

Robots Are Us: Some Economics of Human Replacement*

Seth G. Benzell¹, Laurence J. Kotlikoff², Guillermo LaGarda³, and Jeffrey D. Sachs⁴

¹Initiative on the Digital Economy, MIT

²Department of Economics, Boston University

³Inter-American Development Bank

⁴The Earth Institute, Columbia University

This Version: January 2019

Abstract

Will smart machines do to humans what the internal combustion engine did to horses – render them obsolete? If so, can putting people out of work or, at least, good work leave them unable to buy what smart machines produce? Our model's answer is yes. Over time and under the right conditions, today's supply reduces tomorrow's demand, leaving everyone worse off in the long-run. Carefully crafted redistribution policies can prevent such immiserating growth. But blunt policies, such as limiting intellectual property rights or restricting labor supply, can make matters worse.

*We thank Erik Brynjolfsson, Hernando Zuleta, Dilip Mookherjee, Dietrich Vollrath, Pascual Restrepo, Stephen Terry, John Seater and Kevin Cooke for useful comments. Contact corresponding author at sbenzell@mit.edu

1 Introduction

Whether it's bombing our enemies, steering our planes, fielding our calls, rubbing our backs, vacuuming our floors, driving our taxis or playing Jeopardy, it's hard to think of hitherto human tasks that smart machines can't already or won't soon do. Few smart machines look human. But all combine brains and brawn – artificial intelligence (AI) and physical capital. And they all share one creator – us. Indeed, via supervised machine learning, we are now teaching smart machines to pick our brains.

Will self-precipitated human replacement - the production by ourselves of substitutes for ourselves - deliver an economic utopia with smart machines satisfying our every material need? Or will our self-generated redundancy leave us earning too little to afford what smart machines produce? This paper simulates one vision of human replacement in a bare bones, overlapping generations (OLG) model. It features two types of workers consuming two goods for two periods. Though simple, our model admits a wide range of dynamic reactions to robots, some quite unpleasant.

The model features high- and low-tech workers. Both work full time, but only when young. There are two goods – corn and prayer. Corn is produced with capital and code. Code is the sum of old plus new code. New code is produced exclusively by high-tech workers.¹ Prayers are produced with a combination of high- and low-tech labor. High-tech workers are mobile between sectors. Although their total compensation is the same in each sector, the compensation to coders comes in two forms – a rent for the immediate use of their code and sale of rights to the code's future use.

Code references all rules, instructions, and methods as well as explicit software that help generate output from given inputs. Conceptually, code is measured in efficiency units, not literal lines of instructions, as fewer lines are often more efficient. Old code needs to be retained, maintained, and updated, but is otherwise a perfect substitute for new code. Hence, the stock of old code undermines the demand for new code and, thus, for coders.

The potential for old code to obsolesce new coders is illustrated by *Junior* – 2013's World Computer Chess Champion. Junior can beat all living and, arguably, all future human beings. Consequently, Junior's old code has put chess

¹In the U.S., computer and mathematical occupational income accounted for 5.25 percent of total wages in 2016 (Bureau of Labor Statistics, 2017). But a much larger share of compensation accrues to those engaged in creating code broadly defined.

programmers out of business insofar as their business is beating humans at chess. Junior is, of course, a very smart machine – a robot that combines capital and code. And like other robots, Junior’s code reflects past human brain power. In this sense, Junior as well as all other robots *are us*.

We capture the ability of old code to compete with new technical workers via a code retention rate, denoted δ and examine the model’s reaction to a rise in δ arising from, say, the invention of the silicon chip. The obverse of code retention is code depreciation, equal to $1 - \delta$. Code depreciation can reflect the need to maintain and update code, including recalibrating or retraining to current economic, technological and market conditions. Hands on adjustment of software or further supervised machine learning can play a role in such recalibration, providing a continuous need for coders. If conditions change slowly or the automated system can recalibrate itself, such as by using unsupervised or reinforcement learning, new coders may rarely be needed (corresponding to a high δ).

The response to a rise in the code retention rate is a tech boom that raises the demand for new code and coders. The resulting rise in high-tech and, indeed, low-tech workers’ earnings engenders more national saving and capital formation, reinforcing the boom. But over time, as the stock of legacy code grows, the demand for new code and, thus, coders falls. This depresses the wages of high-tech workers, leading more to migrate to the prayer sector. Depending on the substitutability of high- and low-tech workers in producing prayers, low-tech workers can see their wages rise or fall.

The eventual decline in wages of high- and, potentially, low-tech workers limits what the young save and invest. This can mean less physical capital available for future use. To simplify our model, we consider Cobb-Douglas (CD) preferences in which the share of wages saved by the young each period is insensitive to the interest rate. Positing more realistic, i.e., lower intertemporal elasticity of substitution, preferences would increase the likelihood of immiseration.²

If the capital stock falls by enough, the economy ends up producing less output notwithstanding its higher level of technology. This outcome represents a dynamic counterpoint to Say’s Law as extra short-run supply of one input (code) reduces the long-run demand for, and thus supply of, another (capital). This is the sense in which supply reduces demand. Thus, one of the novelties of

²CD preferences entail an elasticity of substitution of 1 and no saving response to interest rates. Elasticities of less than 1, which appear empirically relevant (see Guvenen2006reconciling) would lead the young to save a smaller share of their wages in response to a higher interest rate. This reflects income effects dominating substitution effects.

our model is that it studies not just how automation may arise and impact the demand for labor. It also considers the impact of automation on saving and the supply of capital.

Our model’s potential for immiserating growth differs from (Sachs and Kotlikoff 2012) and (Sachs, Benzell and LaGarda 2015). In those papers, robots of exogenously specified productivity are treated as capital that can substitute for the labor of young savers. In our model, capital and the output of high-tech labor, namely code, are complements. So the source of immiseration is not machines (i.e., physical capital) replacing workers. Instead, it’s current high-tech workers producing a durable input (code), that substitutes for the supply of code by future high-tech workers.

This said, the selfish OLG framework is common and crucial in the three studies. It permits the reduction in workers’ wages to translate into lower saving and reduced capital formation despite the concomitant rise in the interest rate. If the capital stock falls by enough, the economy’s long run can feature less output and lower welfare notwithstanding improved technology. This potential for immiserating technological progress would not arise in the infinitely-lived single-agent model. In such a model, national saving is driven by the difference in interest and time preference rates, not the intergenerational distribution of resources. Indeed, in such models, capital shortages lead to higher interest rates, which immediately induce greater saving and capital formation. ³

Even in our OLG model, immiseration is not inevitable. It depends on saving preferences as well as the ability to substitute software (code) for hardware. If young workers save at high rates or if software and hardware can readily substitute for one another, immiseration is less likely to arise.

Our main findings assume that code is excludable and rival in its use. This is a reasonable position if software needs to be specialized to each application. But when the same code can not just be copied for free, but also costlessly adapted to multiple applications, it resembles a public good. After considering the case that code is excludable and rival, we take up cases in which code is non-rival and, potentially, non-excludable. This lets us evaluate different intellectual property (IP) regimes, which is of interest to policymakers worried about the effects of emerging AI technologies on market concentration and labor demand (Benzell and Brynjolfsson 2018a). We find that a rise in code retention can still immiserate the economy when code is non-rival, particularly when it is

³There is strong evidence against the operative intergenerational altruism needed to justify the single-agent model. See, for example, Altonji, Hayashi, and Kotlikoff (1992,1997), Hayashi, Altonji and Kotlikoff (1996), and Abel and Kotlikoff (1994).

non-excludable. On the other hand, AI intellectual property protection lowers market concentration by encouraging more entry by small firms.

Our appendix considers whether long-run immiseration can arise when agents have more control over technological change as in Acemoglu (2002) and Acemoglu and Restrepo (2018). In this extension, firms can create either labor-substituting or capital-substituting code. We find that immiseration is still possible provided labor-substituting code technology is more advanced than capital-substituting code technology.⁴

The next section considers the long-standing concern that technical change can hurt workers and reviews a small portion of the relevant recent literature. Section 3 presents the model and demonstrates the potential for immiseration analytically. Section 4 illustrates, numerically, the model’s surprising range of outcomes, including immiseration, arising from a rise in code retention. Section 5 shows that code rivalry and property rights can make a major difference. Section 6 considers evidence supporting the model’s main predictions, including a decline in labor’s share. Section 7 concludes.

2 Background and Literature Review

Concern about new technology dates at least to Ned Ludd’s destruction of two stocking frames in 1779 near Leicestershire, England.⁵ Sixty-five years later, Marx (1867) restated Ned Ludd’s warning about machines replacing humans: “Within the capitalist system all methods for raising the social productivity of labour are put into effect at the cost of the individual worker.” Keynes (1933) also raised technology’s potential for job destruction, writing in the midst of the Great Depression that “We are being afflicted with a new disease ..., namely technological unemployment.”⁶

In the fifties and sixties, with employment high and wage growth rapid, Keynes’

⁴Previous versions of this paper included an extension where firms choose between a range of production technologies. In such a model, immiseration remains possible. Additionally, Kondratiev-type business cycles, like those seen in our non-rival and non-excludable code scenario, can arise.

⁵Ludd, a weaver, was whipped for indolence before wreaking revenge on the machines. More than three decades later, in 1812, 150 armed workers – self-named Luddites – marched on a textile mill in Huddersfield, England. Their purpose was to smash equipment. The British army promptly killed 19. Later that year the British Parliament passed *The Destruction of Stocking Frames Act*, authorizing death for those vandalizing machines. Nonetheless, Luddite riots continued for several years, eventuating in 70 hangings.

⁶But Keynes called this “only a temporary phase of maladjustment” and predicted a future of leisure and plenty one hundred years hence. His contention that short-term pain facilitates long-term gain reinforced Schumpeter’s 1942 encomium to “creative destruction”.

and Schumpeter’s views held sway. Those raising concerns about technology were regularly dismissed as “Luddites.” But in recent years the swift loss of all manner of jobs to smart machines has led economists to rethink Luddism. Erik Brynjolfsson and Andrew McAfee (2014)’s book, *Race Against the Machine*, and Aghion, et. al.’s (2017) recent paper, *Artificial Intelligence and Economic Growth*, are just two examples of a burgeoning literature. Brynjolfsson and McAfee emphasize the ongoing role of machines in changing relative compensation across occupations, increasing inequality and decreasing labor force participation. Aghion and co-authors connect AI to the recent decline in labor’s output share. But they also view AI as potentially just a new form of automation – one likely subject to *Baumol’s Cost Disease*, with long-run outcomes ultimately determined not by what AI can do, but what it can’t do.

The long run can take a long time. Moreover, a key message of OLG models is that where the economy ends up depends on how it gets there. Hence, the focus by Autor, Levy, and Murnane (2003), Acemoglu and Autor (2011), Autor and Dorn (2013) and others on how smart machines are impacting current employment and wage trends is well placed. Each finds significant outsourcing of middle-skilled workers by smart machines. Goos, Manning, and Salomons (2010) offer supporting evidence for Europe. Margo (2013) points to similar *labor polarization* during the early stages of America’s industrial revolution. Many economists are now connecting robotization to the ongoing decline in labor’s share of output. Hemous and Olson (2014) is an example. Their model has capital substituting for low-tech and complementing high-tech labor and explains trends in labor’s share and income inequality since the 1960s.

