

NBER WORKING PAPER SERIES

UNEVEN GROWTH: AUTOMATION'S IMPACT ON INCOME
AND WEALTH INEQUALITY

Benjamin Moll
Lukasz Rachel
Pascual Restrepo

Working Paper 28440
<http://www.nber.org/papers/w28440>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

We thank Daron Acemoglu, Fabrizio Perri, Matt Rognlie, Richard Rogerson, Stefanie Stantcheva, Ivan Werning and seminar participants at various institutions for useful comments and Gabriel Zucman for kindly providing some tabulations from the dataset constructed by Piketty, Saez and Zucman (2018). We also thank Benny Kleinman, Nils Lehr and Max Vogler for outstanding research assistance. The views presented here are solely those of the authors and do not represent the views of the Bank of England or its policy committees. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. Moll received funding for this project from the Leverhulme Trust and the European Union's Horizon 2020 research and innovation programme under grant number No. GA: 865227.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Benjamin Moll, Lukasz Rachel, and Pascual Restrepo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Uneven Growth: Automation's Impact on Income and Wealth Inequality
Benjamin Moll, Lukasz Rachel, and Pascual Restrepo
NBER Working Paper No. 28440
February 2021
JEL No. E21,E22,E24,E25,J31

ABSTRACT

The benefits of new technologies accrue not only to high-skilled labor but also to owners of capital in the form of higher capital incomes. This increases inequality. To make this argument, we develop a tractable theory that links technology to the personal income and wealth distributions – and not just that of wages – and use it to study the distributional effects of automation. We isolate a new theoretical mechanism: automation increases inequality via returns to wealth. The flip side of such return movements is that automation is more likely to lead to stagnant wages and therefore stagnant incomes at the bottom of the distribution. We use a multi-asset model extension to confront differing empirical trends in returns to productive and safe assets and show that the relevant return measures have increased over time. Automation accounts for part of the observed trends in income and wealth inequality and macroeconomic aggregates.

Benjamin Moll
London School of Economics
Houghton Street
London WC2A 2AE
United Kingdom
b.moll@lse.ac.uk

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
pascual@bu.edu

Lukasz Rachel
Bank of England
Threadneedle Street
London, EC2R 8AH
United Kingdom
lukasz.rachel@bankofengland.co.uk

An online appendix is available at: <http://www.nber.org/data-appendix/w28440>

Introduction

Over the past forty years, economic growth in many advanced economies has been unevenly distributed. In the United States, while the aggregate economy has grown at roughly two percent per year, income percentiles corresponding to the lower half of the distribution have stagnated. At the same time, incomes at the 95th percentile have roughly doubled and top 1 percent incomes have roughly tripled.¹

One potential driver of these trends that is often cited by pundits and policy makers alike is technical change, and in particular the automation of tasks performed by labor. A large literature in macro and labor economics has studied how technology and automation affect the distribution of labor incomes.² But not all income is labor income and capital is an important income source, particularly at the top of the distribution where incomes have increased the most. Existing theories therefore paint an incomplete picture of technology’s implications for overall income inequality. This shortcoming is particularly acute when it comes to automation, technical change that substitutes labor with capital, and which therefore increases the importance of capital in the economy.

We argue that the benefits of new technologies—and in particular automation technologies—accrue not only to high-skilled labor but also to owners of capital in the form of higher capital incomes. This increases inequality. To make this argument, we develop a tractable framework that allows us to study the impact of technology on factor prices and the personal income and wealth distributions—and not just that of wages—and use it to study the distributional effects of automation.

Our framework provides a complete characterization of how technology and, more generally, changes in the economy’s production and market structure affect the personal distribution of income, wages, and capital ownership, as well as macroeconomic aggregates.³ We achieve this by assuming that households differ in their skills and household wealth accumulation is subject to random *dissipation shocks* which leave them with zero assets and only their labor income, thereby capturing the hazards of accumulating and maintaining a fortune. Dissipation shocks allow us to deviate in a tractable fashion from models that admit a representative household (like variants of the neoclassical growth model) and introduce two key features relative to such models that are crucial for understanding how automation affects inequality and aggregates.⁴ First, our theory generates well-defined and tractable steady-state distributions for wealth and income. Second, the long-run supply of capital

¹See for example [Census Bureau \(2015\)](#) and [Piketty, Saez and Zucman \(2018\)](#).

²See for example [Katz and Murphy \(1992\)](#), [Krusell et al. \(2000\)](#) and [Autor, Katz and Kearney \(2006\)](#)

³Importantly, ours is a theory of the *personal* income distribution and not just of the *factor* income distribution. The latter type of theory – for example “two class models” with capitalists and workers – cannot speak to a number of empirical regularities in developed countries, for example that individuals at the top of the labor income distribution typically also earn substantial capital incomes (and vice versa).

