitalists rake in more profits and perceive higher returns on robot and non-robot capital. They react prosaically by increasing investment in

both types of capital.¹⁶ Across steady states where $r_z = r_k = \rho + \delta$, equations (2)-(5) give

$$\hat{K} = \hat{Q} = \sigma_2 \frac{\alpha_z}{\alpha_L} \hat{b} > 0, \quad (15)$$

$$\hat{Z} = (\sigma_2/\alpha_L - 1)\hat{b}. \quad (16)$$

Moreover, the positive effect on labor demand of growth in the non-robot capital stock eventually dominates the labor-displacing effect of greater utilization of more productive robots. It follows from the zero-profit condition

$$1 = C[\rho + \delta, w, (\rho + \delta)/b] \quad (5')$$

that the real wage *increases* in the long run

$$\hat{w} = \frac{\theta_Z}{\theta_L} \hat{b} > 0. \quad (17)$$

Figure 1 supplies the intuition for this result. At the initial steady state, the capital rental r_k equals the time preference rate plus the depreciation rate. The increases in b and Z reduce the marginal product of labor (MPL) and raise the return to traditional capital. At $K = K_1$, the wage has recovered to its original level. But since r_z/b has decreased, we know from the zero-profit condition that r_k still exceeds $\rho + \delta$.¹⁷ Capitalists keep investing therefore until K rises to \bar{K} . The additional capital accumulation increases labor demand and bids up the market-clearing wage to \bar{w} .

Two other results deserve attention. First, there are striking tradeoffs between short-run pain and long-run gain for labor and, at the macro level, between inequality and growth.¹⁸ The larger is σ_2 , the more the real wage decreases in the short run and, as shown in the next section, the longer the period of wage stagnation. But larger values of σ_2 also imply larger increases in the total supply of labor services, greater productivity gains, and higher returns to investment in robots and traditional capital. Consequently, both capital stocks, real income, and the real wage increase more in the long run. The positive effect of σ_2 on K , Z , and Q appears directly in (15) and (16). In the solution for the real wage it shows up indirectly as an increase in θ_Z/θ_L for non-infinitesimal changes in b . Importantly, both higher θ_Z and lower θ_L contribute to the increase in w :

$$\hat{\theta}_Z = (\sigma_2 - 1)\hat{b} > 0, \quad (18)$$

$$\hat{\theta}_L = \frac{\alpha_z}{\alpha_L}(1 - \sigma_2)\hat{b} < 0. \quad (19)$$

There is not much doubt that σ_2 exceeds one. In the empirically relevant range, therefore, $\hat{\theta}_Z/\hat{b} > 0$ and $\hat{\theta}_L/\hat{b} < 0$. Labor's share in income declines, and, paradoxically, it declines most when σ_2 is large and the real wage increases most in absolute terms. A rising tide lifts all boats in the long run, but it never lifts labor's boat as much as other boats. This result proves depressingly robust across all the models we investigate.


¹⁶This remark pertains to the medium run. In the short run, higher investment in robot capital may come at the expense of investment in non-robot capital. See the analysis of the transition path in the next section.

The increase in investment is in sharp contrast to the results of Sachs and co-authors, where all saving is done by unskilled (young generation) workers.

¹⁷For $K = K_1$ and $bZ = \bar{b}\bar{Z}$, r_z/b is lower both because $\bar{b} > b_0$ and because $r_z < \rho + \delta$. (On an optimal path, of course, $Z < \bar{Z}$ when $K = K_1$ and $r_z > \rho + \delta$. Figure 1 is a heuristic device.)

¹⁸We quantify the tradeoffs in Section 3.3, assuming Model 2 is more likely to command agreement than Model 1.

Figure 1. Model 1: Real wages in the long run



The other result pertains to methodology. Returning to a point emphasized in the introduction, inclusion of traditional capital complementary to robot capital is critical to understanding the general equilibrium effects of technological advances in robotics. Obviously, non-robot capital accumulation takes all of the credit for the long-run increase in the real wage. In addition, it is the central, driving force behind the large increase in per capita income. Table 1 reports long-run solutions for output, robot capital, and the real wage when the non-robot capital stock is constant. The real wage decreases 9-11 percent, and the increase in income is only 6-25 percent as large as in the full model. Most of the action in the full model revolves around the strong positive feedback effects between robot and non-robot capital accumulation. This is true even of robot investment. Only 8-37 percent of the increase in robot capital is a response to better technology ($b \uparrow$). The rest reflects the impact of non-robot capital accumulation raising the return to robot investment.

2.2 The Transition Path

The real wage decreases in the short run and increases in the long run. What happens in between determines whether Model 1 sides with technology optimism or technology pessimism. If the long run is too far away — if it takes 30+ years for the real wage to increase — we definitely have a problem.

There are a variety of ways to approximate the stable manifold. Given the substantial nonlinearities present in the model, we judged the method in Novales et al. (2008) to offer the best tradeoff between solution speed and minimization of approximation error. The method derives stability conditions from a linear approximation around the steady state, but incorporates the nonlinear structure of the model

Table 1. Model 1: Long-run outcome with constant vs endogenous non-robot capital

Scenario	Model 1 with K constant				Full Model			
	GDP	w	Z	K	GDP	w	Z	K
$\sigma_2 = 2.5, \bar{b} = 1.5$	21	-11	288	0	71	24	788	71
$\sigma_2 = 5, \bar{b} = .75$	12	-9	215	0	47	8	646	47
$\sigma_2 = 10, \bar{b} = .65$	12	-11	240	0	202	12	3099	202

Notes: The entries report the percentage change in the variable.

when tracking the transition path.

To calibrate the model, we assigned the following parameter values in the base case:

$$\delta = .05, \quad \tau = .50, \quad \Omega_K = 1, \quad \rho = .06, \quad \theta_K = .35, \quad \theta_Z = .04, \quad s = .15$$

$$b_o = .50, \quad b(0) = b_o + .1(\bar{b} - b_o), \quad \sigma_1 = .50, \quad \Omega_Z = .50, \quad \sigma_2 = 2.5-20, \infty.$$

The depreciation rate, the intertemporal elasticity of substitution, and the q-elasticity of non-robot investment (Ω_K) all take ordinary values, while the numbers for the time preference rate and the total cost share for capital ($\theta_K + \theta_Z = .39$) match the long-run return on stocks and the data on factor shares in the U.S.¹⁹ With respect to the other choices:

- Improvements in robot productivity arrive continuously, with b covering 90 percent of the ground to its new steady-state level in the first fifteen years.
- As noted earlier, empirical estimates suggest that σ_1 lies somewhere between .4 and 1.2. Our choice of .50 is in line with newer, micro-based estimates²⁰
- Lack of empirical information forced us to rely on semi-educated guesses for the last two parameters. We decided to set the q-elasticity for robot investment (Ω_Z) at .50 on the grounds that installing a new type of capital that embodies new, highly innovative technology is likely to involve greater adjustment costs than construction of another plant identical to plants the firm has built and operated many times before.²¹
- To the best of our knowledge, there are no econometric estimates of σ_2 . Technology experts concur that substitution between robots and human labor (in tasks where substitution is possible) is much easier than substitution between most primary inputs. But it is difficult to translate

¹⁹The income share for labor $\theta_L = 1 - \theta_K - \theta_Z = .61$ equals labor's share in the U.S. in 2013 and the income share for robots implies that $Z/(Z+K)$ is close to the share of information capital in the total capital stock (10.1 percent) reported in Nordhaus (2015). The value assigned to Ω_K pins down the adjustment cost parameter v_k . [The first-order condition for investment is $1 + v_k(I_k/K - \delta) = q$, where $q \equiv \phi_2/\phi_1$ and ϕ_1 and ϕ_2 are multipliers attached to the constraints in (7) and (8). q is Tobin's q , the ratio of the demand price of capital to its supply price. Evaluated at a steady state, $v_k = 1/(\Omega_K \delta)$, where $\Omega_K \equiv \hat{I}_K/\hat{q}$.]

²⁰See Chirinko (2008), Chirinko and Mallick (2017), Raval (2011), Klump et al. (2007), and Oberfield and Raval (2014).

²¹This assumption is not critical. The results change little when $\Omega_k = \Omega_z$ and the q-elasticity varies from .5 to 5. See Table 3 in Section 3.3.

”much easier” into a number for σ_2 . The estimates in Acemoglu and Restrepo (2016*b*) provide some guidance. Their finding that one robot directly eliminates 10.6 jobs suggests that σ_2 might be quite large. We are reluctant, however, to pin too much on one estimate.²² Seeking robust results, we carried out runs for moderately low ($\sigma_2 = 2.5$), moderately high ($\sigma_2 = 5$), very high ($\sigma_2 = 10, 20$), and perfect substitution ($\sigma_2 = \infty$).²³

Scenario 1: Moderately Low Substitution ($\sigma_2 = 2.5$) and Large Increases in Robot Productivity

We start with a relatively optimistic scenario. In Figure 2, σ_2 is a moderately high 2.5 and successive waves of innovation increase the productivity of robots 200 percent. Robot investment surges, and as b and Z increase the market for human labor becomes very soft. To make matters worse, the high return on robots attracts investment away from capital complementary to labor. The non-robot capital stock decreases 4 percent in the first five years and does not surpass its start value until year fourteen. In the short/medium run, everything conspires to depress the demand for labor and the real wage. Since real output rises apace, the income share of capital increases quickly, rising from 39 percent to 43 percent at year ten.

Fortunately, the story does not end here. The diversion of investment from traditional capital into robots is temporary. (K increases across steady states.) Further out on the transition path, firms start investing more in both types of capital and positive wage growth emerges. But it takes time for the positive effect on labor demand of traditional capital accumulation to outweigh the labor-displacing effect of the robot revolution. The path for w does not leave the fourth quadrant until year twenty. In the decades that follow, the real wage increases steadily but never catches up with per capita income growth. The distribution of income continues to worsen; in the long run, labor’s share in national income drops to 44 percent.

The robot revolution produces large gains in per capita income. Inequality worsens, but the real wage increases in the long run. The long-run gains in income and the real wage are greater the larger the elasticity of substitution between robots and human labor; disturbingly, however, inequality worsens more and a full generation may pass before real wage increases show up on the transition path.


Scenario 2: A Troubling Singularity ($\sigma_2 = \infty$)

The above set of results does not carry over to the limiting case where robots substitute perfectly for human labor. One important result disappears and the others become more extreme. A discontinuity separates the cases of high and perfect substitution. In effect, in the perfect substitution case the low-wage/increasing-inequality transition becomes permanent. With perfect substitution, the build-up of robot and traditional capital never brings about a rise in the marginal product of (human) labor, and by the same token the scarcity of human labor never brings about an end to what is otherwise transitional growth.

²²The estimate that one robot directly eliminates 10.6 jobs is not a pure empirical estimate. It depends on the regression coefficient in the employment equation and the values assigned to the inverse of the Frisch elasticity of labor supply and the inverse elasticity of supply of robots.

²³The case of perfect substitution is extreme. However, it claims a fair number of supporters, at least with reference to some time in the (not too distant?) future, provides a useful benchmark, and has received a good deal of attention in the economics literature. The models of Sachs et al. (2015) and Acemoglu and Restrepo (2016*a*), for example, assume perfect substitution in a subset of production tasks.

Figure 2. Model 1: Transition Path for Scenario 1



Notes: Transition path when $\sigma_2 = 2.5$ and b increases from .5 to 1.5 in the long run. The paths for the real wage (w), GDP, and the non-robot capital stock (K) show the percentages deviations from initial values.