The deleterious labor-market impact of smart machines is not without its skeptics. Mishel, Shierholz, and Schmitt (2013) argue that ‘robots’ can’t explain post-1970’s U.S. job polarization given the observed timing of changes in relative wages and employment. Autor (2015) is another skeptic, at least over the long term. He points out that the automobile displaced equestrian drivers but introduced myriad occupations for humans in the auto and other industries. Autor’s argument applies to humans, but not to horses, who suffered permanent job loss. When it comes to AI, the question is whether today’s humans are yesterday’s horses.

Our model features an endogenous technological and growth response to an exogenous technology shock – the aforementioned rise in the rate of code retention). Hence, our study connects, to a degree, to the endogenous growth literature, whose major contributions include Schumpeter (1939), Arrow (1962), Uzawa (1965), Sidrauski (1967), Lucas (1988), Romer (1990), Rebelo (1991),

Ortigueira and Santos (1997), Zeira (1998), Acemoglu (1998), Howitt (1999), Zuleta (2008), and Peretto and Seater (2013).

As for modeling automation, economists have taken a range of approaches. Zeira (1998) posits the availability, at a cost, of labor-substituting machines and shows that countries with high labor costs and low interest rates will industrialize more rapidly than others. This process produces a dispersion in global per capita income. Zuleta (2004, 2008) considers the choice not of labor-replacing machines, but of the degree of capital intensity. As in Zeira (1998), rich economies expand relative to poor economies, which can't afford to increase their degree of capital intensity. Zeira (2004)'s model is OLG. Hence, he too finds that the decline in labor income can cause economic problems over time. In his case, it's the inability to achieve long-run growth absent the presence of bequests.

Acemoglu (1998) features firms that invest in technology that differentially raise the productivity of their least expensive inputs. Rourke, et. al. (2013) examines 18th and 19th technological change in England with special focus on the skill premium. His model, which is similar to Acemoglu's (1998), appears capable of matching the trend in the skill premium over the period.

Peretto and Seater (2013) extend Zuleta (2008). They consider monopolistically competitive firms that invest in particular technologies depending on their relative costs. In their model, firms may specialize in the use of one technology or produce with multiple technologies.

Acemoglu and Restrepo (2017) endogenize the automation of labor as well as the invention of new labor-intensive products. The former (latter) occurs to a greater (lessor) degree when wages are high (low). They show that balanced growth can arise with the demand for labor to perform innovative tasks offsetting job loss due to automation.

Guerreiro, Rebelo, and Teles (2017), Costinot and Werning (2018) and others have examined the use of fiscal policy to offset technology-induced redistribution. Costinot and Werning (2018) tax technology, via a constrained set of taxes, in a complete market with a continuum of agents. The authors find that when it is optimal to tax new technologies, such taxes will likely to be small. Our framework is different. In particular, it doesn't admit contracting between current and future generations. Still, as in Sachs, Benzell and Lagarda (2015), tax and transfer policies could be added to our model to maintain the capital stock and preserve the intergenerational distribution of welfare in the face of AI

innovation.

3 The Model

Agents consume two products – corn and prayers.⁷ Corn, which can be consumed or invested, is produced using capital and code. The CES function governing corn production can be viewed as a smart machine or robot since it combines physical capital and code. Prayers are produced via a CES function of low- and high-tech labor.⁸ Prayers are ephemeral. They are consumed when produced.⁹

Code is durable, with the stock of code equaling the sum of new and existing code. New code is created by high-tech workers not working in the prayer sector. Old code requires maintenance, retention, and updating, which we treat as a form of depreciation. The numbers of high- and low-tech workers remain fixed through time.¹⁰ Both types of workers live and consume for two periods, but work only when young.

Production

Time- t production of corn (goods), Y_t , and prayers (services), S_t , satisfy (1) and (2),

$$Y_t = D_Y [\alpha (K_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1 - \alpha) (A_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}}]^{\frac{\varepsilon_y}{\varepsilon_y - 1}}, \quad (1)$$

$$S_t = D_S [\gamma (H_{S,t})^{\frac{\varepsilon_s - 1}{\varepsilon_s}} + (1 - \gamma) (G_t)^{\frac{\varepsilon_s - 1}{\varepsilon_s}}]^{\frac{\varepsilon_s}{\varepsilon_s - 1}}, \quad (2)$$

where $H_{S,t}$ is the number of high-tech workers in the prayer sector and G_t references low-tech workers. D_S and D_Y are total factor productivity terms, γ and α are CES share parameters, and ε_y and ε_s are CES elasticities. The stock of code A_t grows according to,

$$A_t = \delta A_{t-1} + z H_{A,t}, \quad (3)$$

⁷Frey and Osborne (2013) identify the priesthood, psychotherapy and coaching as among the occupations least subject to automation.

⁸Adding labor to the production function for corn, or capital to the production for prayer, would not alter our qualitative findings.

⁹Corn and prayer are stand-ins for automatable and non-automatable products. They are also rough analogues to mass-produced goods and interpersonal or artisanal services.

¹⁰Our model can, however, accommodate long-run balanced growth arising from population growth or labor-augmenting technological change.

where the “depreciation” factor is $\delta \in [0, 1]$. Higher δ means that legacy code is useful for longer.¹¹ $H_{A,t}$ stands for the supply of high-tech coders and z is their productivity coefficient.¹²

The demands for code, high-tech workers, and capital satisfy

$$\max_{K_t, A_t} Y_t(A_t, K_t) - m_t A_t - r_t K_t, \quad (4)$$

where corn is the numeraire, m_t is the rental rate for code, and r_t is the interest rate. Factor demands for prayer reflect,

$$\max_{H_{S,t}, G_t} q_t S_t(H_{S,t}, G_t) - w_t^G G_t - w_t^H H_{S,t}, \quad (5)$$

where q_t is the price of prayer, w_t^H is a high-tech worker’s wage in the prayer sector, and w_t^G is a low-tech worker’s wage.

Factor prices satisfy

$$w_t^H = q_t D_S [\gamma (H_{S,t})^{\frac{\varepsilon_s - 1}{\varepsilon_s}} + (1 - \gamma) (G_t)^{\frac{\varepsilon_s - 1}{\varepsilon_s}}]^{\frac{1}{\varepsilon_s - 1}} [\gamma (H_{S,t})^{-\frac{1}{\varepsilon_s}}], \quad (6)$$

$$w_t^G = q_t D_S [\gamma (H_{S,t})^{\frac{\varepsilon_s - 1}{\varepsilon_s}} + (1 - \gamma) (G_t)^{\frac{\varepsilon_s - 1}{\varepsilon_s}}]^{\frac{1}{\varepsilon_s - 1}} [(1 - \gamma) (G_t)^{-\frac{1}{\varepsilon_s}}], \quad (7)$$

$$r_t = D_Y [\alpha (K_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1 - \alpha) (A_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}}]^{\frac{1}{\varepsilon_y - 1}} [\alpha (K_t)^{-\frac{1}{\varepsilon_y}}], \quad (8)$$

and

$$m_t = D_Y [\alpha (K_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}} + (1 - \alpha) (A_t)^{\frac{\varepsilon_y - 1}{\varepsilon_y}}]^{\frac{1}{\varepsilon_y - 1}} [(1 - \alpha) (A_t)^{-\frac{1}{\varepsilon_y}}]. \quad (9)$$

¹¹Some of our simulations assume a depreciation rate of 30 percent per period, where a period is roughly 30 to 40 years. This corresponds to a typical company needing to replace approximately 1 percent of its code base annually to maintain the same level of output. The actual rate of code depreciation in the economy is unclear. The IRS allows for a 3 year useful lifespan for licensed software. For software developed in house or purchased bespoke software, costs can be amortized over a 15 year period (as a section 197 intangible). Software that is bundled with hardware is implicitly assumed to depreciate at the rate of the hardware. On the other hand, many programs created over 50 years ago are still in use, such as those written for older nuclear reactors.

¹²The corn production process can be understood by analogy to a firm whose service is making good chess moves. The firm can improve its service either by increasing the quality of its chess program (increasing its efficiency units of code) or using more capital (computer cores) to exploring winning moves.

Households

Whether high-tech or low-tech, households maximize

$$u = (1 - \phi)[(1 - \kappa)\log c_{y,t} + \kappa \log s_{y,t}] + \phi[(1 - \kappa)\log c_{o,t+1} + \kappa \log s_{o,t+1}], \quad (10)$$

subject to,

$$c_{y,t} + q_t s_{y,t} + \frac{c_{o,t+1} + q_{t+1} s_{o,t+1}}{1 + r_{t+1}} = i_{j,t}, \quad (11)$$

where $c_{y,t}$, $c_{o,t}$, $s_{y,t}$, $s_{o,t}$, are consumption of corn and prayer by the young and old, respectively and $i_{j,t}$ is total resources of group j . For low-tech workers,

$$i_{G,t} = w_t^G. \quad (12)$$

For high-tech workers in the prayer sector,

$$i_{(H,S),t} = w_t^H, \quad (13)$$

and for high-tech workers writing code,

$$i_{(H,A),t} = z(m_t + \delta p_t), \quad (14)$$

where zm_t is revenue from renting out newly produced code and $z\delta p_t$ is the proceeds from the sale of rights to future use of newly produced code.

Households save in the form of capital and code. Capital and code accumulation obeys

$$\phi I_t = K_{t+1} + p_t \delta A_t, \quad (15)$$

where I_t is the total resources of those born in t , ϕ is the saving propensity of the young, and $p_t \delta A_t$ is the value of code retained from the current period. In equilibrium the return to both investments is identical.

Figure 1 summarizes the timing of consumption and saving decisions for high-tech workers in the corn industry.

Demands satisfy,

$$s_{y,t} = \frac{\kappa(1 - \phi)i_{j,t}}{q_t}, \quad (16)$$

$$c_{y,t} = (1 - \kappa)(1 - \phi)i_{j,t}, \quad (17)$$

$$s_{o,t+1} = \frac{1 + r_{t+1}}{q_{t+1}}[\kappa \phi i_{j,t}], \quad (18)$$

and

$$c_{o,t+1} = [1 + r_{t+1}][(1 - \kappa)\phi i_{j,t}]. \quad (19)$$

Equilibrium

Since high-tech workers are mobile between sectors,

$$w_t^H = z(m_t + \delta p_t). \quad (20)$$

Asset-market clearing entails equal returns on capital and code, i.e.,

$$p_t = \sum_{s=t}^{\infty} R_{s+1,t}^{-1} \delta^{s-t} m_{s+1}, \quad (21)$$

where $R_{s,t}$ is the compound interest factor between t and s , i.e.,

$$R_{s,t} = \prod_{j=t}^s (1 + r_j). \quad (22)$$

Finally, equilibrium requires

$$Y_t = C_{y,t} + C_{o,t} + K_{t+1} - K_t, \quad (23)$$

$$H_t = H_{A,t} + H_{S,t}, \quad (24)$$

and

$$S_t = S_{y,t} + S_{o,t}, \quad (25)$$

where C_y , C_o , S_y , S_o , are total consumption demand of corn and payer by the young and old respectively.