⁴Our notion of “models that admit a representative household” allows for skill heterogeneity and includes models that assume there are different skill types but these are all members of the same representative household, and models where Gorman aggregation holds (see Theorem 5.2 in [Acemoglu, 2009](#)).

is less than perfectly elastic and therefore determines the long-run effect of technology on returns to wealth (prices) and the amount of capital used in production (quantities).

The framework underscores two key and novel channels through which technology affects income inequality and contributes to the pattern of uneven growth described above. First, automation increases wealth and capital income inequality by raising returns to wealth. Second, relative to theories in which returns are unaffected, automation is more likely to lead to stagnant wages and, therefore, stagnant incomes at the bottom of the income distribution. The key to understanding both results is that the long-run capital supply in our model is upward-sloping. Automation increases the demand for capital relative to labor and, because supply is upward-sloping, this demand shift *permanently* increases returns to wealth. The first part, that automation directly increases wealth and income inequality, then follows because some households receive a higher return on their assets and grow their fortunes more rapidly. The second part—that wages are more likely to stagnate—follows because some of the productivity gains from automation do not accrue to workers but rather to owners of capital in the form of a higher return to their wealth. Neither of these mechanisms would be operational in textbook representative household models, in which the long-run capital supply is perfectly elastic and therefore returns are unaffected by technology.

We first illustrate these mechanisms using a simple baseline model and later extend this model in a number of directions. In the baseline model, capital is the only asset in the economy and all households earn the same return on their wealth. We obtain two main analytical results that illuminate how automation affects inequality.

First, the steady-state return exceeds the discount rate by a premium $p \times \sigma \times \alpha_{net}$, where p is the arrival rate of dissipation shocks, σ is the inverse of the intertemporal elasticity of substitution, and α_{net} is the net capital share – an object that rises with automation. Therefore, in contrast to textbook representative household models in which this return equals the (fixed) discount rate, the return in our model increases with automation.

Second, the steady-state wealth accumulation process of individual households generates an exact Pareto distribution for both effective wealth (the sum of financial and human wealth) and income conditional on wages.⁵ In equilibrium, tail inequality – as measured by the inverse of the Pareto shape parameter – simply equals the net capital share α_{net} , meaning that both top income and wealth inequality increase with automation.⁶ Intuitively, the gap between the return to wealth and the discount rate determines the speed at which *individual households* accumulate wealth in the absence of a dissipation shock, which in turn determines the thickness of the wealth distribution. In equilibrium, this gap equals the premium $p \times \sigma \times \alpha_{net}$ and therefore the net capital share α_{net} shows up in the tail index of these distributions. This result illustrates in a transparent fashion the new mechanism emphasized above that automation contributes to income inequality by permanently increasing returns

⁵Because of dissipation shocks, there is churn at the individual level even though aggregate capital is constant in steady state: households accumulate wealth over time but lose it with some probability.

⁶We also show that the scale parameter of these Pareto distributions is determined by wages which is another channel through which automation affects income and wealth inequality.

to wealth and the concentration of capital ownership.⁷

The flip side of the finding that the return to wealth rises with automation is that wages of displaced workers are more likely to fall, not only *relative* to other skill types (as commonly emphasized in the literature), but also in *levels*. In our theory, automation generates productivity gains but some of these productivity gains accrue to capital owners in the form of a higher return to their wealth. The higher cost of capital permanently limits the expansion of investment and output in response to this technological improvement. As a result, automation can lead to stagnant or falling real wages even in the long run, especially for workers whose skills are more susceptible to automation. This is in contrast to models that admit a representative household. In those models, the supply of capital is perfectly elastic, and automation leads to a substantial increase in capital accumulation and higher average wages in the long run (Acemoglu and Restrepo, 2018; Caselli and Manning, 2018).

Although our baseline model is intentionally stylized, and some of the assumptions are stark, the mechanisms underlying our results are considerably more general. The two key ingredients behind our results are an upward-sloping long-run supply of capital—so that technology persistently affects asset returns—and a nexus between returns to wealth and inequality. The first ingredient, an upward-sloping capital supply, seems to us a more natural and less extreme starting point than the perfectly-elastic capital supply of the textbook growth model and its relatives. It is also a feature of overlapping generations (OLG) models, models with a life-cycle component, and models with labor income risk and precautionary savings (as in Aiyagari-Bewley-Huggett models). The second ingredient, a return-inequality nexus, emerges naturally in models in which stochastic wealth accumulation at the individual level generates a fat-tailed wealth distribution. This includes models with stochastic returns or discount rates (Krusell and Smith, 1998; Benhabib, Bisin and Zhu, 2011, 2015) and models with explosive growth coupled with a birth and death process (Wold and Whittle, 1957; Steindl, 1965; Benhabib and Bisin, 2007; Jones, 2015; Sargent, Wang and Yang, 2020).⁸ Also the argument of Piketty (2014, 2015) that top wealth inequality depends on “ $r - g$ ” highlights precisely this nexus.