The good news is that income rises forever in an economy that would otherwise stop growing. As labor's share in income approaches zero, the production function exhibits constant returns to scale in the two capital stocks and the economy enters the pleasant world of sustained endogenous growth. In Appendix B we show that for $\sigma_1 = 1$

$$g \left\{ 1 + \tau v_k \left[\frac{g}{2\tau} (2 - \tau) + \rho \right] \right\} = \tau (r_k - \rho - \delta), \quad (20)$$

$$g \left\{ 1 + \tau v_z \left[\frac{g}{2\tau} (2 - \tau) + \rho \right] \right\} = \tau (r_z - \rho - \delta), \quad (21)$$

where g is the equilibrium growth rate,

$$r_k = n\theta_K (bz)^{1-\theta_K},$$

$$r_z = \frac{nb(1-\theta_K)}{(bz)^{\theta_K}},$$

and $z \equiv Z/K$. The presence of adjustment costs complicates matters.²⁴ Suppose, purely for expositional purposes, that $v_k = v_z = 0$. We then get the usual solution

$$g = \tau (r_k - \rho - \delta), \quad (22)$$

with

$$r_k = r_z = n\theta_K [b(1-\theta_K)/\theta_K]^{1-\theta_K}. \quad (23)$$

²⁴Either $\tau < 2$ or $\rho > g/2$ are sufficient to ensure sustained growth. In general, sustained growth is possible as long as the q-elasticities of investment are not very close to zero. (Recall that v_z and v_k are tied to the q-elasticities of investment spending via $v_z = 1/\Omega_z\delta$ and $v_k = 1/\Omega_k\delta$. See footnote 19).

Labor’s share in income goes to zero asymptotically. This would not be troubling *if* the real wage increased forever (at a slower rate than GDP) or plateaued at a very high level. There is nothing wrong with a world in which the machines do all the work and all the humans become fabulously rich.²⁵ But not all the humans become rich. In fact, workers suffer a *permanent decrease* in income. The solution for r_z in (23) and the arbitrage condition


$$w = r_z/b$$

give

$$\hat{w}|_{Long\ Run} = -\theta_K \hat{b}. \tag{24}$$

The same qualitative results obtain in the full model that incorporates adjustment costs. Figure 3 plots the transition paths for $\sigma_2 = \infty$ vs. $\sigma_2 = 5$ when $\sigma_1 = 1$ and all other parameters take their base case values.²⁶ Perfect substitution of robot for human labor delivers perpetual growth. But the rich become richer and the poor poorer with every passing year. In the long run, the real wage decreases 13.4 percent while capitalists’ income rises without limit.

Figure 3. Model 1: An Economic Singularity



Notes: Transition paths for the growth rate and the real wage for Model 1 when robots and labor are perfect vs imperfect substitutes.

These results differ sharply from those in Sachs et al. (2015) and Acemoglu and Restrepo (2016a), which assume perfect substitution either in a subset of production tasks or in a separate production function. In Sachs et al.’s two-sector model, there is a strong presumption the economy will fall

²⁵Income envy and class distinctions may also fade away. When low-skill workers earn enough, they will start to save and derive some of their income from capital. Inequality should decrease steadily as everyone joins the “ownership economy.” If capitalists and today’s workers share identical Stone-Geary utility functions, the distribution of income would become perfectly equitable in the long run.

²⁶ b increases from $b_o = .1803$ to $\bar{b} = .30$. When $b_o = .1803$, the return on robots equals the time preference rate for $\sigma_2 = \infty$ (i.e., $b_o = .1803$ is the inflection point) . Hence the initial equilibrium growth rate is zero for both $\sigma_2 = 5$ and $\sigma_2 = \infty$ — the starting point is the same in the two runs.

into the death spiral of never-ending negative growth. The culprit, as usual, is the assumption in the OG framework that automation reduces saving and investment by reducing wage income. The model in Acemoglu and Restrepo (2016a) differs in many ways from ours. Crucially, exogenous technological advances in automation do not affect the growth rate or the return on capital in the long run. Because the capital rental does not change, the results for the long run are the same as when robots and labor are imperfect substitutes in our model: the real wage increases, but labor's share in national income declines. In Acemoglu and Restrepo's preferred ("complete") model, where endogenous directed technological change creates countervailing forces that strengthen labor demand, the outcome is less clear. Labor's income share may decrease or recover to its original value depending on how automation technology improves.

3 Model 2: Robots Cannot Do Everything

Model 1 is not entirely pessimistic. Per capita income rises and strong crowding in of the non-robot capital stock guarantees that overall labor demand and the real wage increase in the long run (assuming $\sigma_2 \neq \infty$). Nevertheless, the model is contrary to the spirit of Autor and other technology optimists. They would presumably take issue specifically with the assumption that robots substitute for human labor in all tasks, arguing that it is unrealistic and strongly biases the results toward pessimism over a lengthy medium run. The right model recognizes that automation reduces the demand for labor in some tasks but increases it in many others. Model 1 needs to be replaced by something more evenhanded.

To accommodate this view, we shift some labor from the CES nest with robots into a CES nest with traditional capital. The production function is $Q = F[H(K, L_2), V(L_1, Z)]$, where L_2 is labor that performs some task complementary to the task performed by L_1 and Z in $V(\bullet)$. $L = L_1 + L_2$, L_1 and L_2 both get paid w , and the firm's unit cost function is $C[h(r_k, w), f(w, r_z/b)]$. In what follows, σ_3 , θ_H , χ_k , and χ_L denote the elasticity of substitution between K and L_2 , the cost share of the composite input H ($\theta_H = \theta_K + \theta_{L_2}$), and the cost shares of capital and labor in production of H ($\chi_k + \chi_L = 1$). We can interpret this model as an optimistic one in which labor moves freely between the two tasks; there is no skill differential and no education or other investment is necessary. Model 3 will present a polar case in which labor is immobile across the two tasks.

3.1 The Direct Effect of Automation ($bZ \uparrow$) on Labor Demand

The new versions of equations (2) - (5) are

$$K = C_r Q, \tag{25}$$

$$Z = C_f f_z Q / b, \tag{26}$$

$$L = (C_f f_w + C_h h_w) Q, \tag{27}$$

$$1 = C(r_k, w, r_z/b), \tag{28}$$

which can again be solved for w , r_z , r_k , and Q as a function of K , Z , and b . The counterpart of equation (13) is

$$\hat{w} = [(\sigma_1 - \sigma_2 \theta_H) \chi_k + \sigma_3 \theta_V (\chi_L + L_2/L_1)] \frac{\alpha_z}{N} (\hat{b} + \hat{Z}) \tag{29}$$

$$+ [\sigma_2 \theta_K (1 + \alpha_L L_2/L_1) + (\sigma_1 - \sigma_3 \theta_V) \chi_k \alpha_z L_2/L_1] \frac{\hat{K}}{N}, \tag{30}$$

where

$$N \equiv \sigma_2(\sigma_1\chi_k\alpha_z + \sigma_3\chi_L) + (\sigma_1\alpha_z + \sigma_2\alpha_L)\sigma_3L_2/L_1.$$

An increase in bZ now reduces the wage when

$$\sigma_2 > \frac{\sigma_1\chi_k + \sigma_3\theta_V(\chi_L + L_2/L_1)}{\theta_K}. \quad (31)$$

There is no support for optimism here. Robots eliminate jobs in the $V(L_1, Z)$ nest and increase productivity of labor in the $H(K, L_2)$ nest. The adjustment in the wage required to absorb displaced L_1 workers in L_2 -type tasks depends on the slope of the MPL_2 schedule and how much the schedule shifts upward. A lot rides therefore on the magnitude of σ_3 relative to σ_1 . The wage is *more* likely to decrease than in the case where robots substitute for all labor when

$$\frac{\sigma_1\chi_k + \sigma_3\theta_V(\chi_L + L_2/L_1)}{\theta_K} < \frac{\sigma_1}{\theta_K},$$

which reduces to

$$\sigma_3 < \frac{\sigma_1}{\theta_V + \theta_H/\alpha_L}. \quad (32)$$

Since $\theta_V + \theta_H = 1$ and $\alpha_L \approx .95$ at the initial equilibrium, this condition holds when σ_3 is slightly below σ_1 . Reversing the inequality makes Model 2 more optimistic than Model 1, but does not weaken the presumption that automation reduces total labor demand. *Ceteris paribus*, robots increase the wage only if

$$\sigma_3 > \sigma_3^* = \frac{\chi_k(\sigma_2\theta_H - \sigma_1)}{\underbrace{\chi_L(\theta_V + \theta_H/\alpha_L)}_{\approx 1}}. \quad (33)$$

For plausible parameter values, σ_3^* is much greater than σ_1 and considerably above high-end estimates of the elasticity of substitution between capital and labor (Table 2). At present, the neutral assumption is $\sigma_3 \approx \sigma_1$. The case for optimism over the short/medium term is no greater in Model 2 than in Model 1.

Table 2. Model 2: Value of σ_3^*
 $(\sigma_3 > \sigma_3^* \implies$ rising wages in short run)

	$\sigma_1 = .5$			$\sigma_1 = 1$		
	$\sigma_2 = 2.5$	$\sigma_2 = 5$	$\sigma_2 = 10$	$\sigma_2 = 2.5$	$\sigma_2 = 5$	$\sigma_2 = 10$
$L_2/L_1 = .5$	1.44	3.70	8.23	NA	2.89	7.41
$L_2/L_1 = 1$	1.20	2.93	6.39	NA	2.40	5.87
$L_2/L_1 = 2$	1.04	2.46	5.29	NA	2.08	4.92

Note: The total cost share of labor is 61 percent at the initial equilibrium. All other parameters take the values specified in section 2.2. "NA" indicates that the real wage always rises (by a minuscule amount) when $\sigma_1 = 1$ and $|\sigma_2 = 2.5$.

3.2 The Long Run

The long run is also not kind to Model 2. Compared to model 1, the effects are similar but attenuated. As with model 1, inequality worsens, the wage falls at first, and output and the real wage eventually grow. Consider first the impact on GDP. When robots substitute for labor in some but not all tasks, the productivity gain associated with the innovation in b is smaller. There is a straight line from this to a smaller increase in the return to non-robot investment, less crowding in of non-robot capital, and a smaller multiplier effect on real output. In the benchmark case where $\sigma_3 = \sigma_1$,

$$\hat{K} = \hat{Q} = \left[\sigma_2 \left(1 + \frac{\theta_{L2}}{\theta_V} \right) + \sigma_1 \theta_{L2} \frac{\alpha_z}{\theta_{L1}} \right] \frac{L_1 \theta_Z}{L \theta_L} \hat{b}, \quad (34)$$

where $\theta_L \equiv \theta_{L1} + \theta_{L2}$. K and Q increase less than in Model 1 provided

$$\sigma_2 > \sigma_1 \frac{\alpha_z \theta_{L1}}{\theta_L + \alpha_z \theta_{L1}}, \quad (35)$$

which holds even when σ_2 is much smaller than σ_1 .

The wage is pinned down by the zero-profit condition (28). As before,

$$\hat{w} = \frac{\theta_Z}{\theta_L} \hat{b} > 0. \quad (36)$$

For small changes, the solution is the same as in Model 1. But for large (i.e., non-infinitesimal) changes, w increases less because θ_Z/θ_L increases less. The explanation for this result is straightforward. When the inequality sign in (33) is reversed, job losses directly attributable to automation are smaller than in Model 1. But the positive impact on labor demand of traditional capital accumulation is also smaller. The jobs lost from the weakening of this positive effect exceed (ex ante) jobs recovered from the weakening of the adverse automation effect. Consequently, regardless of what happens in the short run, the real wage always increases less in the long run (Table 3).

Table 3. Model 1 vs. Model 2 in the long run

	GDP	w	K	θ_L	Z	I_Z/GDP
Model 1	71	24	71	44	788	9
Model 2						
$\sigma_3 = .5$	45	19	45	58	446	7
$\sigma_3 = .25$	43	19	40	59	436	7
$\sigma_3 = 1$	50	19	56	56	465	7

Note: The entries for GDP , w , K , and Z are the percentage change in the variables. In all cases $\sigma_1 = .5, \sigma_2 = 2.5$, and b increases from 0.5 to 2.