The Steady State

If production functions are Cobb-Douglas, the steady state is implicitly defined by the following two equations in $k = \frac{K}{A}$, the capital to code ratio, and q , the relative price of prayer.

$$\begin{aligned}
D_y k^\alpha = & \left[\frac{(1-\phi)(1-\kappa)}{\phi} \right] \left[k + \frac{(1-\alpha)D_y k^\alpha \delta}{1 + \alpha D_y k^{\alpha-1} - \delta} \right] \\
& + (1-\kappa) \left[k + \frac{(1-\alpha)D_y k^\alpha \delta}{1 + \alpha D_y k^{\alpha-1} - \delta} \right] [1 + \alpha D_y k^{\alpha-1}]
\end{aligned} \tag{26}$$

and

$$k + p\delta = \phi \left[z(m + p\delta)H + (1-\gamma)G \left(\frac{\gamma}{z(m+p\delta)} \right)^{\frac{\gamma}{1-\gamma}} (qD_s)^{\frac{1}{1-\gamma}} \right], \tag{27}$$

where,

$$m = (1-\alpha)D_y k^\alpha, \tag{28}$$

$$r = \alpha D_y k^{\alpha-1}, \tag{29}$$

$$p = \frac{(1-\alpha)D_y k^\alpha}{1 + \alpha D_y k^{\alpha-1} - \delta}. \tag{30}$$

The mechanism for immiseration in the model is low wages. Real wages for high-tech workers are lowered when the code to capital ratio is lowered (keeping the price of prayer constant). In the Cobb-Douglas case, we have

$$\frac{dk}{d\delta} = - \frac{aD_y k^{\alpha-1} + \alpha(1-\alpha)[D_y k^{\alpha-1}]^2 - b}{2\alpha(1-\alpha)cD_y^2 k^{2\alpha-3} + (1-\alpha)eD_y k^{\alpha-2}} \tag{31}$$

where $a = [1 + \frac{(1-\kappa)(1-2\alpha)}{\phi}]$, $b = \frac{(1-\kappa)}{\phi}$, $c = 1 - (\alpha + \delta(1-\alpha))$, and $e = (1-\delta)(1-\alpha b) - b(1 - (1-\alpha)\delta)$. One can readily choose reasonable parameter values that make this derivative negative (i.e. such that an increase in code retention technology lowers the capital to code ratio).

Analytical Results in a Simplified Model

Only simplified versions of our model are analytically tractable. As an illustration of the immiseration mechanism, consider the special case of only one good (corn) produced using Leontief technology (i.e. $\kappa = 0$, $G = 0$, $\epsilon_y \rightarrow 0$).¹³ Under the right conditions, a sufficiently large increase in the code retention rate will

¹³Unlike the main model, we also assume no physical capital depreciation.

totally and irreversibly collapse the economy. This will occur either after two periods or at a specific date in the future.

This simplified model exhibits many key features of the full model. These include the potential for a boom-bust and the importance of the saving rate ϕ to the potential for and timing of economic collapse. However, unlike the full model, the economy shifts gears much more abruptly from boom to bust when this occurs.

Assume production satisfies

$$Y_t = D_y \min[(1 - \alpha)A_t, \alpha K_t]. \quad (32)$$

Since all workers are coders and their period-specific labor supply, H , is fixed, A_t is deterministic. Assuming δ is zero prior to time zero and fixed at $\hat{\delta}$ afterwards,

$$A_t = \frac{zH(1 - \hat{\delta}^{t+1})}{1 - \hat{\delta}}. \quad (33)$$

When code is scarce, production is linear in code and the marginal product of code and the wage are positive. The return to capital is zero. The opposite holds when capital is scarce.¹⁴ I.e.,

$$m_t = \begin{cases} D_y(1 - \alpha) & \text{if } (1 - \alpha)A_t < \alpha K_t \\ 0 & \text{if } (1 - \alpha)A_t > \alpha K_t \end{cases}$$

and

$$r_t = \begin{cases} 0 & \text{if } (1 - \alpha)A_t < \alpha K_t \\ D_y \alpha & \text{if } (1 - \alpha)A_t > \alpha K_t. \end{cases}$$

Economic collapse, if it occurs, arises in the period after code switches from scarce to abundant. Once the switch occurs, wages fall to zero, eliminating saving and thus ensuring that capital will be in zero supply in all subsequent periods. Output will be fixed at zero. Let T reference the first date at which code becomes abundant, assuming this occurs. Then $m_t = D_y(1 - \alpha) \forall t < T$

¹⁴We do not consider the implausible knife-edge case where capital and code are precisely balanced.

and $m_t = 0 \forall t \geq T$. The economy permanently collapses at $T + 1$.

The condition for code abundance is $\alpha K_T < (1 - \alpha)A_T$. Inserting equations (13) and (14) and rearranging yields

$$\alpha(\phi H z m_t + p_{T-1} \hat{\delta}(\phi z H - A_{T-1})) < (1 - \alpha)A_T. \quad (34)$$

Note from equation (33) that A_{T-1} is greater than or equal to zH . Hence, the second term on the left-hand-side of equation (34) is negative. Because only a fraction ϕ of newly produced code's value is saved, a positive price of code serves, on balance, to crowd out capital.

Also note that $m_T = 0$ implies $p_{T-1} = 0$ since code is valueless if it is expected to be zero after $T - 1$. Setting $p_{T-1} = 0$ and incorporating formulas for m_t and A_t yields a condition for the economy to collapse at time $T + 1$ as a function of parameters

$$\alpha \phi D_y < \frac{1 - \hat{\delta}^{T+1}}{1 - \hat{\delta}}, \quad (35)$$

As equation (35) indicates, if collapse occurs, it occurs faster the smaller the product of ϕ , α , and D_y , and the larger the value of $\hat{\delta}$. Output rises until the economy collapses since A_t is strictly increasing over time.¹⁵

What happens to welfare? Those born at $t = 0$ through $t = T - 2$ are better off than those born prior to period 0. The reason is their ability to sell their durable code raises their total wage. Those born in $T - 1$ are also better off. They receive the $t = 0$ wage while young, but a high interest rate on their physical capital investments while old. For generations born in $t = T$ or later, the marginal product of code will be zero throughout their lives. They will enjoy no income or consumption.

Equation (35) also specifies the condition for permanent scarcity of code, namely,

$$1 - \frac{1}{\alpha \phi D_y} > \hat{\delta} \quad (36)$$

¹⁵Neither the timing nor occurrence of economic collapse is affected by coder productivity, z .

If code is permanently scarce, then $w_{H,t} = zm_t + \frac{zm_t}{1-\delta} = z(1-\alpha)D_y + \frac{z(1-\alpha)D_y}{1-\delta}$. Hence, so long as $\hat{\delta}$ is sufficiently small (or D_y , ϕ , and α are sufficiently large), a rise in δ leads to an immediate and permanent increase in compensation and welfare.¹⁶ The size of the increase in wages is increasing in δ , z and D_y .

This simplified model also illustrates the role of the elasticity of substitution in the corn sector on the possibility of immiseration. If code and capital were perfect substitutes rather than complements, an increase in δ from zero to positive will never lower the marginal product of code, capital or the wage.

4 Numerical Solution and Simulation Results

The models' main novelty is the inclusion of the stock of code in the production of goods. We calculate the economy's perfect foresight transition path following an immediate and permanent increase in the rate of code retention (decline in the code depreciation rate). The solution is via Gauss-Seidel iteration. First, we calculate the economy's initial and final steady states. This yields initial and final stocks of capital and code. These steady-state values provide, based on linear interpolation, our initial guesses for the time paths of the two input stocks. Next, we calculate associated guesses of the time paths of factor prices as well as the price paths of code and prayer. Step three uses these price paths and the model's demand, asset arbitrage, and labor market conditions to derive new paths of the supplies of capital and code. The new paths are weighted with the old paths to form the iteration's next guesses of capital and code paths. The convergence of this iteration, which occurs to an arbitrarily high degree of precision, implies market clearing in each period.

Simulation Results

When the code retention rate, δ equals zero, corn production is conventional – based on contemporaneous inputs of capital and labor (code writers). But when δ rises, corn production depends not just on capital and current labor, but also, implicitly, on deceased high-tech workers.

¹⁶Compensation of coders jumps immediately to a higher level where it remains indefinitely if the economy never collapses (if the economy does collapse in the future, the present discounted value of code falls as collapse nears). Interestingly, although new coders get paid the same through time and capital income is zero, output rises, albeit at a decelerating rate. In the steady state, total wage income is equal to total output. Along the transition, the high wage is financed by decumulation of physical capital, which is increasingly crowded out.

The increase in δ initially raises the compensation of coders. This draws more high-tech workers into coding, raising high-tech worker compensation in both sectors. In most parameterizations, the concomitant reduction in the supply of prayers raises the price of prayers. Depending on the degree to which high-tech workers complement low-tech workers in producing prayer, the wages of low-tech workers will rise or fall. When the two forms of labor are close substitutes the wages of low-tech workers track those of the high tech. When low-tech workers require high-tech workers to complement them, their wages fall as high-tech workers depart for the corn industry.

The situation of high-tech workers degrades over time. As more durable code comes on line, the marginal productivity of code falls, making new coders increasingly redundant. Eventually the demand for coders is limited to those needed to cover the depreciation of legacy code, i.e., to retain, retrain, maintain, and update the AI. The remaining high-tech workers find themselves working in the prayer sector. The upshot is that high-tech workers can end up potentially earning far less than in the initial steady state.

What about low-tech workers? The price of prayers peaks and then declines thanks to the return of high-tech workers to the sector. This puts downward pressure on low-tech workers' wages and, depending on the complementarity of the two inputs in producing prayer, low-tech workers may also see their wages fall. When low-tech workers are close substitutes for high-tech workers, the boom-bust in high-tech workers' compensation generates a boom-bust in low-tech compensation. In the special case where high and low-tech workers are perfect substitutes their wages move in lock step.

The economy's dynamic reaction to a higher δ operates in part through the impact on capital formation. The initial rise in earnings of at least the high-tech workers can engender more aggregate saving and investment. The increased capital makes code and, thus, high-tech workers more productive. But if the total compensation of workers eventually falls, so too will the saving of the young and the economy's supply of capital. Less capital lowers the marginal productivity of code and raises interest rates. This lowers the price of code and the wages of those who produce it.

We next consider alternative transition paths that emerge based on different assumptions about technology or preferences. We first consider paths that feature immiserating growth. Next we show that the opposite is possible - long-run, welfare improving growth. The third type of path involves a change in the relative income positions of high- and low-tech workers. I.e., we consider paths in which

low-tech workers end up with higher wages than high-tech workers despite the opposite initial situation. Finally, we show that if the code retention shock is accompanied by an increase in corn's share of final consumption, immiseration can be magnified.