Although real output and the real wage increase less in Model 2, those who care more about inequality than growth may not agree that the long-run equilibrium in Model 2 is inferior to the long-run equilibrium in Model 1. For small changes, the wage increases the same amount in Model 2 as in Model 1 while output rises less. Thus, at the new steady state, labor's share in national income

is higher in Model 2 than in Model 1. But inequality still worsens. From (34 and (36),

$$\hat{\theta}_L < 0 \quad \text{iff} \quad \sigma_2 > \frac{1 + (1 - \sigma_1 \alpha_z) \theta_{L2} / \theta_{L1}}{1 + \theta_{L2} / \theta_V}. \quad (37)$$

The term on the right side is smaller than unity for $\sigma_1 > 1$. When $\sigma_1 < 1$, θ_L decreases under the weak condition

$$\sigma_2 > \frac{\theta_L}{\theta_L - \theta_{L2} \alpha_z}. \quad (38)$$

$\theta_{L2} \alpha_z$ is on the order of .05, so values of σ_2 slightly above unity satisfy (38).

3.3 The Transition Path

Income per capita and the real wage increase less in the long run in Model 2 than in Model 1. In scenarios where $\sigma_3 \geq \sigma_1$, however, the wage decreases less in Model 2 in the short run. If the difference is large and persists for decades, the argument that labor fares better when robots and humans compete in fewer production tasks has merit, even though the economy ends up at an inferior steady-state equilibrium. We evaluate the argument in two scenarios that differ in the degree of substitution between robots and human labor and the size of the increase in robot productivity. First, we compare Models 1 and 2 in Scenario 1 (low σ_2 , large increase in b). Following this, we analyze a new scenario (Scenario 3) with small increases in robot productivity and high but not perfect substitution between robots and human workers.

Scenario 1 Redux: Moderately Low Substitution ($\sigma_2 = 2.5$) and Large Increases in Robot Productivity

Figure 4 compares impulse responses in Models 1 and 2 in three cases which vary with respect to σ_3/σ_1 .²⁷ Model 2 does not score well in any of the runs. In the two runs where $\sigma_3 \leq \sigma_1$, real output is slightly higher in Model 2 than in Model 1 for the first 50 years because the transfer of labor from the $V(\bullet)$ nest to the $H(\bullet)$ nest increases the marginal product of both capital stocks, spurring faster capital accumulation. ($L_1 \downarrow$ increases MPZ and $L_2 \uparrow$ increases MPK). On the other hand, the low-wage phase lasts 4-8 years longer and inflicts deeper wage cuts, especially in the run for $\sigma_3 < \sigma_1$ (Figure 4b). The most favorable case (4c), where $\sigma_3 > \sigma_1$, also yields mixed results. The ranking of the paths for w , as well as the overall ranking, are largely a matter of taste. (The long-run increase in w is 71 percent in Model 1 vs. 50 percent in Model 2.)


Taking stock, what is the general import of these results? To be clear, the argument for optimism, based on complementarity of labor and robots in some production tasks, is not necessarily wrong. It is shaky, however, and certainly does not work as advertised. The argument gets the long-run outcome completely wrong. In the short/medium run, it cuts the right way when $\sigma_3 > \sigma_1$ and the wrong way when $\sigma_3 < \sigma_1$; but no empirical estimates justify $\sigma_3 > \sigma_1$, the condition for a better result (w decreases less), let alone $\sigma_3 > \sigma_3^*$, the stringent condition for a strictly positive result ($w \uparrow$ at $t = 0$).

Short- vs. Long-Run Tradeoffs


We demonstrated analytically in Model 1 that there are sharp tradeoffs between inequality, growth, real wage cuts in the short run, and real wage increases in the long run. Figures 5-7 quantify the

²⁷To calibrate the model, we set $\theta_{L2} = .30$ and $\theta_{L1} = .31$. (The total cost share of labor is the same as in Model 1.) All other parameters take the same values as in the baseline calibration.


Figure 4. Model 1 vs Model 2 for various values of σ_3



(a) $\sigma_3 = .5$



(b) $\sigma_3 = .25$




(c) $\sigma_3 = 1$

Note: In all cases $\sigma_1 = .5$, $\sigma_2 = 2.5$, and b increases from .5 to 2.

tradeoffs in Model 2. As σ_2 increases from 2 to 10 in Figure 5, the long-run increases in w and GDP rise from 6.7 percent and 12.9 percent to 52.8 percent and 151.3 percent, respectively. The much larger long-run gains for $\sigma_2 = 10$ come at a steep price, however: the real wage cut at the trough of the transition path jumps from -.8 percent to -21.8 percent.

Figure 5. Model 2 in the short and long run for various values of σ_2



Notes: The short-run solution for the first panel shows the largest decrease in w on the transition path (which generally occurs between years 10 and 15). The right-hand panel shows the long-run solutions for GDP and the labor share. In all cases b increases from .5 to 1 and all other parameters take their base-case calibration values (see section 2.2).

Figure 6 provide additional information on the tradeoff for labor. Reducing the increase in robot productivity from 200 percent to 50 percent greatly reduces investment in robots and the long-run increases in real output and the real wage (9.4 percent vs. 50 percent for GDP and 3.7 percent vs. 19.8 percent for w). By way of compensation, the low-wage phase ends seven years earlier and real wage losses are 2-5 percentage points smaller in the short/medium run ($t = 0-20$). Increasing σ_2 from 2.5 to 5 makes an even bigger difference. Real output and the real wage increase much more in the long run (204 percent vs. 50 percent for GDP and 70.7 percent vs. 19.8 percent for w), but workers suffer huge, brutal wage cuts during a low-wage phase that persists for almost two generations. The problem, clearly, is the composition of investment; labor gets hit front and back as firms finance rapid increases in robot capital by cutting investment in traditional capital.²⁸


The robot productivity parameter b and the substitution elasticity σ_2 are the only parameters that significantly affect the depth and duration of the low-wage phase. Individually, none of the other parameters have much impact. Table 3 presents results for runs with alternative values of σ_1 , Ω_k , Ω_z , and τ . It is possible to reduce the low-wage phase to 12 years, but *only* if one gets a lucky draw from the parameter space.

Scenario 3: Small Increases in Robot Productivity, But Very High Substitution Between Robots and Human Labor ($\sigma_2 = 10, 20$)


In Scenario 1 the robot revolution is powered by technological innovations that greatly increase robot productivity. This is not the only path to revolution. When robots and labor are almost perfect

²⁸Large cuts in non-robot investment are concentrated in the first ten years. At year x , K is 6.8% below its initial level in the run for $\sigma_2 = 5$ vs. 3.6% in the run for $\sigma_2 = 2.5$. The difference — equal to 10.2% of initial GDP — helps fuel the much faster growth in Z when $\sigma_2 = 5$.

Figure 6. Model 2: The tradeoff between the short and long run



(a) For various values of \bar{b} , $\sigma_2 = 2.5$.



(b) For various values of σ_2 , $\bar{b} = 1.5$.

Note: In all cases $\sigma_1 = .5$, $\sigma_2 = 2.5$, and b increases from .5 to 2.

substitutes, the MPZ schedule in Figure 7 is extremely flat. Even a small upward shift in the schedule then leads to a huge increase in utilization of robots and the destruction of millions of jobs. Extremely high substitution, in other words, puts the economy near a tipping point. The robots are almost ready to take over; they just need to become a little smarter, a little better at doing what humans do.

Figures 8a and 8b depict this scenario.²⁹ The small 20 percent increase in b in Figure 8a does not seem like a big deal, but it excites a tremendous response. Robot investment skyrockets, jumping from 1.8 percent to 6.2 percent of GDP at $t = 0$. Moreover, since the MPZ schedule is very flat and shifts upward every time the traditional capital stock increases, the surge in investment continues far into the future. At $t = 50$, robot investment is 8.4 percent of GDP and rising.

The forces operating on the labor market are the same as in Scenario 1: robots destroy jobs, growth of traditional capital creates them. When σ_2 is large, it takes much longer for the second effect to dominate the first. In Figure 4a, the low-wage phase lasts 27 years; in Figure 8b it lasts 142 years and in Figure 8a it stretches well into the third century (w pulls within 1 percent of its initial value at $t = 283$).

²⁹In these two runs, b jumps at $t = 0$ from .5 to .6.

Table 4. Model 2: Depth and duration of the low-wage phase

Scenario	$\sigma_2 = 2.5$		$\sigma_2 = 5$	
	w decrease	Duration	w decrease	Duration
base case: $\sigma_1 = .50, \Omega_k = 1, \Omega_z = .50, \tau = .50$	-5.3	27	-20.6	38
variations:				
$\sigma_1 = .75$	-2.1	20	-16.2	45
$\sigma_1 = 1$	NA2	NA	-14.8	51
$\Omega_k = \Omega_z = .5$	-4.8	33	-20.0	46
$\Omega_k = \Omega_z = 1$	-5.6	27	-22.8	33
$\Omega_k = \Omega_z = 2$	-6.6	23	-27.5	28
$\Omega_k = \Omega_z = 5$	-7.9	21	-36.0	24
$\tau = .33$	-6.6	32	-22.7	42
$\tau = .75$	-4.3	24	-18.8	34
$\tau = 1$	-3.7	21	-17.9	33
$\sigma_1 = .75, \tau = 1$	-1.0	13	-12.3	35
$\sigma_1 = .75, \tau = 1, \Omega_k = \Omega_z = 2$	-1.4	12	-15.6	24


Note: Notes: "w decrease" is the percentage decrease in the real wage at the trough of the low-wage phase. "Duration" is the length of the low-wage phase in years. "NA" applies to cases where automation has a (minuscule) positive impact on labor demand, which occurs when $\sigma_1 = 1$ and $\sigma_2 = 2.4$. For variations, parameters are as in base case unless otherwise specified.

4 Model 3: Robots Do Not Substitute for Skilled Labor

Model 1 implicitly assumes that robots substitute for low- and high-skill labor to the same degree. Some smart people believe this will soon be true, but, for now, it is very much a minority view. Model 2 assumes that when robots substitute for some labor workers can readily shift into other production tasks. A still-dominant view, however, emphasizes the interaction of automation and skill. The skill-biased technological change hypothesis holds not only that high-skill jobs are less vulnerable to automation than low-skill jobs, but that automation has increased the productivity of skilled labor and contributed significantly to widening wage inequality over the past thirty years (Autor, 2014). The machine-learning crowd might respond that yesterday's world is gone and that in the future skilled labor will find itself struggling in the same way as unskilled labor. Be that as it may, the majority view deserves a hearing.