Each simulation features an immediate and permanent rise in the code-retention rate. But the dynamic impact of this technological breakthrough depends on the size of the shock and the choice of parameters. In addition presenting the above cases in detail, we perform a sensitivity analysis of outcomes to parameter assumptions.

Illustrating Long-Run Immiseration

In this section, we assume that workers save 20 percent of their wages, i.e., that the current consumption share in the CD utility function is 0.2. This may seem very low, but it entails an initial steady-state marginal product of capital of 1.66. Taking a period as 30 years, this is 1.02 percent on an annual basis. This is far below empirical estimates of the U.S. economy's marginal product of capital (pre-tax, but post-depreciation).¹⁷ Thus saving is, if anything, calibrated to be over abundant in this illustration.

Figure 2 shows that technological growth, namely the code-retention rate, δ , rising from 0 to .7, can have negative long-term consequences.¹⁸ The simulation assumes Cobb-Douglas production of corn and linear production of prayer; i.e., both types of workers are perfect substitutes in producing prayers ($\varepsilon_S = \infty$).

As the top left panel indicates, real national income rises several percent, peaking at 7.8 percent above baseline a period after the shock.¹⁹ But it ultimately declines, ending up 4.2 percent below its initial steady-state value. In the long-run, corn output decreases 28.0 percent.

¹⁷see, for example, (Poterba 1998)

¹⁸Note, given our CD utility function, the economy would not respond in advance to the increase in δ were the increase known in advance.

¹⁹Unless otherwise noted, national income, personal income, and wages are reported in real terms. The price index used is a geometric mean of the relative price of corn and prayer. This is the correct price index to use when utility is Cobb-Douglas. The weights used are their corresponding shares in consumption. The price index is

$$\Pi_t = q_t \frac{q_t(S_{y,t} + S_{o,t})}{q_t(S_{y,t} + S_{o,t}) + C_{y,t} + C_{o,t}} = q_t^\kappa \quad (37)$$

Output of corn and prayer are reported in their own units. Other prices (the price of code, p , interest rate, r , and price of prayer, q) are reported in units of the numeraire (corn) unless otherwise noted. National income is the sum of all payments to workers and capital.

The output of prayer dips after the shock, as workers migrate to lucrative coding jobs, then increases steeply as the automated workers return. The price of prayers does the opposite. While production of prayers is 27.4 percent higher in the long run, prayers, at that point, aren't worth very much at the margin. In fact, the long-run price of prayers is 43.4 percent lower than before the technological breakthrough. Both types of workers are perfect substitutes in the spirituality, which, in our simulations, always employ both. Their common compensation initially jumps 7.6 percent and then falls gradually. In the long run all workers earn 24.8 percent less than was originally the case.

What happens to the welfare of different agents through time? The initial elderly are essentially indifferent to the tech boom. Both interest rates, benefiting them, and the price of code, hurting them, rise slightly. The initial young experience a 14.2 percent rise in lifetime utility, measured as a compensating differential relative to their initial steady-state utility. But those born in the long run are 16.5 percent worse off and national income is 4.2 percent lower.

The top right chart helps explain why good times presage bad times. The stock of code shoots up and stays high. But the stock of capital immediately starts falling. After six periods there is over 50 percent more code, but 65 percent less capital. The marginal product of capital skyrockets, increasing the long-run interest rate 110.2 percent above its initial steady-state value.

The huge long-run decline in the capital stock and associated rise in its marginal product has two causes. First, as just stated, wages, which finance the acquisition of capital, fall almost in half by the implicit competition with deceased workers. Second, the advent of a new asset – durable code – crowds out asset accumulation in the form of capital. When δ rises, all workers immediately enjoy an increase in their compensation. This leads to more saving, but not necessarily more saving in the form of capital. Instead, much of the short-lived extra saving is used to acquire claims to legacy code. Initially, when the stock of code is small, its price is quite high. Later, when the stock of code is large, its price is quite low. But the product of code's price and its quantity are always sufficiently high to crowd out investment in capital.

What happens to labor's share of national income? Initially it rises slightly from 75.0 percent in the initial steady state to 75.7 percent. It then declines, falling to 57.5 percent in the long run. This reflects the higher share of output paid to legacy code. This long-run decline in both physical capital and labor's shares of national income arise in all our simulations.

Illustrating Long-Run Welfare Improvement

As figure 3 shows, the tech boom need not auger long-term economic decline. A higher saving propensity is the key. In the immiserating growth case, we assumed a saving propensity, ϕ , equal to .2. Here, keeping all other parameters fixed, we assume ϕ equals .85. With this far higher saving rate, workers are better off in the long run. However, in the first period, when the code-retention shock hits, the old generation is worse off due to the short-term rise in the price of prayer.

National income peaks in the period after the technological shock, rising 25.6 percent above baseline. But in the long-run, national income is only 16.8 percent higher. This reflects a rise and then fall in the capital stock, but not one that is sufficient to reduce long-run welfare. In the prior simulation the capital stock immediately declined. Here the capital stock temporarily increases 8.7 percent above its initial value.

A short-run rise in both the capital stock and prayer price boost the common wage in the short run and leaves it at roughly its initial value in the long run. After peaking 20.5 percent above its initial value, the wage falls, ending up 2.2 percent lower. The stock of code ends up more than doubling. But the capital stock, notwithstanding the high rate of saving, ultimately declines by 35.8 percent.

The respective increase and decrease in the stocks of code and capital produce a significant rise in the economy's interest rate – 78.0 percent in the long run. Although the common wage of high and low-tech workers decreases slightly from its initial level, the rise in the interest rate permits future generations to consume significantly more.

Why does a high enough saving rate lead the δ shock to increase long-run welfare? The answer is that whatever happens to the stock of code, a higher saving rate entails a higher capital stock and, therefore, less of a decline in wages after code accumulates.

At the end of this section we explore the sensitivity of welfare changes to the combination of the saving propensity and code retention rates. The combination of low saving and high code retention rates lead to the most negative outcomes. Conversely, a high saving preference and high code retention rate lead to the best outcomes.

Illustrating a Reversal in Inequality

If high and low-tech workers are complements in production of prayer, their relative wage and welfare levels can flip. Consider, for example, simulating the model based on table 2's parameters. The results are displayed in figure 4. This simulation assumes production in the prayer sector is Cobb-Douglas. We also assume that the saving propensity, ϕ , equals .7. We normalize high- and low-tech efficiency units such that their wages are the same in the no-shock steady state.

As always occurs, the initial effect for high-tech workers is positive. Indeed, the shock immediately raises the wages of high-tech workers by 26.0 percent. It also immediately lowers the share of high-tech workers in the prayer sector from 50.0 to 39.8 percent. Consequently, in the short run, low-tech workers, who, given our complements in production assumption, need high-tech workers to be productive, see their wages increase only .2 percent.

However, as code accumulates and capital decumulates, high-tech workers earn less and less writing code. In response, they move in great number back to the prayer sector. Ultimately, 67.6 percent of high-tech workers end up working in the prayer sector, earning less than was originally the case. Indeed, in the final steady state, high-tech workers earn 18.1 percent less than in the initial steady state. Low-tech workers, in contrast, earn 10.8 percent more. Interestingly, the low-tech wage peaks two periods after the shock at 23.2 percent above its original value. This rise and fall in the wage of low-tech workers reflects, in part, the rise and fall in the price of prayer.

Although, the wages of high-tech workers end up below their initial values, they, like low-tech workers, experience long-run increases in welfare. But the increase is very small for high-tech workers and very large for low-tech workers.

Increase in Corn's Share in Final Consumption

Our model's assumption that corn and prayer are Cobb-Douglas complements in final consumption is important. Assuming that the shares of spending on automatable versus non-automatable goods are fixed is reasonable, given that there is no strong evidence on the elasticity of substitution between them. However, due to changes in preferences, relative prices, or technology, corn could command a larger or smaller share of final consumption than corn.

This subsection examines the consequences of a change in the corn preference-share parameter, κ , arising in conjunction with the rise in the code retention rate. Figure 5 shows results for the case that κ decreases from .5 to .25 at the same time the code retention rate increases. All other parameters are those in table 2. This alternative assumption has a dramatic and surprising impact on the path of national income. Unlike figure 4’s results, long-run national income falls by 30.3 percent rather than rising. As in previous cases, immiseration is caused by capital decumulation. Capital stocks in this case decrease 34.6 percent in the period immediately after the shock and 85.8 percent in the long-run.

Capital decumulation is greater when κ is larger for three reasons. First, the higher demand for corn increases the share of high-tech workers working as coders. This translates, after one period, into more legacy code and lower labor compensation – the source of saving and capital formation. Second, the increase in immediate corn consumption reduces the amount of corn left over to invest. Third, the higher preference toward corn limits the rise in the price of prayers. This, too, has a negative impact on wages and capital formation.

Sensitivity Analysis

As just demonstrated, the model’s reaction to the δ shock is highly sensitive to parameter values. We now consider this sensitivity in more detail. Figure 6 displays our previous results. Table 3 shows additional findings for other parameter combinations. The table’s baseline simulation (row one) assumes intermediate parameter values. Subsequent rows show the impact of sequentially modifying one parameter. Figure 7 plots the path of national income for each row of the table.

These simulations teach several new things. First, high-tech workers benefit from higher substitutability between capital and code in the corn sector. This makes sense. Indeed, when capital and code are perfect substitutes, corn production is linear in the sum of the two inputs. Hence, the marginal product of code is unchanged in response to a rise in the stock of code. Consequently, compensation to coders is unaffected by a rise in code retention and there is no mechanism for a fall in workers’ saving and investment or for long-run immiseration.

Second, with both Cobb-Douglas production and preferences, the path of the capital-to-code ratio in response to a rise in delta, starting from $\delta = 0$, is

independent of the absolute and relative numbers of each type of worker.²⁰

Third, a positive δ shock always produces a tech boom with increases in both the price of code and the wage of high-tech workers.²¹ In most simulations, the boom is short lived, auguring a major tech and saving bust. Finally, the δ shock generally raises labor's share in the short run and lowers it in the long run.

Figure 8 presents a contour map of the long-run compensating differential. Its top half considers combinations of saving preference parameters ϕ and shocks to δ assuming table 1's values of the other parameters. Because the two types of workers are perfect substitutes, the compensating differential for both types is the same. Redder areas reference higher long-run utilities relative to the initial steady state. Bluer areas reference the opposite. Long-run utility increases most when δ is large and the saving rate is high. It decreases the most when the δ shock is high and the saving rate is low.

Figure 8's lower half considers joint shocks to the saving rate and code-writing productivity z).²² Higher values of each reinforces their individual positive impacts on long-run utility. As opposed to δ shocks, increases in code-writing productivity (z) enhance all agents' welfare. The reason is simple – this shock makes living, but not deceased high-tech workers more productive. Increasing labor's productivity in other tasks has the same result. As this model posits no disutility from labor, reducing labor's productivity is isomorphic to restricting its supply. Policies that attempt to raise wages by reducing labor supply - such as increasing the minimum wage - will, therefore, backfire.