We bring skilled labor S into the model by adapting Model 2. S replaces L_2 in the CES nest with traditional capital [i.e., $H(K, S)$ replaces $H(K, L_2)$]. The wage paid to skilled labor is w_s and $C[h(r_k, w_s), f(w, r_z/b)]$ is the new unit cost function. Skilled labor saves and invests in the same

Figure 7. Model 2: Response of the robot capital stock
when σ_2 is very high



manner as capitalists. The budget constraint of the representative Ramsey agent thus includes the additional income term $w_s S$. Motivated by the empirical evidence that the supply of skilled labor responds little, if at all, to the skill premium, we abstract from investment in human capital, treating S as perfectly inelastic in both the short and long run.³⁰

The static structure of the model contains an extra market-clearing condition (for S) and an extra endogenous variable (w_s). Formally,

$$K = C_h h_r Q, \quad (39)$$

$$Z = C_f f_z Q/b, \quad (40)$$

$$L = C_f f_w Q, \quad (41)$$


$$S = C_h h_w Q, \quad (42)$$

$$1 = C(r_k, w_s, w, r_z/b), \quad (43)$$

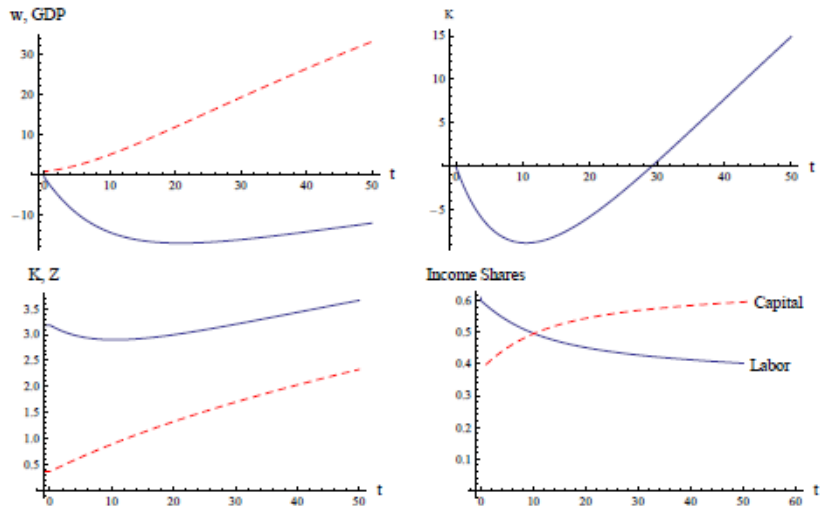
and the optimization problem in (6) - (9) [with income totaling $r_k K + r_z Z + w_s S$ in the budget constraint (7)] define the complete model.

³⁰College graduates' share of total hours worked by males with less than ten years of experience stayed constant at 40 percent in the U.S. as the skill premium rose 80 percent between 1975 and 2015 (Autor, 2015). At the other extreme from the assumption in this section of inelastic supply of skilled labor, the notion that unskilled workers can costlessly become skilled is captured by Model 2.

Figure 8. Model 2: Transition path under alternative scenarios for σ_2 and b



(a) $\sigma_2 = 20$, b increases from .5 to .6



(b) $\sigma_2 = 10$, b increases from .5 to .75

Note: The paths for the real wage and GDP show percentage deviation from initial values

Low-skill labor has a rough time of it in Model 3. The usual algebra delivers

$$\hat{w} = \frac{\sigma_1 - \sigma_2 \theta_H}{\sigma_1 \sigma_2} \alpha_z (\hat{b} + \hat{Z}) + \frac{\theta_K}{\sigma_1} \hat{K}, \quad (44)$$

$$\hat{w}_s = \frac{\theta_Z}{\sigma_1} (\hat{b} + \hat{Z}) + \frac{\sigma_1 - \sigma_3 \theta_V}{\sigma_1 \sigma_3} \chi_k \hat{K}. \quad (45)$$

Automation always increases w_s . It reduces w iff

$$\sigma_2 > \frac{\sigma_1}{\theta_H}. \quad (46)$$

Since $\theta_H = \theta_K + \theta_S$ and $\theta_S \approx .30$ in the data,³¹ this condition is appreciably weaker than $\sigma_2 > \sigma_1/\theta_K$, the corresponding condition in Model 1.³²

The results for the long run are also an order of magnitude worse. Across steady states,

$$\hat{K} = \frac{\sigma_3 \sigma_2 \theta_Z}{J} \hat{b} > 0, \quad (47)$$

$$\hat{Q} = (\sigma_3 \chi_k + \sigma_1 \chi_s) \frac{\sigma_2 \theta_Z}{J} \hat{b} > 0, \quad (48)$$

$$\hat{w} = \left[\frac{\chi_s}{\theta_V} (\sigma_1 - \sigma_2 \theta_H) + \sigma_3 \chi_k \right] \frac{\theta_Z}{J} \hat{b} \geq 0, \quad (49)$$

$$\hat{w}_s = \frac{\sigma_2 \theta_z}{J} \hat{b} > 0, \quad (50)$$

$$\theta_S \hat{w}_s + \theta_L \hat{w} = \theta_Z \hat{b} > 0 \quad \text{sgn} \quad \frac{S}{L+S} dw_s + \frac{L}{L+S} dw, \quad (51)$$

where

$$J \equiv \sigma_2 \theta_s \alpha_z + \sigma_1 \chi_s \alpha_L + \sigma_3 \chi_k \theta_L.$$

The zero-profit condition tells us only that the (employment or cost share) weighted-average wage increases. There is no assurance that both wages increase. The skilled wage rises, but when

$$\sigma_2 > \frac{\sigma_1}{\theta_H} \left(1 + \frac{\sigma_3 \chi_k \theta_V}{\sigma_1 \chi_s} \right) \quad (52)$$

the low-skill wage stays lower forever. Skilled labor then claims all of the productivity gains along with some of the wages formerly paid to poor, low-skill workers [$\hat{w}_s = (\theta_Z \hat{b} - \theta_L \hat{w})/\theta_S$, where $\hat{w} < 0$].

This ugly outcome is almost a foregone conclusion. Empirical estimates suggest $\sigma_3 \approx .5\sigma_1$, while $\chi_k \approx \chi_s$ and $\theta_V \approx .35$ in the raw data.³³ The "mean" value of the term on the right side in (52) is thus $1.82\sigma_1$. Even if we throw in a couple of standard deviations so that $\sigma_3 = \sigma_1$ joins the potentially relevant parameter space, the borderline value of σ_2 is only 1-2.1 for $\sigma_1 = .5-1$. Econometricians have yet to estimate σ_2 , but the findings in Acemoglu and Restrepo (2016b) and narrative accounts in case studies argue that values of σ_2 in this range are unrealistically low. In all likelihood, σ_2 satisfies the condition in (52) with room to spare.

4.1 Multi-Dimensional Inequality

In contrast to Models 1 and 2, inequality in Model 3 is multi-dimensional. Four variables matter: the income shares of low- and high-skill labor, θ_L and θ_S ; the aggregate income share of labor, $\theta_{LS} = \theta_L + \theta_S$; and the skill premium, w_s/w .

³¹For $\theta_s = .30$ and an income share of .61 for all labor, the share of skilled labor in the wage bill is 49 percent vs. 41-54 percent in Appendix 1 in Autor et al. (1998).

³²In the unlikely event that $\sigma_2 < \sigma_1/\theta_H$, w increases. Wage inequality worsens, however ($\hat{w}_s > \hat{w}$, assuming $\sigma_2 > \sigma_1$).

³³See Griliches (1969), Fallon and Layard (1975), Hamermesh (1993), Krusell et al. (2000), and Raval (2011) for estimates of how substitution between capital and skilled labor compares to substitution between capital and unskilled labor. $\theta_s \approx .30$ in the raw data. This together with $\theta_K = .35$ implies $\theta_H = .65$, $\theta_V = .35$, $\chi_k = .538$ and $\chi_s = .462$.

Under weak conditions, inequality worsens in all dimensions. From (44) and (45),

$$\hat{w}_s > \hat{w} \quad \text{iff} \quad \sigma_2 > \frac{\sigma_1\chi_s + \sigma_3\chi_k\theta_V}{\theta_s + \theta_V}. \quad (53)$$

The term on the right side is smaller than σ_1/θ_H if $\sigma_1 > \sigma_3\theta_K$, which is sure to hold. Wage inequality worsens in both the short and long run, therefore, whenever automation reduces the demand for low-skill labor (for given K).³⁴

The long-run income shares for labor can be computed from (48)-(45). After simplification, the solutions read

$$\hat{\theta}_L = \frac{\theta_Z}{J} \left[\frac{\theta_S}{\theta_V}(\sigma_1 - \sigma_2) + \sigma_3\chi_k(1 - \sigma_2) \right] \hat{b}, \quad (54)$$

$$\hat{\theta}_S = \frac{\sigma_2\theta_Z}{J}(1 - \sigma_1\chi_s - \sigma_3\chi_k)\hat{b}, \quad (55)$$

$$\hat{\theta}_{LS} = \{\sigma_1\chi_s + \sigma_3\chi_k\theta_V - \sigma_2[(\sigma_1\chi_s + \sigma_3\chi_k)\theta_{LS} - \theta_S\alpha_z]\} \frac{\theta_Z}{\theta_{LS}J} \hat{b}. \quad (56)$$

Low-skill labor loses ground: its income share decreases assuming $\sigma_2 > \text{Max}\{\sigma_1, 1\}$. The share for skilled labor rises or falls depending on whether the weighted average of σ_1 and σ_3 is above or below unity. Given the empirical evidence that $\sigma_3 < \sigma_1 < 1$, an increase in θ_S is much more likely than not. Finally, the total income share for labor decreases if

$$\sigma_1\chi_s + \sigma_3\chi_k > \frac{\theta_S\alpha_z}{\theta_{LS}} \quad (57)$$

and

$$\sigma_2 > \alpha_L \frac{\sigma_1\chi_s + \sigma_3\chi_k\theta_V}{\theta_{LS}(\sigma_1\chi_s + \sigma_3\chi_k) - \theta_S\alpha_z}. \quad (58)$$

The first condition holds ($\theta_S\alpha_z/\theta_{LS}$ is close to zero). The second condition is not automatic, but it holds easily in the relevant parameter space. Suppose $\theta_K = .35$, $\theta_S = \theta_L = .30$, $\sigma_1 = .5-2$, and $\sigma_3 = .5-2$. Despite the inclusion of wildly biased and irrelevant parameter combinations ($\sigma_1 = 2?$, $\sigma_3 = 2?$), the highest borderline value of σ_2 is only .98.

4.2 The Transition Path: Distributional Carnage

It is obvious from the preceding results that skilled labor and capitalists make out very well. The only unresolved question is whether the transition path and the long-run outcome are bad or very bad for low-skill labor.

Both are *very* bad. There is no need to comment at length on Figures 9a-9c.³⁵ From a distributional standpoint, they are a disaster, with 26-56 percent reductions in the low-skill wage while GDP increases 30-105 percent and the share of national income paid to capital and skilled labor rises from 69 percent to 82-92 percent.


³⁴Positing $\sigma_2 < \sigma_1/\theta_H$ does not overturn this conclusion unless σ_3 is absurdly large. When $\sigma_3 = \sigma_1$, for example, the skill premium increases in the short and long run if $\sigma_2 > \sigma_1$.

³⁵We calibrate the model for $\theta_S = .30$, $\theta_L = .31$, and $w_s/w = 2$. All other parameters take the same values as in the baseline calibration ($\rho = .06$, $\Omega_K = 1$, etc.). The values assigned to θ_S , θ_L , and w_s/w are consistent with the data on the share of skilled wages in the total wage bill, the wage ratio for high- vs. low-skill labor, the college wage premium, and the ratio of hours worked for high- vs. low-skill labor (see Johnson (1997), Acemoglu and Autor (2011b), Blankenau and Cassou (2011), and Autor (2014)).


Figure 9. Model 3: The transition path for various elasticities



(a) $\sigma_2 = 2.5, \sigma_1 = .67, \sigma_3 = .33$



(b) $\sigma_2 = 5, \sigma_1 = \sigma_3 = .5$



(c) $\sigma_2 = 5, \sigma_1 = 1, \sigma_3 = .5$

Note: b increases from .5 to 2. See footnote 35 for explanation of scenario calibration.

5 Model 4: Adding a Non-Automatable Sector

We consider one final way in which we could be setting things up for failure. Because they have only one sector, Models 1-3 assume perforce that automation penetrates production processes throughout the economy. Some sectors, however, seem largely immune to automation. Home nursing care, the arts and entertainment, and a variety of service-related industries such as elementary education, yoga, and psychotherapy come to mind. When automation increases real income, these sectors will expand to meet higher demand. If they expand enough and if factor-intensity conditions take the right form, the real wage for low-skill labor may increase in general equilibrium despite layoffs in the automatable sectors. Autor (2015) propounds this hypothesis, speculating that the income and price elasticities of demand for non-automatable (NA) goods and substitutability of labor across sectors will play a large role in determining the outcome.