Figure 9 considers combinations of the saving rate, ϕ , and the good sector's elasticity of substitution, ε_y . It visualizes the aforementioned sensitivity of long-run utility to the substitutability of code for capital. High-tech workers, in particular, benefit from both higher levels of their scarce complement and from

²⁰Consider a doubling of H . This will double H_Y in the $\delta = 0$ economy. But if H_Y also doubles along the entire transition path, the path of k will remain unchanged. One can see this by combining the equation for market-clearing in capital with that for market-clearing in code. This, of course, requires that the path of H_S be twice as large as well. But this outcome as well as a doubling of q_t is implied by (16). This k -path invariance to initial levels of H and G is somewhat surprising and suggests that transforming more low-tech into high-tech workers may have less impact on the economy than one might have thought. Still, such a policy, if enacted before the rise in delta, would lower the real wages of skilled workers. (Their wages valued in corn wouldn't change, but the higher price of prayer would lower their real wage.) It would also improve the relative welfare of those who remain unskilled workers since their wage measured in units of capital will rise thanks to the higher marginal revenue of their labor. Additional effects would arise were H or G to vary once delta had risen and the economy was in transition. In this case, the k path would temporarily fall making code and coding less valuable. However, in the long run, the real wages of each type of worker are independent of such transition effects on the path of k .

²¹This and the previous result can both be shown analytically.

²²Other parameters are those in the immiserating growth scenario.

being more substitutable for it. The figure also indicates that this sensitivity is greater for low than for high saving rates. Higher substitutability moderates the negative effects of capital's crowding out that occurs with a low saving rate. Low-tech workers also benefit from both monotonically. However, the relationship is slightly more complex. Low-tech workers benefit when some high-tech workers cross sectors to complement them. When both the saving rate and elasticity of substitution in the corn sector are high, the marginal product of high-tech workers in the corn sector declines very little as code accumulates. Therefore, few leave to complement low-tech workers, limiting low-tech workers' gains.

5 The Role of Property Rights and Rivalry

To this point we've assumed that code is private and rival. Specifically, we've assumed that when one firm uses code it is unavailable for rent or use by others. But unlike physical capital, code represents stored information, which may be non-rival in its use. Non-rivalry does not however necessarily imply non-excludability. Patents, copyrights, trade secrets, and other means can be used to limit code's unlicensed distribution. On the other hand, the government can turn code into a public good by mandating it be open source. There are significant open questions about the legality and desirability of IP protection for AI and original works created by AI systems (Center for the Fourth Industrial Revolution 2018).

This section modifies the baseline model to investigate these questions. To do so we add a firm entry decision. Firm entry is important in the context of IP because when code can be simultaneously used by multiple parties, the number of these parties must be determined. Corn producing firms enter by paying a fixed cost each period, and gain access to the amount of free code available. This fixed cost corresponds to both the overhead necessary to run a business and the cost of discovering a new idea for applying AI. When code is excludable, firms may also rent an additional supply of it at a market price. Strong IP protections for AI incentives the creation of more code but limit the use of AI that already exists.²³

We first present the modified model in the baseline case of private (rival and excludable) code. We then explore two alternative scenarios. In the first, code

²³We incorporate some amount of free public code in all institutional scenarios to ensure entry.

is non-rival and non-excludable, i.e., it is a public good. In the second, code is non-rival, but excludable. In other words, those who develop AI can rent it out to as many companies as they like without friction.

Rival, Excludable (Private) Code

Corn firms maximize their profit, which is equal to

$$\pi_{j,t} = F(k_{j,t}, zH_{j,t} + a_{j,t} + \bar{A}) - C - r_t k_{j,t} - m_t a_{j,t}, \quad (38)$$

where $\pi_{j,t}$ are profits for firm j at time t , $F(\bullet)$ is a constant elasticity of substitution production function, $k_{j,t}$ is the amount of capital rented by the firm, $a_{j,t}$ is the amount of code rented by the firm, $H_{j,t}$ is the amount of high-tech labor hired by the firm, \bar{A} is the exogenously set amount of free code in the economy, and C is the per-period fixed operating cost. In equilibrium all firms have zero profits.

$$0 = F(k_{j,t}, zH_{j,t} + a_{j,t} + \bar{A}) - C - r_t k_t - m_t a_{j,t}. \quad (39)$$

Market clearing conditions are,

$$\sum a_{j,t} = \delta A_{t-1}, \quad (40)$$

$$\sum k_{j,t} = K_t, \quad (41)$$

$$\sum H_{j,t} = H_{A,t}, \quad (42)$$

$$Y = c_{o,t} + c_{y,t} - K_t + K_{t+1} - NC, \quad (43)$$

where N is the number of firms. All other equations are as in the baseline model. Since all firms are identical, (43) can be converted into an equation for N , the number of firms.

$$0 = NF\left(\frac{K_t}{N}, zH_t + \frac{1}{N}\delta A_{t-1} + \bar{A}\right) - NC - r_t K_t - m_t \delta A_{t-1} \quad (44)$$

Firms enter up to the point that the value of the public code they obtain for free, namely \bar{A} , equals their fixed cost of production. Thus,

$$\bar{A}F_{a,t} = C. \quad (45)$$

This fixes the marginal product of code at $\frac{C}{\bar{A}}$ in every period. Intuitively, new firms can acquire a perfect substitute for new code, and, thus, new coders at

a fixed cost by setting up shop and gaining access to \bar{A} in free code. Given that corn's production obeys constant returns to scale, fixing code's marginal product means fixing the ratio of capital to code. This, in turn, fixes the interest rate. Hence, the rental rates of code and capital are invariant to the increase in δ .

To solve the model an additional step is added to the iteration procedure. Given a guess of prices and stocks in a period, (44) is used to calculate N . This guess of N in each period is included in the next iteration to calculate new prices.

Figure 10 shows transition paths for this economy, with excludable, non-rival code, after an increase in the code retention rate. Parameter values are provided in table 4. Although an increase in δ does not change the marginal productivity of code, it does raise coder compensation. The reason is that coders can now sell property rights to the future use of their invention. Immediately after the shock, the wage of high-tech workers increases 10.2 percent, decreasing to 6.5 percent higher in the long run due to a decrease in the relative price of corn.

The number of firms in the economy decreases as a result of the technology shock, by 11.3 percent in the long run. Were the number of firms to remain fixed, the jump in δ would entail a higher code to capital ratio (in the short term the capital stock is fixed, and in the long run it increases by less than the code stock). This would mean a lower marginal productivity of code, which equation (45) precludes. It would also mean a negative payoff to setting up a new firm. Another way of viewing this relationship is that as the rental price of code decreases, the attractiveness of acquiring code by setting up a new firm decreases. As the latter is fixed, the rental price of code is fixed, with the margin of adjustment being fewer firms created.

In the period of the shock, welfare for the old decreases by 1.8 percent. This is because the cost of prayer increases 8.0 percent while the interest rate remains fixed. In the long run, high and low-tech workers are 6.5 percent better off. As the wage of high-tech workers can only increase as a function of the technology shock, this model variant does not admit long-run immiserating growth absent additional assumptions. For example, if the number of firms were to be fixed due to oligopolization of the industry, or if the fixed cost of firm entry were increasing in the number of firms, (45) would not hold, in which case the marginal productivity of code would decrease as code accumulates. This would reintroduce the possibility of immiserating growth.

Non-Rival, Non-Excludable (Public) Code

Consider next the case that code, in the period after it is produced, becomes a pure public good used simultaneously by every firm. This could arise by government edict, the wholesale pirating of code, or reverse engineering.

Profits are now

$$\pi_{j,t} = F(k_{j,t}, zH_{j,t} + a_{j,t} + \bar{A}) - C - r_t k_{j,t}, \quad (46)$$

as firms no longer need to rent their stock of code ($a_{j,t}$), where

$$a_{j,t} = \delta A_{t-1} \forall j \quad (47)$$

As before, firm entry and exit imply zero profits,

$$0 = NF\left(\frac{K_t}{N}, zH_t + \delta A_{t-1} + \bar{A}\right) - NC - r_t K_t. \quad (48)$$

and, because the amount of free code available to newly set up firms changes over time, (45) is modified to

$$(\delta A_{t-1} + \bar{A})F_{a,t} = C. \quad (49)$$

Finally, with investment in code no longer crowding out investment in capital,

$$K_{t+1} = \phi I_t. \quad (50)$$

Figure 11 shows the transition path after a δ increase for the case of non-excludable code with table 4's parameter values. The initial steady state is the same as in the prior case of excludable rival code. However, the response to the jump in δ is dramatically different. It has no immediate effect on the high-tech wage because workers no longer hold copyright to their code. They, therefore, have no incentive to move to the corn sector, leaving the economy unresponsive to the shock in the short term.

In the period after the shock, the economy begins to react. The stock of free public code, which now includes both \bar{A} plus all of the economy's legacy code, is larger. This induces more firm entry. The number of firms immediately more than doubles to 112.4 percent above its initial level in the short run and is 43.4 percent higher in the long run. As indicated in (49), with more free code available, the break even condition entails a lower rental rate of code. In

equilibrium, this entails more firms operating with less capital per unit of code.

The lower marginal product of code and, thus, of coders leads 30.3 percent of high-tech workers to move from coding into prayer in the period after the shock. National income peaks at 10.8 percent above its initial level in this period. The interest rate rises by 35.5 percent and the wage of low-tech workers increases by 15.0 percent.

The economy's transition is characterized by a series of damped oscillations. Periods of relatively high coder hiring and fewer firms is followed by periods of plentiful free code, more firm entry, and relatively low coder hiring. In the long run, the share of high-tech workers coding is 12.6 percent higher than its initial level and the high-tech wage is 23.4 percent lower. Welfare in the long run for high-tech workers is 9.7 percent lower. For low-tech workers, welfare is 8.0 percent higher. It is easy to select parameters such that both groups are worse off. As in the baseline model, the main mechanism for immiseration is the reduction of the high-tech wage leading to less capital accumulation. A contributing factor is the inefficiency introduced due to coders no longer being able to internalize the full value of their work.

Non-Rival, Excludable (Private) Code

Another possibility is that code is excludable, but non-rival in its use, permitting high-tech workers to license all their code to all firms in the period after it is produced. The equations for the rival, excludable model hold with the following exceptions. First, profits are given by

$$\pi_{j,t} = F(k_{j,t}, zH_{j,t} + \delta A_{t-1} + \bar{A}) - C - r_t k_{j,t} - m_t \delta A_{t-1} \quad (51)$$

Second, the price of code reflects its use by all firms.

$$p_t = \sum_{s=t}^{\infty} R_{s+1,t}^{-1} \delta^{s-t} m_{s+1} N_{s+1}. \quad (52)$$

As figure 12 shows, the δ shock produces a long-run, welfare-improving growth path, indeed a significantly better path than in the rival, excludable case. As in the rival, excludable case, firm entry satisfies equation (45). Hence, the interest rate and marginal product of new code are fixed, the wage of high-tech workers must increase, and a long-term welfare improvement is ensured. Welfare for the old in the period of the shock decreases by 2.6 percent, slightly more than in

the rival, excludable case. Thereafter, however, households are much better off, enjoying 11.5 percent higher utility in the long run, measured, as always, as a compensating differential.

This case features less entry. In the long-run, there are 20.8 percent fewer firms than before the δ shock. In the public code case, there is a 43.4 percent increase in the number of firms in response to the shock. In the private code case, the shock reduces the number of firms by 11.3 percent. Intuitively, since each firm can use all available code, fewer firms are needed. More surprising is the decrease relative to the rival code case. The reason is that with non-rival, excludable code, the effective supply and value of code is very high. The higher value crowds out capital investment. As can be seen from equation (45), the relatively small number of firms entering is due to a relatively higher effective stock of code and lower stock of capital.

6 Testable Implications and Supportive Evidence

Each of our simulations feature a temporary rise followed by a decline in labor's share of national income as well as a rise in code as a share of total assets. U.S. labor-share data going back four decades provides support for these trends.²⁴ There is also recent evidence of a decline in capital per worker, consistent with our model's immiseration scenarios.²⁵

Figure 13 displays three measures of labor's share of U.S. income based on three approaches to handling labor's unknown share of proprietorship and partnership income. The orange and gray curves use Bureau of Economic Analysis (BEA) data. The orange curve charts labor's share of total non-proprietorship national income assuming that labor's share of proprietorship income is the same as that of national income.²⁶ The blue curve displays labor's share of corporate income, i.e., it simply ignores the non-corporate sector.

The yellow curve displays labor's share of private businesses including proprietorships as calculated by the Bureau of Labor Statistics (BLS). The BLS imputes labor's share in proprietorship income by assuming proprietors (and part-

²⁴Other models, without smart machines, deliver this conclusion. Karabarbounis and Neiman (2013) attribute the decline to capital accumulation and their finding of gross substitutability between capital and labor. Rather than capital abundance, Rognlie (2015) argues that the decrease in the labor share is due to the scarcity of land. He attributes the decline in labor's share to an increase in property values and imputed rents.

²⁵Capital-hours ratio; BLS multifactor productivity series, Table PG-2-3. Records date back to 1949.

²⁶National income is measured at producer prices.

ners) earn the annual average wage in their industry. Proprietor income above this amount is considered capital income. This measure is smaller than the others because the BLS's income measure is not net of depreciation.

By all three measures, labor's share of income is lower in 2015 than in the mid 1970's. In the yellow curve, labor's share peaks in the mid-1970s with the two lowest shares recorded in 2014 and 2015. The precise percentage-point decline in labor's share between 1975 and 2014 are 5.96 percentage points, 5.88 percentage points, or 4.88 percentage points according to the orange, gray, and yellow curves, respectively.

Other authors, including Karabarbounis and Neiman (2013) and Brigdman (2014), report similar findings using related labor-share measures. The consensus view is that labor's share has decreased significantly since peaking in the mid 1970's. Armenter (2015) considers the possibility that the decrease in the BLS's measure is driven by the assumption that the proprietors pay themselves the average wage in their industry. When he instead fixes labor's share of proprietor's income at 85 percent, labor' share since 1975 still falls, but by less.

Our model predicts both a rise in code relative to other economic input and an increase in the share of output attributable to intangibles, i.e., inputs that are neither physical capital or labor. Code stocks have certainly increased since the invention of the digital computer and the silicon chip. Figure 14 reports stocks of R&D and software as a share of total U.S. fixed assets. According to the BEA, software grew from essentially zero percent of capital in 1960 to over 1.5 percent today. Combined software and R&D stocks have grown as a share of capital by about 3.5 percentage points over the same period.²⁷

Many papers suggest that the BEA underestimates the stock of organizational capital and code complementary to computers. Brynjolfson, Hitt, and Yang (2002) find that firms with large investments in computer capital have much higher valuations, that computer-capital investments lead to disproportionately large increases in firm valuations, and that firms that make such investments tend to be more productive in future years. Similarly, Hulten and Hao (2008) find that the book value of R&D-intensive firms in 2006 explains only 31 percent

²⁷These numbers are likely underestimates of the increasing importance of programmers, scientists and engineers in the economy. Software is decomposed in NIPA table 2.1 into own account, prepackaged and custom software. The true value of prepackaged software in the economy is likely undercounted because it is often pirated. It is also often free or sold at a discount in order to cross subsidize some other product or subscription (see Parker and Van Alstyne, 2005). BEA estimates of firms' internal creation of their own software are based on very conservative assumptions about the share of programmers who are developing new code, rather than maintaining old code, and the rate at which the software stock decays.

of their valuation. Both these papers argue that only firms who have made large investments in organizational and technological capital are able to implement innovative technologies.

Code and software controlled by firms that are not counted as assets by the BEA still increase the productivity of firms. Such firms would be more valuable than they should be based on only their observed assets. Figure 15 shows the value of the U.S. corporate sector less the replacement cost of its physical and financial assets.²⁸ This measure of the stock of intangible assets is highly cyclical due to the volatility of the stock market. Despite this, it shows a dramatic secular increase starting in the mid 1970s. For firms in the S&P 500, intangible assets increased from 17 percent of market value in 1975 to 84 percent in 2015 (Ocean Tomo, 2015).

Hall (2001) argues that the increase in the value of economy-wide intangible assets, and therefore Tobin's (average) q , is due to the creation of code and organizational capital within firms, which he calls 'e-capital.' Barkai (2016) and Barkai and Benzell (2018) also note that US firms' output per unit of observed capital has increased even as the marginal cost of capital (as measured by the real interest rate) has decreased since the mid-1980s. Simultaneous decreases in both capital and labor's share of income are consistent with an increasing share for intangibles like code. Barkai (2016) argues that the stock of intangible assets needed to explain the wedge between the observed average product of capital and its marginal cost is implausibly large. The level of intangible assets in 2014 would need to be 42 Trillion (or 54% of U.S. wealth) in order to explain the discrepancy. However, an extremely rapid increase in the share of intangibles in total assets is a phenomena implied by our model.

Long-run immiseration in our model hinges on a long-run decline in capital per worker and a corresponding increase in the interest rate. While capital per worker increased at an average rate of 2.5 percent from 1985 to the present, it has been decreasing since 2011 at .5 percent per year on average. Whether this short-term trend continues remains to be seen. But this measure significantly underestimates the extent to which physical capital per person has decreased. Capital services as measured by the BLS include accumulation of intellectual property and capital quality increases (through the deflator) that are attributable in our model not to physical capital per worker but to larger stocks of code.

²⁸U.S. Corporate intangible assets are calculated as U.S. corporate equity less corporate net worth from Federal Reserve series Z.1.

On the other hand, real interest rates have decreased steadily and dramatically since the mid-1980s. Almost all models of automation predict an increase in interest rates, at least in the short run, as a result of an advance in automation technology. Benzell and Brynjolfsson (2018) discuss this riddle. They propose that digital abundance has increased scarcity of an inelastically supplied complement to capital, code and ordinary labor. Integrating this additional factor of production to our model would generate decreasing interest rates while preserving all other qualitative results.

7 Conclusion

Will smart machines, which are rapidly replacing workers in a wide range of jobs, produce economic misery or prosperity? Our two-period, OLG model admits both outcomes. But it does firmly predict three things - a long-run decline in labor's share of income (apparently underway in OECD countries), tech-booms followed by tech-busts, and a growing dependency of current output on past software investment.

In our simple model, long-run immiseration is caused by a reduction in labor income and, thereby, saving and capital formation. Yes, the economy has better technology. But it has less capital. With the right parameters, the latter factor can outweigh the former. Immiseration is more likely the smaller the propensity to save and the smaller the elasticity of substitution of capital for code.²⁹ As the appendix shows, our results can be generalized to consider directed technological change in which AI can replace labor or capital. If AI is better at replacing labor, immiseration can readily follow.

Making higher code retention a win win for all current and future workers as well as initial elderly requires taxing high-skilled workers who benefit from the models technological breakthrough and investing the tax proceeds. This keeps the capital stock and wages from falling. Other policies for managing the rise of smart machines may backfire. For example, policies restricting labor supply will reduce total labor income. While this may temporarily raise wages, it will reduce saving, investment and the capital formation on which wages depend.

To the extent that AI is non rival, countries must weigh several factors in determining whether to grant AI developers property rights to their creations.

²⁹Alternative intertemporal preferences that make saving of the young depend on the interest rate could also influence the saving rate, making immiseration more (less) likely when saving is a decreasing (increasing) function of the interest rate.

Excludability increases the return to software development, increasing long-run levels of code, wages, national income, and welfare. It also lowers the potential for technologically driven business cycles. On the other hand, treating AI as a public good by, for example, denying AI software patents, ignoring AI copyrights, failing to prosecute software piracy, or mandating that new code be open-source, produces worse results along these dimensions. This said, such a policy increases short-term welfare as well as firm entry. This last implication is an increasingly important desideratum as countries become concerned about the concentration of power in the hands of a small number of superstar technology companies.

Our simple model illustrates the range of things that smart machines can do for us as well as do to us. Its central message is disturbing. Absent appropriate fiscal policy, which redistributes from winners to losers, smart machines can mean collective long-term misery. Ironically, the same AI that helps us produce more in the present can limit our production in the future.

References

- Abel, Andrew and Laurence J Kotlikoff**, “Intergenerational Altruism and the Effectiveness of Fiscal Policy: New Tests Based on Cohort Data,” in “Savings and bequests” 1994, chapter 7, pp. 31–42.
- Acemoglu, Daron**, “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality,” *The Quarterly Journal of Economics*, nov 1998, *113* (4), 1055–1089.
- , “Directed technical change,” *The Review of Economic Studies*, 2002, *69* (4), 781–809.
- **and David Autor**, “Chapter 12 – Skills, Tasks and Technologies: Implications for Employment and Earnings,” in “Handbook of Labor Economics,” Vol. 4 2011, pp. 1043–1171.
- **and Pascual Restrepo**, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 2018, *108* (6), 1488–1542.
- , **David Autor, David Dorn, Gordon H Hanson, and Brendan Price**, “Return of the Solow Paradox? IT, Productivity, and Employment in U.S. Manufacturing,” jan 2014.
- Aghion, Philippe, Benjamin F Jones, and Charles I Jones**, “Artificial Intelligence and Economic Growth,” 2017.
- Agrawal, Ajay K, Joshua Gans, and Avi Goldfarb**, “Prediction, Judgment, and Uncertainty,” in “Economics of Artificial Intelligence,” University of Chicago Press, 2017.
- Altonji, Joseph G., Fumio Hayashi, and Laurence J. Kotlikoff**, “Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data,” *American Economic Review*, 1992, *82* (5), 1177–98.
- Altonji, Joseph G, Fumio Hayashi, and Laurence Kotlikoff**, “Parental Altruism and Inter Vivos Transfers: Theory and Evidence,” dec 1995.
- Armenter, Roc**, “A bit of a miracle no more: the decline of the labor share,” *Business Review; Federal Reserve Bank of Philadelphia Research Department*, 2015, (Third Quarter), 1–8.
- Arrow, Kenneth J.**, “The Economic Implications of Learning by Doing,” *The Review of Economic Studies*, 1962, *3* (June), 155–173.