We now have the apparatus to examine this idea carefully. In particular, we test this case for technology optimism by adding a NA sector to Model 3. Firms in sector 2 (the NA sector) operate a two-tiered CES production function in which low-skill labor combines in the upper tier with a composite input formed by capital and skilled labor in the lower tier.³⁶ Labor is intersectorally mobile, while capital is sector-specific. The unit cost function in sector 2 is thus $C^2[h^2(r_{k2}, w_s), w]$. Good 1 may be used either for consumption or investment, good 2 only for consumption. Households derive utility from a CES aggregate of the two consumption goods.³⁷ The exact price index for the composite consumer good is $p = [m + (1 - m)p_2^{1-\epsilon}]^{1/(1-\epsilon)}$, where m is a distribution parameter, ϵ is the elasticity of substitution, and p_2 is the relative price of good 2. See Appendix C for other details and a full statement of the model.

Although analytical results are out of reach for Model 4, it is useful to consult certain intermediate-stage solutions. Consider the zero-profit conditions

$$1 = C^1(r_{k1}, w_s, w, r_z/b), \quad (59)$$

$$p_2 = C^2(r_{k2}, w_s, w). \quad (60)$$

Across steady states, $\hat{r}_{k1} = \hat{r}_{k2} = \hat{r}_z = 0$ and

$$\hat{w} = \frac{\theta_S^1 \hat{p}_2 - \theta_S^2 \theta_Z \hat{b}}{\Delta}, \quad (61)$$

$$\hat{w}_s = \frac{\theta_L^2 \theta_Z \hat{b} - \theta_L^1 \hat{p}_2}{\Delta}, \quad (62)$$

where

$$\Delta = \theta_S^1 \theta_L^2 - \theta_S^2 \theta_L^1 \quad \text{which has the same sign as } L_2/S_2 - L_1/S_1.$$

For the real wages $\omega = w/p$ and $\omega_s = w_s/p$,

$$\hat{\omega} = \hat{w} - \gamma \hat{p}_2 = \frac{[\theta_S^1(1 - \gamma \theta_L^2) + \gamma \theta_S^2 \theta_L^1] \hat{p}_2 - \theta_S^2 \theta_Z \hat{b}}{\Delta}, \quad (63)$$

$$\hat{\omega}_s = \hat{w}_s - \gamma \hat{p}_2 = \frac{\theta_S^2 \theta_Z \hat{b} - [\theta_L^1(1 - \gamma \theta_S^2) + \gamma \theta_S^1 \theta_L^2] \hat{p}_2}{\Delta}, \quad (64)$$

³⁶We could suppose that the lower tier composite is of capital and unskilled labor, which would then combine with skilled labor in the upper tier. However, the literature and associated empirical evidence follows the approach we assume here. See Krusell et al. (2000) and Papageorgiou and Saam (2008).

³⁷This implies unitary income elasticity of demand for output in both sectors.

where γ is the consumption share of the NA good.

These solutions are incomplete (p_2 is a function of b), but informative. Two points merit emphasis. First, factor-intensity conditions are important. The low-skill wage increases with p_2 and decreases with b when employment in the NA sector is relatively intensive in low-skill labor (i.e., $\Delta > 0$ or equivalently $L_2/S_2 - L_1/S_1 > 0$). Second, it is hard to get p_2 and b to collaborate and pull strongly in the desired direction. An increase in b raises real income and demand for good 2. The net effect on supply in sector 2 is positive and large when employment in the sector is relatively low-skill intensive. In this case, $\Delta > 0$ and $b \uparrow$ tends to lower ω while $p_2 \uparrow$ has the opposite effect. But the large rightward shift of the supply curve in sector 2 weakens the positive effect by limiting the increase in p_2 . (Indeed, p_2 could decrease.) Conversely, when sector 2 is relatively skill intensive, $\Delta < 0$ and p_2 increases sharply because the supply curve shifts left or only a small amount to the right. The negative effect on ω of p_2 is then large relative to the positive effect of $b \uparrow$. Many more general equilibrium effects are at play in Model 4 than in Model 3. But since the effects are offsetting rather than reinforcing, it is unlikely their net impact will be large enough to reverse or significantly ameliorate the harsh distributional results in Model 3.

Figure 10 confirms this conjecture.³⁸ In both runs, high-skill labor gains more and low-skill labor loses more in Model 4 than in Model 3.³⁹ Because GDP growth is higher in Model 4, the income share of low-skill labor also decreases more.

We generated many other solutions with widely varying values for τ , ϵ , σ_1 , σ_2 , σ_3 , β_1 , and β_3 (elasticities of substitution in sector 2 analogous to σ_1 and σ_3 in sector 1). The runs always came back looking like either Figure 10a or Figure 10b.⁴⁰ After investigating further, we managed to find some scenarios in which real wages increase for both low- and high-skill labor. But the scenarios require that b increase only a little, that σ_2 not be too large, and that other parameters take just the right values.⁴¹ And since σ_2 and the increase in b are small, real wages increase only a few percent. Technology optimism is confined to a small, uninteresting part of the parameter space that produces small, uninteresting results.

6 Concluding Remarks

The premise of this paper is that we are in the midst of a technological inflection point, a new "machine age" in which artificial intelligence and robots are rapidly developing the capacity to do the cognitive as well as physical work of large fractions of the labor force. This idea is dominating discussions of technologists, business elites, and policymakers. The macroeconomic literature has lagged behind, with few papers coming to grips with this possibility and analyzing the consequences. This has left informal and policy discussions to be based (often implicitly) on inappropriate models.

Here we present a general equilibrium model to analyze the implications of robots for output, wages, and inequality. Despite its complexities, this model is the simplest we could think of with the critical ingredients: a second type of capital ("robots") that operates alongside traditional capital and labor;


³⁸The differences in relative skill intensity of sectors 1 and 2 in Figure 10 are in line with the data on low- vs. high-skill sectors in Blankenau and Cassou (2011).

³⁹The qualitative comparisons change sign in the very long run. Across steady states, the decrease in the low-skill wage and the increases in the high-skill wage and GDP are all greater in Model 3 than in Model 4.


⁴⁰Most notably, the results are very insensitive to ϵ , the elasticity of substitution between goods 1 and 2. Raising ϵ to five has almost no effect on the runs in Figures 10a and 10b.

⁴¹The NA sector gets driven out of business in runs with modest values for σ_2 or small (as opposed to very small) increases in b . Model 4 then produces the same results as Model 3.

Figure 10. Model 4: Transition paths for differing factor intensities in the non-automatable sector



(a) The ratio of unskilled to skilled labor in sector 2 is 1.6



(b) The ratio of unskilled to skilled labor in sector 2 is 0.44

Note: $\sigma_2 = 5$, $\sigma_1 = \sigma_3 = \beta_1 = \beta_3 = .5$, and b increase from .5 to 2. The ratio of skilled to unskilled labor is 0.967 in sector 1.

and plausible savings behavior in which owners of capital (and skilled workers, in the variants that include them) do the saving. Most of the economics literature on inequality and technology has focused on the wage differential between high and low-skill workers. Or, it has analyzed an undifferentiated capital stock such that technological advances always raise the real wage. Often, the conclusion has been essentially optimistic. As with past technological change, some categories of workers or tasks may be harmed, at least in the short run. But output and overall wages should increase, as new technologies open up more opportunities than they close. As long as labor force skills keep up, we have little to worry about.

We show that this time may indeed be different. In our benchmark model, robots are close substitutes for human workers, combining with them in a CES aggregate to produce labor services. These labor services then combines with traditional capital in a CES production function to produce final output. Even a small increase in the level of robot productivity can increase output enormously if

the robots and humans are sufficiently close substitutes. The basic mechanism is that the introduction of more productive robots initially lowers the wage and raises the return to both robots and traditional capital. Large quantities of traditional capital have to be accumulated before a scarcity of human labor raises wages and the return on capital declines to normal levels.

All this is thus very good for output. It is also very bad for distribution. For the real wage, understanding the dynamics is critical; the short run—which can last for generations—is very different from the long run. At first, the real wage is likely to fall in absolute terms, even as the economy grows. Eventually, the real wage will rise above initial levels, but there are two distributional problems. First, “eventually” can take a long time, typically 20 to 50+ years in our baseline calibrations. Second, even in the long run, the labor share declines substantially and overall inequality rises.

The benchmark model glosses over many issues that have been central in the policy debate about the new machine age. We address these issues in extensions of the baseline model in which robots (i) compete only for some tasks (Model 2); (ii) substitute only for unskilled labor while complementing skilled labor (Model 3); and (iii) contribute to production only in one sector — elsewhere firms use only labor and traditional capital (Model 4). It turns out, perhaps surprisingly, that our main results are robust across the different models. In all variants, for reasonable parameter values, automation is very good for growth and very bad for equality, including those reputed to be conducive to technological optimism.

When robots only compete for some tasks (Model 2), both the impact on growth and on inequality tend to be attenuated, but it remains the case that inequality worsens in the short run as wages fall and in the long run, as the capital share rises. Allowing for tasks that complement robots does not help as much as one might think, partly because more and more workers compete for those jobs, driving down the overall wage. When, alternatively, we close off labor mobility between tasks, assuming that robots complement “skilled” labor and substitute for “unskilled” labor (Model 3), the distributional outcome is much more uneven. In addition to the fall in the average wage and the rise in the capital share, unskilled workers suffer large decreases in absolute and relative wages. Finally, we assume a non-automatable sector where robots cannot compete. This also does not really help, again because there are only so many of those jobs to go around, and labor chased out of the automatable sector tends to drive down wages. Of course, it is not impossible to overturn our main results. If robots do only a very small fraction of tasks, contribute to output in only a very small fraction of the economy, or are poor substitutes for human workers, then they will not increase inequality. But this is another way of restating the premise of the paper: if there is no technological revolution underway of the sort we have been discussing, then there will only be small effects.

We see three natural extensions of this work as being of high priority. The first is to discuss policy responses. Education is often touted as the main solution to the challenges of automation (The White House, 2016, e.g.). We can already see the limits of this view. Education can be viewed as investment to convert workers from “unskilled” to “skilled.” Doubtless this would reduce wage inequality and strengthen the demand for unskilled labor. But can it offset the huge real wage cuts unskilled labor suffers and the decrease in labor’s overall income share at an acceptable cost? And if the answer is yes, how long will it take for wages to increase for those who remain unskilled?

Another set of policy responses involves transfers from the owners of capital to workers, perhaps in the form of a universal basic income, widely touted by industrialists and technologists as well as some economists. It would be interesting to examine in our framework the implications of taxing capital initially (when returns are especially high) to finance ownership shares in capital that could be distributed to or held on behalf of workers. Of course, there are many practical and conceptual chal-

lenges. Among the latter, one worth noting at the outset is that growth in our framework derives not directly from technological change or total-factor-productivity but rather from capital accumulation, especially of traditional capital, as robots keep returns to this capital higher for longer. A conjecture is that this makes the potential negative implications of capital taxation for growth especially strong. Among the more practical considerations, capital taxation is becoming more and more difficult with globalization and the increasing capital mobility.

A second set of extensions could use this framework to examine the experience of recent decades in the United States and other countries on the technological frontier. Along the lines of Nordhaus (2015), but with a more appropriate model, it might be revealing to consider whether the recent experience with stagnant wages and falling labor shares can be usefully explained with a framework such as ours.

Finally, we have considered a closed economy at the technological frontier. However, a technological revolution such as we have examined here has profound implications for the international division of labor and prospects for developing countries. The introduction of robots is by no means limited to frontier economies. For example, nearly full automation of China's iPhone factories is apparently underway, with a benchmark of 30 percent replacement by robots by 2020, while technology is threatening call center industries in places such as India and the Philippines.⁴² Rodrik (2016) documents the phenomenon of "premature deindustrialization" and links it to both globalization and technological change. Extending our models to consider these international dimensions could be useful.

⁴²On iPhone production, see for example The South China Morning Post (2015). On call centers, see for example Economist (2016). See also World Bank (2016).

Appendix A: Model 1

Equations (2)-(5) can be solved for r_k , r_z , w , and Q as a function of K , Z , and b . To express the results in an intelligible form, note that the Allen-Uzawa partial elasticity of substitution between inputs i and j is

$$\sigma_{ij} = \frac{C_{ij}C}{C_i C_j}.$$

Applying the formula to equation (2) gives

$$\begin{aligned} \hat{K} &= \underbrace{\left(\frac{C_{rr}C}{C_r^2}\right)}_{\sigma_{KK}} \underbrace{\left(\frac{r_k C_r}{C}\right)}_{\theta_K} \hat{r}_k + \underbrace{\left(\frac{C_{rf}C}{C_r C_f}\right)}_{\sigma_1} \underbrace{\left(\frac{f C_f}{C}\right)}_{\theta_V} \hat{f} + \hat{Q} \\ \implies \hat{K} &= \sigma_{KK} \theta_K \hat{r}_k + \sigma_1 \theta_V \hat{f} + \hat{Q}. \end{aligned} \quad (\text{A1})$$

The adding-up condition $\sum_j \theta_j \sigma_{ij} = 0$ says that $\sigma_{KK} \theta_K = -\sigma_1 \theta_V$. Thus⁴³

$$\begin{aligned} \hat{K} &= \sigma_1 \theta_V \left[\underbrace{(f_w w / f)}_{\alpha_L} \hat{w} + \underbrace{(f_z r_z / b f)}_{\alpha_z} (\hat{r}_z - \hat{b}) - \hat{r}_k \right] + \hat{Q}, \\ \implies \hat{K} &= \sigma_1 [\theta_L \hat{w} + \theta_Z (\hat{r}_z - \hat{b}) - \theta_V \hat{r}_k] + \hat{Q}. \end{aligned} \quad (\text{A2})$$

Equations (3) and (4) are more involved. From (3),

$$\begin{aligned} \hat{Z} &= \underbrace{\left(\frac{C_{fr}C}{C_r C_f}\right)}_{\sigma_1} \underbrace{\left(\frac{r_k C_r}{C}\right)}_{\theta_K} \hat{r}_k + \underbrace{\left(\frac{C_{ff}C}{C_f^2}\right)}_{\sigma_{VV}} \underbrace{\left(\frac{f C_f}{C}\right)}_{\theta_V} \hat{f} + \underbrace{\left(\frac{f_{zw}f}{f_z f_w}\right)}_{\sigma_2} \underbrace{\left(\frac{w f_w}{f}\right)}_{\alpha_L} \hat{w} \\ &+ \underbrace{\left(\frac{f_{zz}f}{f_z^2}\right)}_{\sigma_{ZZ}} \underbrace{\left(\frac{f_z r_z / b}{f}\right)}_{\alpha_z} (\hat{r}_z - \hat{b}) + \hat{Q} - \hat{b}. \end{aligned}$$

After making use of the adding-up conditions $\theta_V \sigma_{VV} + \sigma_1 \theta_K = 0$ and $\alpha_z \sigma_{ZZ} + \alpha_L \sigma_2 = 0$, this simplifies to

$$\begin{aligned} \hat{Z} &= \sigma_1 \theta_K [\hat{r}_k - \alpha_L \hat{w} - \alpha_z (\hat{r}_z - \hat{b})] + \sigma_2 \alpha_L [\hat{w} - (\hat{r}_z - \hat{b})] + \hat{Q} - \hat{b}, \\ \implies \hat{Z} &= \sigma_1 \theta_K \hat{r}_k + \alpha_L (\sigma_2 - \sigma_1 \theta_K) \hat{w} - (\sigma_1 \alpha_z \theta_K + \sigma_2 \alpha_L) (\hat{r}_z - \hat{b}) + \hat{Q} - \hat{b}. \end{aligned} \quad (\text{A3})$$

Executing the same manipulations in equation (4) yields

$$(\sigma_1 \alpha_L \theta_K + \sigma_2 \alpha_z) \hat{w} = \sigma_1 \theta_K \hat{r}_k + \alpha_z (\sigma_2 - \sigma_1 \theta_K) (\hat{r}_z - \hat{b}) + \hat{Q}. \quad (\text{A4})$$

⁴³Recall that $f_z = \partial f / \partial (r_z / b)$.

Last but not least, the zero-profit condition provides

$$\theta_K \hat{r}_k + \theta_L \hat{w} + \theta_Z (\hat{r}_z - \hat{b}) = 0. \quad (\text{A5})$$

Solving (A2)-(A5) for w , r_k , r_z , and Q is a straightforward, albeit tedious, business. First substitute $\hat{r}_z - \hat{b} = -(\theta_K \hat{r}_k + \theta_L \hat{w})/\theta_Z$ in (A2) and (A4). This produces

$$\hat{Q} = \hat{K} + \sigma_1 \hat{r}_k \quad (\text{A6})$$

and

$$\begin{aligned} \sigma_2 \hat{w} &= (\sigma_1 - \sigma_2) \frac{\theta_K}{\theta_V} \hat{r}_k + \hat{Q}, \\ \implies \sigma_2 \hat{w} &= \frac{\sigma_1 - \sigma_2 \theta_K}{\theta_V} \hat{r}_k + \hat{K}, \end{aligned} \quad (\text{A7})$$

as $\theta_K + \theta_V = 1$.⁴⁴

Turn next to equation (A3). Substitute for \hat{Q} and $\hat{r}_z - \hat{b}$ and solve for \hat{r}_k :

$$\begin{aligned} \{\sigma_1[1 + \theta_K(1 + \theta_K/\theta_V)] + \sigma_2 \alpha_L \theta_K / \theta_Z\} \hat{r}_k &= \hat{b} + \hat{Z} - \hat{K} - [\alpha_L(\sigma_2 - \sigma_1 \theta_K) \\ &\quad + (\sigma_1 \alpha_z \theta_K + \sigma_2 \alpha_L) \theta_L / \theta_Z] \hat{w}, \end{aligned} \quad (\text{A8})$$

$$\frac{a_2}{\theta_V} \hat{r}_k = \hat{b} + \hat{Z} - \hat{K} - \frac{\alpha_L}{\alpha_z} \sigma_2 \hat{w}, \quad (\text{A9})$$

where

$$a_2 \equiv \sigma_1 + \sigma_2 \alpha_L \theta_K / \alpha_z.$$

Solve (A7) and (A9) for w and r_k :

$$\hat{w} = \frac{\sigma_1 - \sigma_2 \theta_K}{\sigma_1 \sigma_2} \alpha_z (\hat{b} + \hat{Z}) + \frac{\theta_K}{\sigma_1} \hat{K}, \quad (\text{A10})$$

$$\hat{r}_k = \frac{\theta_Z (\hat{b} + \hat{Z}) - \theta_V \hat{K}}{\sigma_1}. \quad (\text{A11})$$

Substitute for \hat{w} and \hat{r}_k in (A5) to get the solution for \hat{r}_z . After considerable simplification, the solution reads

$$\hat{r}_z = \left[1 - \frac{\sigma_2 \alpha_z \theta_K + \sigma_1 \alpha_L}{\sigma_1 \sigma_2} \right] \hat{b} - \frac{\sigma_2 \alpha_z \theta_K + \sigma_1 \alpha_L}{\sigma_1 \sigma_2} \hat{Z} + \frac{\theta_K}{\sigma_1} \hat{K}. \quad (\text{A12})$$

⁴⁴Here and in subsequent derivations we make use of the links between the aggregate cost shares and the cost shares in the composite labor input ($\theta_Z = \alpha_z \theta_V$ and $\theta_L = \alpha_L \theta_V$).

The Long-Run Outcome

Across steady states where $\hat{r}_k = \hat{r}_z = 0$, equations (A5) and A6) yield

$$\hat{w} = \frac{\theta_Z}{\theta_L} \hat{b} = \frac{\alpha_z}{\alpha_L} \hat{b}, \quad (\text{A13})$$

$$\hat{Q} = \hat{K}. \quad (\text{A14})$$

From (A7) and (A13),

$$\begin{aligned} \hat{K} &= \hat{Q} = \sigma_2 \hat{w}, \\ \implies \hat{K} &= \hat{Q} = \sigma_2 \frac{\alpha_L}{\alpha_z} \hat{b}. \end{aligned} \quad (\text{A15})$$

The steady-state version of (A9) is

$$\begin{aligned} \hat{Z} &= \hat{K} + \sigma_2 \frac{\alpha_L}{\alpha_z} \hat{w} - \hat{b}, \\ \implies \hat{Z} &= (\sigma_2 / \alpha_L - 1) \hat{b}. \end{aligned} \quad (\text{A16})$$

Solving for the Transition Path

The core dynamic system is comprised of

$$\dot{K} = I_k - \delta K \quad (8)$$

$$\dot{Z} = I_z - \delta Z \quad (9)$$

$$\dot{b} = s(\bar{b} - b), \quad s > 0 \quad (12)$$

and the two Euler equations for investment

$$\frac{v_k}{K} \dot{I}_k = \psi_k \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 - r_k \quad (10)$$

$$\frac{v_z}{Z} \dot{I}_z = \psi_z \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 - r_z \quad (11)$$

where $\psi_k \equiv 1 + v_k(I_k/K - \delta)$ and $\psi_z \equiv 1 + v_z(I_z/Z - \delta)$.

To get the system in the right form, we need pseudo reduced-form solutions that relate r_k , r_z , and \dot{c} to K , Z , b , I_k , and I_z . Equations (A11) and (A12) supply the requisite solutions for r_k and r_z . For \dot{c} , the private agent's budget constraint gives

$$\begin{aligned} \dot{c} &= r_k \dot{K} + r_z \dot{Z} + K \dot{r}_k + Z \dot{r}_z - \dot{I}_k - \dot{I}_z - \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 \dot{K} - v_k \left(\frac{I_k}{K} - \delta \right) \dot{I}_k \\ &\quad + v_k \left(\frac{I_k}{K} - \delta \right) \frac{I_k}{K} \dot{K} - \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 \dot{Z} - v_z \left(\frac{I_z}{Z} - \delta \right) \dot{I}_z + v_z \left(\frac{I_z}{Z} - \delta \right) \frac{I_z}{Z} \dot{Z}. \end{aligned}$$

Substitute for \dot{r}_k and \dot{r}_z from (A11) and (A12):

$$\begin{aligned}
\dot{c} &= r_k \dot{K} + r_z \dot{Z} + r_k K [n_2 (\dot{Z}/Z + \dot{b}/b) - n_1 \dot{K}/K] + r_z Z [(1 - n_3) \dot{b}/b - n_3 \dot{Z}/Z + n_4 \dot{K}/K] \\
&\quad - \dot{I}_k - \dot{I}_z - \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 \dot{K} - v_k \left(\frac{I_k}{K} - \delta \right) \dot{I}_k + v_k \left(\frac{I_k}{K} - \delta \right) \frac{I_k}{K} \dot{K} \\
&\quad - \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 \dot{Z} - v_z \left(\frac{I_z}{Z} - \delta \right) \dot{I}_z + v_z \left(\frac{I_z}{Z} - \delta \right) \frac{I_z}{Z} \dot{Z}, \\
\implies \dot{c} &= (r_k(1 - n_1) + r_z n_4 Z/K)(I_k - \delta K) + (r_z(1 - n_3) + r_k n_2 K/Z)(I_z - \delta Z) \\
&\quad + [r_k n_2 K/b + r_z(1 - n_3)Z/b]s(\bar{b} - b) - \dot{I}_k - \dot{I}_z - \frac{v_k}{2} \left(\frac{I_k}{K} - \delta \right)^2 \dot{K} - v_k \left(\frac{I_k}{K} - \delta \right) \dot{I}_k \\
&\quad + v_k \left(\frac{I_k}{K} - \delta \right) \frac{I_k}{K} \dot{K} - \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 \dot{Z} - v_z \left(\frac{I_z}{Z} - \delta \right) \dot{I}_z + v_z \left(\frac{I_z}{Z} - \delta \right) \frac{I_z}{Z} \dot{Z}, \tag{A17}
\end{aligned}$$

where

$$\begin{aligned}
n_1 &= \theta_V/\sigma_1, \\
n_2 &= \theta_Z/\sigma_1, \\
n_3 &= \frac{\sigma_2 \alpha_z \theta_K + \sigma_1 \alpha_L}{\sigma_1 \sigma_2}, \\
n_4 &= \theta_K/\sigma_1.
\end{aligned}$$