- Autor, David**, “Why are there still so many jobs? The history and future of workplace automation,” *The Journal of Economic Perspectives*, 2015, 29 (3), 3–30.
- Autor, David H. and David Dorn**, “The growth of low-skill service job and the polarization of the U.S. labor market,” *The American Economic Review*, 2013, 103 (5), 1553–1597.
- , **Frank Levy, and Richard J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, nov 2003, 118 (4), 1279–1333.
- Barkai, Simcha**, “Declining Labor and Capital Shares,” 2016.
- and **Seth G. Benzell**, “70 Years of US Corporate Profits,” 2018.
- Benzell, Seth G. and Erik Brynjolfsson**, “Centralization of Information and Authority in the Digital Economy,” in “Forthcoming in World Intellectual Property Organization Technology Trends: Report on Artificial Intelligence” 2018.
- and —, “Digital Abundance and Scarce Genius: Implications for Wages, Interest Rates, and Growth,” 2018.
- Bridgman, Benjamin et al.**, “Is Labor’s Loss Capital’s Gain? Gross versus Net Labor Shares,” *Washington: US Bureau of Economic Analysis*, 2014.
- Brynjolfsson, Erik and Andrew McAfee**, *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*, WW Norton & Company, 2014.
- Bureau of Labor Statistics**, “Occupational Employment Statistics,” March 2017.
- Center for the Fourth Industrial Revolution**, “White paper: Artificial Intelligence Collides with Patent Law,” April 2018.
- Cockburn, Iain M, Rebecca Henderson, and Scott Stern**, “The Impact of Artificial Intelligence on Innovation,” in “NBER Economics of Artificial Intelligence Toronto Conference,” University of Chicago Press, 2017.
- Costinot, Arnaud and Iván Werning**, “Robots, Trade, and Luddism,” *mimeo*, 2018.
- Frey, Carl Benedikt and Michael A Osborne**, “The future of employment: how susceptible are jobs to computerisation,” *Retrieved September*, 2013.

- Goos, Maarten, Alan Manning, and Anna Salomons**, “Explaining Job Polarization in Europe: The Roles of Technology, Globalization and Institutions,” *CEP Discussion Papers*.
- Gordon, Robert J**, *The rise and fall of American growth: The US standard of living since the civil war*, Princeton University Press, 2017.
- Guerreiro, Joao, Sergio Rebelo, and Pedro Teles**, “Should Robots be Taxed?,” 2017.
- Guvenen, Fatih**, “Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective,” *Journal of Monetary Economics*, 2006, *53* (7), 1451–1472.
- Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff**, “Risk-Sharing between and within Families,” *Econometrica*, 1996, *64* (2), 261–94.
- Hemous, David and Morten Olsen**, “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality,” *CEPR Discussion Papers*.
- Howitt, Peter**, “Steady endogenous growth with population and r. & d. inputs growing,” *Journal of Political Economy*, 1999, *107* (4), 715–730.
- Internal Revenue Service**, “Publication 535 (2015), Business Expenses.”
- Karabarbounis, Loukas and Brent Neiman**, “The Global Decline of the Labor Share.”
- Katz, Lawrence F. and Robert A. Margo**, “Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective.”
- Kelley, Kevin**, “The Three Breakthroughs That Have Finally Unleashed AI on the World,” *Wired Online Edition*, oct 2014.
- Keynes, John Maynard**, “Economic Possibilities for our Grandchildren,” in “Essays in Persuasion,” New York: W.W.Norton & Co, 1963, pp. 358–373.
- Kondratieff, Nikolai D**, “The Long Waves in Economic Life,” *The Review of Economics and Statistics*, 1935, *17* (6), 105–115.
- Lucas, Robert E.**, “On the mechanics of economic development,” *Journal of Monetary Economics*, jul 1988, *22* (1), 3–42.
- Madrigal, Alexis C.**, “Inside Google’s Secret Drone Delivery Program,” *The Atlantic Online Edition*, aug 2014.

- Mansfield, Edwin**, “Long Waves and Technological Innovation,” *American Economic Review*, 1983, 73 (2), 141–45.
- Marx, Karl**, *Capital: A Critique of Political Economy*, 1992 ed., Penguin Classics, 1887.
- Mishel, Lawrence, John Schmitt, and Heidi Shierholz**, “Don’t blame the robots. Assessing the job polarization explanation of growing wage inequality,” 2013.
- Nelson, Richard R. and Edmund S. Phelps**, “Investment in Humans, Technological Diffusion, and Economic Growth,” *American Economic Review*, 1966, 56 (1/2), 69–75.
- Ocean Tomo LLC**, “Annual Study of Intangible Asset Market Value Ocean Tomo, LLC,” 2015.
- of Economic Analysis, U.S. Bureau**, “National Income and Product Accounts.”
- of Governors of the Federal Reserve System, Board**, “Financial Accounts of the United States - Series Z.1.”
- Ortigueira, Salvador and Manuel S Santos**, “On the speed of convergence in endogenous growth models,” *The American Economic Review*, 1997, pp. 383–399.
- O’Rourke, Kevin, Ahmed Rahman, and Alan Taylor**, “Luddites, the industrial revolution, and the demographic transition,” *Journal of Economic Growth*, 2013, 18 (4), 373–409.
- Peretto, Pietro F. and John J. Seater**, “Factor-eliminating technical change,” *Journal of Monetary Economics*, 2013, 60 (4), 459–473.
- Poterba, James M.**, “The rate of return to corporate capital and factor shares: New estimates using revised national income accounts and capital stock data,” in “Carnegie-Rochester Conference Series on Public Policy,” Vol. 48 Elsevier 1998, pp. 211–246.
- Rebelo, Sergio**, “Long-run policy analysis and long-run growth,” *Journal of political Economy*, 1991, 99 (3), 500–521.
- Rognlie, Matthew**, “Deciphering the fall and rise in the net capital share: accumulation or scarcity?,” *Brookings papers on economic activity*, 2016, 2015 (1), 1–69.

- Romer, Paul M.**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, 98 (5), S71–S102.
- Rosenberg, Nathan and Claudio R. Frischtak**, “Long Waves and Economic Growth: A Critical Appraisal,” *American Economic Review*, 1983, 73 (2), 146–151.
- Sachs, Jeffrey D and Laurence J Kotlikoff**, “Smart Machines and Long-Term Misery,” dec 2012.
- , **Seth G Benzell, and Guillermo LaGarda**, “Robots: Curse or blessing? A basic framework,” Technical Report, National Bureau of Economic Research 2015.
- Schumpeter, Joseph A.**, *Business Cycles: A Theoretical, Historical and Statistical Analysis of the Capitalist Process*, New York: McGraw Hill Book Co., 1939.
- Sidrauski, Miguel**, “Rational choice and patterns of growth in a monetary economy,” *The American Economic Review*, 1967, 57 (2), 534–544.
- Stiglitz, Joseph E and Anton Korinek**, “Artificial Intelligence, Worker-Replacing Technological Change, and Income Distribution,” in “Economics of Artificial Intelligence,” University of Chicago Press, 2017.
- U.S. Bureau of Labor Statistics**, “May 2013 OES National Industry-Specific Occupational Employment and Wage Estimates.”
- , “Multifactor Productivity and Related Measures.”
- Uzawa, Hirofumi**, “Optimum technical change in an aggregative model of economic growth,” *International economic review*, 1965, 6 (1), 18–31.
- Website of the National Archives**, “Luddites: The Growth of Political Rights in Britain in the 19th Century.”
- Zeira, Joseph**, “Workers, Machines, and Economic Growth,” *The Quarterly Journal of Economics*, 1998, 113 (4), 1091–1117.
- Zuleta, Hernando**, “Factor saving innovations and capital income share in OLG models,” 2004.




Figure 1: Overlapping generations budget constraint summary for high-tech worker in the corn industry who is young in period t .

Table 1
Parameters for Immiserating Growth

Parameter Description	Model Parameter	Value
Elasticity in Service/Prayer Sector	ε_s	∞
Elasticity in Goods/Corn Sector	ε_y	1
Prayer High-Tech Input Share Param.	γ	0.5
Good Capital Input Share Param.	α	0.5
Code Retention Rate	δ	0 shocked to 0.7
Saving Preference Param.	ϕ	0.2
High-Tech Worker Quantity	H	1
Low-Tech Worker Quantity	G	1
Prayer Consumption Share	κ	0.5
Code Writing Productivity	z	1
TFP in Corn	D_y	1
TFP in Prayer	D_s	1

Table 1: This table gives parameter values for the first pair of illustrations of the effects of a permanent increase in the retention rate, δ , from zero to .7. We take the intermediate value of .5 for κ , α , and γ . The productivity terms z , D_Y , and D_S , are set to one.

Figure 2
Immiserating Growth




Figure 2: Transition paths based on table 1. Compensating Differential references the percentage change increase in steady-state consumption needed for initial steady-state households' lifetime utility to equal the lifetime utility of households who are old in a given period on the transition path. Output, capital and code stocks are in units. Wage and national income are deflated. Wage is total income in youth from labor in any occupation. Price of code, interest rate, and price of prayer are in units of the numeraire, corn.

Figure 3
Welfare Improving Growth
 (higher saving rate, $\phi = .85$)




Figure 3: Transition paths based on table 1, with the exception of a higher saving rate ($\phi = .85$). Compensating Differential references the percentage change increase in steady-state consumption needed for initial steady-state households' lifetime utility to equal the lifetime utility of households who are old in a given period on the transition path. Output, capital and code stocks are in units. Wage and national income are deflated. Wage is total income in youth from labor in any occupation. Price of code, interest rate, and price of prayer are in units of the numeraire, corn.

Figure 4
Inequality-Flipping Growth




Figure 4: Transition paths based on table 2. Compensating Differential references the percentage change increase in steady-state consumption needed for initial steady-state households' lifetime utility to equal the lifetime utility of households who are old in a given period on the transition path. Output, capital and code stocks are in units. Wage and national income are deflated. Wage is total income in youth from labor in any occupation. Price of code, interest rate, and price of prayer are in units of the numeraire, corn.

Figure 5
 Changed Corn Preference




Figure 5: Transition paths based on table 2, except in addition to the δ shock, κ is simultaneously shocked from .5 to .25. Compensating Differential references the percentage change increase in steady-state consumption needed for initial steady-state households' lifetime utility to equal the lifetime utility of households who are old in a given period on the transition path. Output, capital and code stocks are in units. Wage and national income are deflated. Wage is total income in youth from labor in any occupation. Price of code, interest rate, and price of prayer are in units of the numeraire, corn.