Linearize (A17) around a steady state. The result is⁴⁵

$$\dot{c} = (\rho + \delta)u_1(dI_k - \delta dK) + (\rho + \delta)u_2(dI_z - \delta dZ) - u_3 s db - \dot{I}_k - \dot{I}_z, \tag{A18}$$

where the differentials denote deviations from the steady state and

$$\begin{aligned}
u_1 &= 1 + n_4 Z/K - n_1, \\
u_2 &= 1 + n_2 K/Z - n_3, \\
u_3 &= \frac{\rho + \delta}{b} [n_2 K + (1 - n_3)Z].
\end{aligned}$$

Now linearize (10). The first step generates

$$\frac{v_k}{K} \dot{I}_k = \frac{\dot{c}}{c\tau} + (\rho + \delta) \frac{v_k}{K} (dI_k - \delta dK) - (\rho + \delta) \hat{r}_k.$$

Substituting for \hat{r}_k and \dot{c} from (A11) and (A17) leads to

⁴⁵The adjustment cost terms drop out because $I_k/K - \delta$, $I_z/Z - \delta$, \dot{K} , \dot{Z} , \dot{I}_k , and \dot{I}_z all equal zero evaluated at a steady state.

$$\begin{aligned} \left(1 + \frac{v_k c\tau}{K}\right) \dot{I}_k + \dot{I}_z &= (\rho + \delta) \left(u_1 + \frac{v_k c\tau}{K}\right) dI_k + (\rho + \delta) u_2 dI_z + (\rho + \delta) \left(\frac{n_1 c\tau}{K} - u_1 \delta - \frac{\delta v_k c\tau}{K} dK\right) \\ &\quad - (\rho + \delta) \left(\frac{n_2 c\tau}{Z} + u_2 \delta\right) dZ - \left[\frac{(\rho + \delta) n_2 c\tau}{b} + u_3 s\right] db. \end{aligned} \quad (\text{A19})$$

Linearizing (11) yields

$$\frac{v_z}{Z} \dot{I}_z = \frac{\dot{c}}{c\tau} + (\rho + \delta) \frac{v_z}{Z} (dI_z - \delta dZ) - (\rho + \delta) \hat{r}_z$$

and the symmetric solution

$$\begin{aligned} \left(1 + \frac{v_z c\tau}{Z}\right) \dot{I}_z + \dot{I}_k &= (\rho + \delta) u_1 dI_k + (\rho + \delta) \left(u_2 + \frac{v_z c\tau}{Z}\right) dI_z - (\rho + \delta) \left(\frac{n_4 c\tau}{K} + u_1 \delta\right) dK \\ &\quad + (\rho + \delta) \left(\frac{n_3 c\tau}{Z} - u_2 \delta - \frac{\delta v_z c\tau}{Z}\right) dZ - \left[\frac{(\rho + \delta) c\tau (1 - n_3)}{b} + u_3 s\right] db. \end{aligned} \quad (\text{A20})$$

The linearized versions of (8), (9), and (12) are simply

$$\dot{K} = dI_k - \delta dK, \quad (\text{A21})$$

$$\dot{Z} = dI_z - \delta dZ, \quad (\text{A22})$$

$$\dot{b} = -s db. \quad (\text{A23})$$

The complete linearized system is thus

$$\begin{bmatrix} e_1 & 1 & 0 & 0 & 0 \\ 1 & e_2 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \dot{I}_k \\ \dot{I}_z \\ \dot{K} \\ \dot{Z} \\ \dot{b} \end{bmatrix} = \begin{bmatrix} (\rho + \delta) h_1 & (\rho + \delta) u_2 & e_3 & -e_4 & -g_1 \\ (\rho + \delta) u_1 & (\rho + \delta) h_2 & -e_5 & e_6 & -g_2 \\ 1 & 0 & -\delta & 0 & 0 \\ 0 & 1 & 0 & -\delta & 0 \\ 0 & 0 & 0 & 0 & -s \end{bmatrix} \begin{bmatrix} I_k - \bar{I}_k \\ I_z - \bar{I}_z \\ K - \bar{K} \\ Z - \bar{Z} \\ b - \bar{b} \end{bmatrix}, \quad (\text{A24})$$

where

$$\begin{aligned} e_1 &= 1 + v_k c\tau / K, \\ e_2 &= 1 + v_z c\tau / Z, \\ h_1 &= u_1 + v_k c\tau / K, \\ h_2 &= u_2 + v_z c\tau / Z, \\ e_3 &= (\rho + \delta) (n_1 c\tau / K - u_1 \delta - \delta v_k c\tau / K), \\ e_4 &= (\rho + \delta) (n_2 c\tau / Z + u_2 \delta), \\ e_5 &= (\rho + \delta) (n_4 c\tau / K + u_1 \delta), \\ e_6 &= (\rho + \delta) (n_3 c\tau / Z - u_2 \delta - \delta v_z c\tau / Z), \\ g_1 &= (\rho + \delta) n_2 c\tau / b + u_3 s, \\ g_2 &= (\rho + \delta) (1 - n_3) c\tau / b + u_3 s. \end{aligned}$$

Since b , K , and Z are state variables, the steady state is saddle-point stable iff the system in (A23) has three negative eigenvalues. This condition held in all runs. The solutions for the transition path were generated by substituting the solutions for $K(t)$, $Z(t)$, and $b(t)$ into (2)-(5) and solving for w , Q , r_k , and r_z with

$$C = \frac{[ar_k^{1-\sigma_1} + (1-a)f^{1-\sigma_1}]^{1/(1-\sigma_1)}}{n},$$

$$f = [ew^{1-\sigma_2} + (1-e)z(r_z/b)^{1-\sigma_2}]^{1/(1-\sigma_2)}.$$

Space does not allow us to present detailed solutions for Models 2, 3, and 4. These models are more complicated and their solutions involve more algebra. The solution procedure, however, is essentially the same as in Model 1.

Appendix B: The Singularity Scenario

When robots and human labor are perfect substitutes in Model 1, the economy achieves sustained, endogenous growth. To solve this case, we set σ_1 equal to unity and work with the firm's production function⁴⁶

$$Q = nK^\theta(L + bZ)^{1-\theta}. \quad (\text{B1})$$

On the balanced growth path, ratios of variables converge to stationary values. Accordingly, we scale all variables by K . Let $q \equiv Q/K$, $\ell \equiv L/K$, $z \equiv Z/K$, $i_z \equiv I_z/K$, $i_k \equiv I_k/K$. and $g_k \equiv \dot{K}/K$. Then

$$q = n(\ell + bz)^{1-\theta} \quad (\text{B2})$$

and the new versions of the private agent's budget constraint and the accumulation equations for K and Z are

$$g_k = i_k - \delta, \quad (\text{B3})$$

$$\begin{aligned} \dot{Z}/K &= i_z - \delta, \\ \implies \dot{z} &= i_z - (\delta + g_k)z. \end{aligned} \quad (\text{B4})$$

Transforming the Euler equations for investment involves a little more work. First rewrite equations (10) and (11) as

$$v_k \frac{\dot{I}_k}{K} = \psi_k \left(\rho + \delta + \frac{\dot{c}/K}{\tilde{c}\tau} \right) + \frac{v_k}{2} (i_k - \delta)^2 - r_k, \quad (\text{B5})$$

$$\frac{v_z}{z} \frac{\dot{I}_z}{K} = \psi_z \left(\rho + \delta + \frac{\dot{c}/K}{\tilde{c}\tau} \right) + \frac{v_z}{2} \left(\frac{i_z}{z} - \delta \right)^2 - r_z, \quad (\text{B6})$$

⁴⁶For simple production functions, there are no advantages to working with the firm's cost function. Also, the mechanics of solving endogenous growth models with Cobb-Douglas technology are familiar from the existing literature.

where $\tilde{c} \equiv c/K$, $\psi_k \equiv 1 + v_k(i_k - \delta)$, and $\psi_z \equiv 1 + v_z(i_z/z - \delta)$, and

$$\begin{aligned} r_k &= Q_K = n\theta(\ell + bz)^{1-\theta}, \\ r_z &= Q_Z = \frac{bn(1-\theta)}{(\ell + bz)^\theta}. \end{aligned}$$

Now

$$\begin{aligned} \dot{I}_k/K &= di_k/dt + i_k g_k, \\ \dot{I}_z/K &= di_z/dt + i_z g_k, \\ \dot{c}/K &= d\tilde{c}/dt + \tilde{c} g_k. \end{aligned}$$

Substitute these expressions into (B5) and (B6):

$$v_k \left(\frac{di_k}{dt} + i_k g_k \right) = \psi_k \left(\rho + \delta + \frac{d\tilde{c}/dt}{\tilde{c}\tau} \right) + \frac{g_k}{\tau} + \frac{v_k}{2} (i_k - \delta)^2 - r_k, \quad (\text{B7})$$

$$\frac{v_z}{z} \left(\frac{di_z}{dt} + i_z g_k \right) = \psi_z \left(\rho + \delta + \frac{d\tilde{c}/dt}{\tilde{c}\tau} \right) + \frac{g_k}{\tau} + \frac{v_z}{2} \left(\frac{i_z}{z} - \delta \right)^2 - r_z. \quad (\text{B8})$$

The Long-Run Equilibrium Growth Rate

The economy converges asymptotically to a balanced growth path where $di_k/dt = di_z/dt = d\tilde{c}/dt = \dot{z} = \ell = 0$ and $i_k - \delta = i_z/z - \delta = g_k$. Imposing these conditions in (B2), (B7), and (B8) gives

$$q = (bz)^{1-\theta}, \quad (\text{B9})$$

$$v_k g_k (g_k + \delta) = \psi_k (\rho + \delta + g_k/\tau) + \frac{v_k}{2} g_k^2 - r_k, \quad (\text{B10})$$

$$v_z g_k (g_k + \delta) = \psi_z (\rho + \delta + g_k/\tau) + \frac{v_z}{2} g_k^2 - r_z, \quad (\text{B11})$$

with

$$\begin{aligned} r_k &= n\theta(bz)^{1-\theta}, \\ r_z &= \frac{nb(1-\theta)}{(bz)^\theta}. \end{aligned}$$

Equations (B9)-(B11) can be solved for q , g_k , and z as a function of b . Since z and q are constant in the long run, $g_z = g_k = g_q = g$. After replacing g_k with g and combining terms in (B10) and (B11), we have

$$g \left\{ 1 + \tau v_k \left[\frac{g}{2\tau} (2 - \tau) + \rho \right] \right\} = \tau (r_k - \rho - \delta), \quad (\text{B12})$$

$$g \left\{ 1 + \tau v_z \left[\frac{g}{2\tau} (2 - \tau) + \rho \right] \right\} = \tau (r_z - \rho - \delta). \quad (\text{B13})$$