Table 2
Inequality-Flipping Growth

Parameter Description	Model Parameter	Value
Elasticity in Service/Prayer Sector	ε_s	1
Elasticity in Goods/Corn Sector	ε_y	1
Prayer High-Tech Input Share Param.	γ	0.5
Good Capital Input Share Param.	α	0.5
Code-Retention Rate	δ	0 shocked to 0.7
Saving Preference Parameter	ϕ	0.7
High-Tech Worker Quantity	H	2
Low-Tech Worker Quantity	G	1
Prayer Consumption Share	κ	0.5
Code Writing Productivity	z	1
TFP in Corn	D_y	1
TFP in Prayer	D_s	1

Table 2: This table gives parameter values for the second pair of illustrations of the effects of a permanent increase in the retention rate, δ , from zero to .7.

Figure 6
Comparing Four Case Studies




Figure 6: Transition paths from the first 4 cases presented (immiserating growth, etc.) superimposed.

Figure 7
National Income, Sensitivity Analysis




Figure 7: Illustration of the 10 sensitivity analysis cases superimposed. ‘Subs’ refer to cases in which the production technology of a sector is more substitutable. ‘Com’ refer to cases in which the production technology is more complementary.

Figure 8
Long-Run Compensating Differential for Alternative Saving and Code-Retention and Productivity Shocks




Figure 8: Compensating Differential references the ratio of consumption needed to achieve lifetime utility of households in the long run to initial steady-state consumption less 1. Parameters not on axes are given in table 1. X's denote parameter combinations with transition paths discussed in the text.

Figure 9
 Long-Run Compensating Differential for Alternative Saving and
 Elasticity of Substitution for Low and High-Tech Workers




Figure 9: Compensating Differential references the ratio of consumption needed to achieve lifetime utility of households in the long run to initial steady-state consumption less 1. Parameters not on axes are given in table 2. X's denote parameter combination in the inequality flipping case study.

Figure 10
Rival, Excludable (Private) Code




Figure 10: Transition paths based on table 4's parameters. Compensating Differential references the percentage change in initial steady-state consumption needed for the utility levels of workers in the initial steady state to equal their respective transition-path utility levels. Prayer sector output is measured in units of prayers. Corn output is net of firm entry fixed costs.

Table 4
Parameters for Institutional Simulations

Model Parameter	Role	Value
ε_s	Elasticity in Service/Prayer Sector	1
ε_y	Elasticity in Goods/Corn Sector	1
γ	Prayer High-Tech Input Share Param.	0.5
α	Good Capital Input Share Param.	0.5
δ	Code Retention Rate	0 shocked to 0.25
ϕ	Saving Rate	0.5
H	High-Tech Worker Quantity	1
G	Low-Tech Worker Quantity	1
κ	Prayer Consumption Share	0.5
z	Code Writing Productivity	1
D_y	TFP in Good Sector	1
D_s	TFP in Prayer Sector	1
C	Firm Setup cost	.055
\bar{A}	Exogenous Free Code	.25

Table 4: This table gives parameter values for illustrations of the effects of a one-time, permanent increase in the retention rate, δ , from zero to .25 given different institutional settings.

Figure 11
Non-Rival, Non-Excludable (Public) Code




Figure 11: Transition paths based on table 4's parameters. Compensating Differential references the percentage change in initial steady-state consumption needed for the utility levels of workers in the initial steady state to equal their respective transition-path utility levels. Prayer sector output is measured in units of prayers. Corn output is net of firm entry fixed costs.

Figure 12
Non-Rival, Excludable (Private) Code




Figure 12: Transition paths based on table 4's parameters. Compensating Differential references the percentage change in initial steady-state consumption needed for the utility levels of workers in the initial steady state to equal their respective transition-path utility levels. Prayer sector output is measured in units of prayers. Corn output is net of firm entry fixed costs.

Figure 13
Three Measures of Labor's Share of Income in the U.S.




Figure 13: Three measures of the U.S. labor share. The orange curve, labor's share of non-proprietorship income, is calculated as employee compensation divided by national income at producer prices less proprietorship income (NIPA table 1.12, lines 2/(1-25+26-18)). The gray curve, labor's share of income in the corporate sector, is calculated as corporate employee compensation divided by corporate business income less corporate taxes net of subsidies (NIPA table 1.13 lines 4/(3-9)). The yellow curve is the BLS's measure of labor share in the private business sector (from the BLS multi-factor productivity series). Dashed lines are fitted third-degree polynomials.

Figure 14
The Stock of Software and Software and R&D as a Share of U.S. Fixed Assets




Figure 14: The stock of software (solid line) and software plus R&D assets (dashed line) as a share of total fixed assets (authors' calculation based on NIPA table 2.1).

Figure 15
U.S. Corporate Intangible Assets as a Share of U.S. Wealth



Figure 15: U.S. corporate intangible assets as a share of U.S. wealth is calculated by subtracting the net worth of U.S. corporations from their equity value. Net worth is the replacement cost of fixed assets plus the market value of other assets less liabilities apart from owners' equity. This imputed value of intangible corporate assets (goodwill) is divided by total U.S. wealth (authors' calculation based on Federal Reserve financial accounts series Z.1).

Appendix

A Directed Technological Change

An important critique of the above models is that technological change is purely (directly) labor-substituting. Recent models of technological change, such as Acemoglu (2002) and Acemoglu and Restrepo (2018) have emphasized the directedness of technological change. That is, rather than only deciding how much technology to develop, individuals and firms can choose between a suite of research projects that each enter the production function in different ways.

Here we relax both assumptions in a stripped down, one-sector model, with only high-tech workers who can write both labor-substituting and capital-substituting software. This extension captures the spirit of the directed technological change literature in that workers can develop either a labor substituting or labor complementing technology. As we'll show, immiserizing technological change due to higher code retention can still occur, at least in cases when the code-retention shock disproportionately impacts labor-substituting software.

Production of the economy's single good, corn, satisfies

$$Y_t = D((1 - \gamma)(W_t + A_t^L)^\alpha + (\gamma)(K_t + A_t^K)^\alpha)^{\frac{1}{\alpha}}. \quad (53)$$

This production function deviates from the baseline model in two ways. First, we introduce a type of labor, W that is a perfect substitute for software but that doesn't create software as a byproduct. Second, we add a type of software A^K which is a substitute for capital.

Producers face the following profit maximization problem.

$$\pi_t = Y_t(W_t + A_t^L, K_t + A_t^K) - w_t W_t - r_t K_t - A_t^L m_t^L - A_t^K m_t^K, \quad (54)$$

where w_t is the wage of non-coders and m_t^L and m_t^K are the rental prices of labor substituting and capital substituting software, respectively.

Both types of code are written by software companies with fixed costs of entry and decreasing returns to scale. Production of code satisfies

$$A_{t,i}^L = z_L (L_{t,i}^L)^{\beta_L} \quad (55)$$

$$A_{t,j}^K = z_K (L_{t,j}^K)^{\beta_K}, \quad (56)$$

where $L_{t,i}^L$ and $L_{t,j}^K$ are demands for coders by labor substituting and capital substituting software companies, respectively and z_K , z_L , β_L and β_K are parameters. Both β_L and β_K are between zero and 1.

Software firms maximize profits – revenues from renting code net of paying coders and covering each period’s fixed cost of operation, F^L or F^K . Firms rent out their software at its marginal product and then sell their software after depreciation at prices p_t^L and p_t^K .

$$\pi_{t,i}^L = A_{t,i}^L (m_t^L + \delta^L p_t^L) - w_t L_{t,i}^L - F_t^L \quad (57)$$

$$\pi_{t,j}^K = A_{t,j}^K (m_t^K + \delta^K p_t^K) - w_t L_{t,j}^K - F_t^K \quad (58)$$

Total software of each type accumulates as in the baseline model.

$$A_t^L = \sum A_{t,i}^L + \delta_L A_{t-1}^L \quad (59)$$

$$A_t^K = \sum A_{t,i}^K + \delta_K A_{t-1}^K. \quad (60)$$

Firms enter until profits are zero. This implies

$$\pi_{t,i}^L = 0 \quad (61)$$

$$\pi_{t,i}^K = 0 \quad (62)$$

Combining equations gives

$$0 = z_L (L_t^L / N_t^L)^{\alpha_L} (m_t^L + \delta^L p_t^L) - w_t L_t^L / N_t^L - F_t^L \quad (63)$$

$$0 = z_K (L_t^K / N_t^K)^{\alpha_K} (m_t^K + \delta^K p_t^K) - w_t K_t^K / N_t^K - F_t^K, \quad (64)$$

where N_t^K and N_t^L are the number of K and L software companies, respectively.

Both types of software are priced as the present discounted value of their marginal product. The capital stock equals saving of the young net of their purchase of ownership rights to software.

$$K_{t+1} = \phi I_t - p_t^L \delta A_t^L - p_t^K \delta A_t^K \quad (65)$$

Labor market equilibrium requires

$$\bar{L} = L_t^L + L_t^K + W_t. \quad (66)$$

Since workers can move freely between the tasks of ordinary worker, capital-substituting coder, or labor substituting coder, the wage, w_t , is the same across all task. Hence, the total income of the young is

$$I_t = w_t \bar{L}. \quad (67)$$

There is no depreciation of capital. Output is either consumed, invested or used to cover fixed costs.

$$Y_t = C_{y,t} + C_{o,t} + K_{t+1} - K_t - F_t^K N_t^K - F_t^L N_t^L, \quad (68)$$


Appendix figure 1 reports steady-state results based on the parameters in appendix table 1. As can be seen, steady-state outcomes for an economy with a higher δ^L are worse, while those with relatively higher δ^K are better. The mechanisms for immiseration are similar to the baseline model. The increase in A^L stocks decreases wages and increases interest rates. The decrease in wages decreases capital stocks, which decreases wages yet further. Immiseration can arise if the exogenous shock makes it disproportionately easier to accumulate labor- than capital-substituting software.

Appendix Table 1
Parameters for Directed Technical Change Simulation

Model Parameter	Role	Value
α	Elasticity Param. in Final Good Prod.	.952
γ	Capital Input Share Param.	.7
\bar{L}	Total Labor Supply	25
δ^L	Labor Subs. Code Retention Rate	0 initially
δ^K	Capital Subs. Code Retention Rate	0 initially
ϕ	Saving Rate	.8
F^L	L Code Firm Fixed Cost	$5 * 10^{-13}$
F^K	K Code Firm Fixed Cost	$5 * 10^{-13}$
Z_L	L Code Productivity Param.	.00013
Z_K	K Code Productivity Param.	.00013
β_L	L Decreasing Returns to L-Software Coding Param.	.6
β_K	K Decreasing Returns to K-Software Coding Param.	.6
D	TFP in Final Production	1

Appendix Table 1: This table presents parameter values for illustrations of the effects of a one-time, permanent increase in the depreciation rate of labor substituting and capital substituting software as displayed in appendix figure 1.

Appendix Figure 1
 Long-Run Compensating Differential for Alternative δ Shocks



Compensating Differential references the ratio of consumption needed to achieve life-time utility of households in the long run to initial steady-state consumption less 1. Parameters not on axes are given in appendix table 1.