The Transition Path

For

$$b_o = \frac{\rho + \delta}{\theta_L},$$

$r_z = r_k = \rho + \delta$ and $g = 0$ at the initial equilibrium. Any small increase in b then leads to permanent growth. The transition path is governed by (B4), (B7), (B8), the law of motion for b

$$\dot{b} = s(\bar{b} - b), \quad s > 0 \quad (12)$$

and

$$\dot{\ell} = -\ell g_k, \quad (B14)$$

a self-contained system of five differential equations in z , i_k , i_z , b , and ℓ . (r_z and r_k depend on b , z , ℓ , while the private agent's budget constraint links $d\tilde{c}/dt$ to the variables in the system). The solution procedure is the same as in Appendix A.⁴⁷

Appendix C: Model 4

Model 4 is Model 3 plus a second sector that produces a non-automatable good using capital, skilled labor, and unskilled labor. The factor demands and zero-profit conditions in the two sectors are

$$K_1 = C_h^1 h_r^1 Q_1, \quad (C1)$$

$$Z = C_f^1 f_z Q_1 / b, \quad (C2)$$

$$L_1 = C_f^1 f_w Q_1, \quad (C3)$$

$$S_1 = C_h^1 h_w^1 Q_1, \quad (C4)$$

$$K_2 = C_h^2 g_r Q_2, \quad (C5)$$

$$L_2 = C_w^2 Q_2, \quad (C6)$$

$$S_2 = C_h^2 g_w Q_2 \quad (C7)$$

$$1 = C^1(r_{k1}, w_s, w, r_z/b), \quad (C8)$$

$$p_2 = C^1(r_{k2}, w_s, w), \quad (C9)$$

where p_2 is the relative price of good 2 and

⁴⁷Two points should be noted. First, the adjustment cost terms do not drop out in the linearized system because $i_k - \delta = i_z/z - \delta = g \neq 0$ on the balanced growth path. Second, the linearized version of (B14) is $\dot{\ell} = -g\ell$.

$$\begin{aligned}
C^1 &= \frac{[ah^{1-\sigma_1} + (1-a)f^{1-\sigma_1}]^{1/(1-\sigma_1)}}{n_1}, \\
f &= [ew^{1-\sigma_2} + (1-e)(r_z/b)^{1-\sigma_2}]^{1/(1-\sigma_2)}, \\
h &= [jr_{k1}^{1-\sigma_3} + (1-j)w_s^{1-\sigma_3}]^{1/(1-\sigma_3)}, \\
C^2 &= \frac{[ug^{1-\beta_1} + (1-u)w^{1-\beta_1}]^{1/(1-\beta_1)}}{n_2}, \\
g &= [vr_{k2}^{1-\beta_3} + (1-v)w_s^{1-\beta_3}]^{1/(1-\beta_3)}.
\end{aligned}$$

Wages adjust to equate demand to the fixed supplies of low- and high-skill labor:

$$L_1 + L_2 = L, \quad (\text{C10})$$

$$S_1 + S_2 = S. \quad (\text{C11})$$

Capitalists and skilled labor belong to an extended family that chooses c , I_{k1} , I_{k2} , and I_z to maximize

$$U = \int_0^\infty \frac{c^{1-1/\tau}}{1-1/\tau} e^{-\rho t}, \quad (\text{C12})$$

subject to

$$\begin{aligned}
pc + I_z + I_{k1} + I_{k2} &= r_{k1}K_1 + r_{k2}K_2 + r_zZ + w_sS - \frac{v_k}{2} \left(\frac{I_{k1}}{K_1} - \delta \right)^2 K_1 \\
&\quad - \frac{v_k}{2} \left(\frac{I_{k2}}{K_2} - \delta \right)^2 K_2 - \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 Z, \quad (\text{C13})
\end{aligned}$$

$$\dot{K}_1 = I_{k1} - \delta K_1, \quad (\text{C14})$$

$$\dot{K}_2 = I_{k2} - \delta K_2, \quad (\text{C15})$$

$$\dot{Z} = I_z - \delta Z, \quad (\text{C16})$$

where

$$\begin{aligned}
c &= [m^{1/\epsilon} c_1^{(\epsilon-1)/\epsilon} + (1-m)^{1/\epsilon} c_2^{(\epsilon-1)/\epsilon}]^{\epsilon/(\epsilon-1)}, \\
p &= [m + (1-m)p_2^{1-\epsilon}]^{1/(1-\epsilon)}.
\end{aligned}$$

c is a CES aggregate of the two consumer goods, with substitution elasticity ϵ . The exact consumer price index is p and the automatable good serves as the numeraire. Good 2 is a pure consumption good while good 1 may be used either for consumption or investment. Hence p multiplies c but not I_{k1} or I_{k2} in the budget constraint (C13).

The solution to the optimization problem in (C12) - (C16) generates three Euler equations for I_z , I_{k1} , and I_{k2} :

$$\frac{v_k}{K_1} \dot{I}_{k1} = [1 + v_k(I_{k1}/K_1 - \delta)] \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_k}{2} \left(\frac{I_{k1}}{K_1} - \delta \right)^2 - r_{k1}, \quad (\text{C17})$$

$$\frac{v_k}{K_2} \dot{I}_{k2} = [1 + v_k(I_{k2}/K_2 - \delta)] \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_k}{2} \left(\frac{I_{k2}}{K_2} - \delta \right)^2 - r_{k2}, \quad (\text{C18})$$

$$\frac{v_z}{Z} \dot{I}_z = [1 + v_z(I_z/K_z - \delta)] \left(\rho + \delta + \frac{\dot{c}}{c\tau} \right) + \frac{v_z}{2} \left(\frac{I_z}{Z} - \delta \right)^2 - r_z. \quad (\text{C19})$$

Finally, the model is closed by the market-clearing condition for the non-automatable good. Since low-skill workers do not save,

$$Q_2 = (1 - m) \left(\frac{p_2}{p} \right)^{-\epsilon} (c + \omega L), \quad (\text{C20})$$

where $\omega \equiv w/p$.⁴⁸

⁴⁸Capitalists, low- and high-skill labor have identical homothetic preferences, qua consumers. Demand depends therefore only on aggregate consumption, not the distribution of consumption spending across agents.

References

- Acemoglu, D. (2002), ‘Technical change, inequality, and the labor market’, *Journal of economic literature* **40**(1), 7–72.
- Acemoglu, D. and Autor, D. (2011a), ‘Skills, tasks and technologies: Implications for employment and earnings’, *Handbook of labor economics* **4**, 1043–1171.
- Acemoglu, D. and Autor, D. (2011b), ‘Skills, tasks and technologies: Implications for employment and earnings’, *Handbook of labor economics* **4**, 1043–1171.
- Acemoglu, D. and Restrepo, P. (2016a), The race between machine and man: Implications of technology for growth, factor shares and employment, Technical report, National Bureau of Economic Research.
- Acemoglu, D. and Restrepo, P. (2016b), ‘Robots and jobs: Evidence from u.s. labor markets’, *unpublished*.
- Autor, D. (2014), ‘Skills, education, and the rise of earnings inequality among the ‘other 99 percent.’’, *Science* **344**, 843–851.
- Autor, D. H. (2015), ‘Why are there still so many jobs? the history and future of workplace automation’, *The Journal of Economic Perspectives* **29**(3), 3–30.
- Autor, D. H., Katz, L. F. and Krueger, A. B. (1998), ‘Computing Inequality: Have Computers Changed the Labor Market?’, *The Quarterly Journal of Economics* **113**(4), 1169–1213.
- Blankenau, W. F. and Cassou, S. P. (2011), ‘Industry estimates of the elasticity of substitution and the rate of biased technological change between skilled and unskilled labour’, *Applied Economics* **43**(23), 3129–3142.
- Brynjolfsson, E. and McAfee, A. (2014), *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company.
- Chirinko, B. and Mallick, D. (2017), ‘The substitution elasticity, factor shares, long-run growth, and the low-frequency panel model’, *American Economic Journal: Macroeconomics* **9**(4), 224–253.
- Chirinko, R. S. (2008), ‘[sigma]: The long and short of it’, *Journal of Macroeconomics* **30**(2), 671–686.
- Chui, M., Manyika, J. and Miremadi, M. (2015), ‘Four fundamentals of workplace automation’, *McKinsey Quarterly* pp. 1–9.
- DeCanio, S. J. (2016), ‘Robots and humans—complements or substitutes?’, *Journal of Macroeconomics* **49**, 280–291.
- Economist (2016), Call centres—the end of the line!, Technical report, The Economist.
- Fallon, P. R. and Layard, P. R. G. (1975), ‘Capital-skill complementarity, income distribution, and output accounting’, *Journal of political economy* **83**(2), 279–302.
- Ford, M. (2015), *Rise of the Robots: Technology and the Threat of a Jobless Future*, Basic Books.

- Freeman, R. B. (2015), ‘Who owns the robots rules the world’, *IZA World of Labor* .
- Frey, C. B. and Osborne, M. A. (2017), ‘The future of employment: how susceptible are jobs to computerisation?’, *Technological Forecasting and Social Change* **114**, 254–280.
- Grace, K., Salvatier, J., Dafoe, A., Zhang, B. and Evans, O. (2017), When will ai exceed human performance? evidence from ai experts, Technical report.
- Griliches, Z. (1969), ‘Capital-skill complementarity’, *The review of Economics and Statistics* pp. 465–468.
- Hamermesh, D. (1993), *Labor Demand*, Princeton University Press, Princeton, NJ.
- Hemous, D. and Olsen, M. (2015), The rise of the machines: Automation, horizontal innovation and income inequality, 2015 Meeting Papers 456, Society for Economic Dynamics.
- Johnson, G. E. (1997), ‘Changes in earnings inequality: the role of demand shifts’, *The Journal of Economic Perspectives* **11**(2), 41–54.
- Karabarbounis, L. and Neiman, B. (2014), ‘The global decline of the labor share’, *The Quarterly Journal of Economics* **129**(1), 61–103.
- Klump, R., McAdam, P. and Willman, A. (2007), ‘Factor substitution and factor-augmenting technical progress in the united states: a normalized supply-side system approach’, *The Review of Economics and Statistics* **89**(1), 183–192.
- Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V. and Violante, G. L. (2000), ‘Capital-skill complementarity and inequality: A macroeconomic analysis’, *Econometrica* **68**(5), 1029–1053.
- Nordhaus, W. D. (2015), Are we approaching an economic singularity? information technology and the future of economic growth, Technical Report 21547, National Bureau of Economic Research.
- Novalés, A., Fernández, E. and Ruíz, J. (2008), *Economic growth: theory and numerical solution methods*, Springer Science & Business Media.
- Oberfeld, E. and Raval, D. (2014), Micro data and macro technology, Technical Report 20452, National Bureau of Economic Research.
- Papageorgiou, C. and Saam, M. (2008), ‘Two-level ces production technology in the solow and diamond growth models’, *The Scandinavian Journal of Economics* **110**(1), 119–143.
- Rader, T. (1968), ‘Normally, factor inputs are never gross substitutes’, *Journal of Political Economy* **76**(1), 38–43.
- Raval, D. (2011), Beyond Cobb-Douglas: Estimation of a CES Production Function with Factor Augmenting Technology, Working Papers 11-05, Center for Economic Studies, U.S. Census Bureau.
- Rodrik, D. (2016), ‘Premature deindustrialization’, *Journal of Economic Growth* **21**(1), 1–33.
- Sachs, J. D., Benzell, S. G. and LaGarda, G. (2015), Robots: Curse or blessing? a basic framework, Technical Report 21091, National Bureau of Economic Research.

Sachs, J. D. and Kotlikoff, L. J. (2012), Smart machines and long-term misery, Technical Report 18629, National Bureau of Economic Research.

Summers, L. (2016), 'A disaster is looming for american men', *Washington Post*, September 29 .

The Atlantic (2016), The next industrial revolution could put millions out of work, Technical report, TheAtlantic.

The South China Morning Post (2015), "foxconn's foxbot army close to hitting the chinese market, on track to meet 30 percent automation target, Technical report.

The White House (2016), Artificial intelligence, automation, and the economy, Technical report.

World Bank (2016), *World Development Report 2016*, The World Bank, Washington, D.C.