In our baseline model, automation has important distributional consequences because it raises *the* return to wealth. One of the main challenges in applying this theory to the data is that returns on different assets have displayed divergent empirical trends, in particular treasury rates have declined over time (e.g. Rachel and Summers, 2019) while the return to other assets, including US business capital and equity have increased (e.g. Caballero, Farhi

⁷This new channel differs from the common argument that a rise in the capital share leads to higher inequality because capital income is more unequally distributed than labor income (Meade, 1964; Piketty, 2014). As we discuss in detail, such compositional effects are small relative to the data and to the changes in capital ownership generated by our model.

⁸See Benhabib and Bisin (2018) for a review of models capable of generating fat-tailed wealth distributions. Among the models surveyed, only models where the tail of the wealth distribution is induced by the tail of the distribution of labor earnings (“models with skewed earnings”) lack a return inequality nexus. These include models with finite lives and no inheritances, and simple versions of Aiyagari-Bewley-Huggett models (Stachurski and Toda, 2018).

and Gourinchas, 2017; Gomme, Ravikumar and Rupert, 2011; Reis, 2020). What is then the relevant return to evaluate our theory? And how should we account for changes in the growth rate of the economy and other forces affecting returns?

To answer these questions, we extend our model to feature multiple assets with different returns, richer capital income risk, markups, taxation and long-run growth. Despite the added complexity, the model remains tractable, demonstrating the potential of our baseline model to form the core of a more elaborate and realistic framework. In our extended model, the return to risky capital exceeds the return to safe bonds because of compensation for risk or financial frictions that prevent arbitrage.⁹ We show that the key return measure determining wealth inequality is the overall return to investors' wealth, r_W , which depends on investors' portfolio allocation. The safe rate is not necessarily informative about this return. In fact, there are equilibria in which investors are leveraged, and lower safe rates *increase* this return and hence wealth inequality. Moreover, automation leads to a higher return on investors' wealth and a widening spread between risky and safe returns, with ambiguous implications for the level of the safe rate. Finally, the model also shows that, when there are changes in the growth rate of the economy, wealth inequality depends on $r_W - \sigma \times g$, which determines how fast successful investors accumulate wealth relative to the growth rate of the economy, and that this gap always rises with automation.

These results show that, when assessing the validity of our mechanism, one should focus on the return to investors' wealth and the gap between this return and $\sigma \times g$. The level of safe rates offers little information on the extent to which technology is increasing the demand for capital and how this might be contributing to rising wealth inequality. Building on this insight, we provide a series of estimates for the return to investors' wealth r_W , both in levels and relative to $\sigma \times g$. Our estimates suggest that, contrary to the observed trend in safe rates and in line with our theory, the relevant return to wealth r_W rose over the 1980–2007 period, and $r_W - \sigma \times g$ increased by 0.7–2 percentage points.

The addition of risky capital in the extended model also generates higher levels of inequality and churn at the top of the wealth distribution. In our baseline model with dissipation shocks only, top wealth inequality as measured by the inverse of the Pareto shape parameter equals $\alpha_{net}^* \approx 0.2$ which is much lower than the corresponding empirical values of around 0.7. In contrast, the extended model with capital income risk can generate the higher tail inequality that we see in the data. Moreover, the thicker tail also exhibits more frequent and sizable churn, with investors rapidly reaching the top of the wealth distribution after receiving a series of high returns (or leaving it after a low-return spell).

After presenting our main analytical results, we turn to a numerical evaluation of the model. Our objective here is not to conduct a full quantitative exercise, but to explore the range of implications of this class of models for aggregates and distributions under plausible

⁹The observed divergence between real risk-free interest rates and the returns to capital has been previously emphasized by Caballero, Farhi and Gourinchas (2017) and Farhi and Gourio (2018) who propose various candidate explanations ranging from rising market power to rising discounts on safe assets.

parameterizations, and to contrast these implications with salient empirical trends. In our numerical exercise, we study how the automation of routine jobs contributed to overall income inequality. To do so, we infer changes in automation by percentile of the wage distribution using exposure to routine jobs, which the literature singles out as jobs that can be easily automated using computer software or other equipment (see [Autor, Levy and Murnane, 2003](#); [Autor, Katz and Kearney, 2006](#)). To benchmark the extent to which these jobs have been automated, we calibrate the automation of routine jobs since 1980 to match the declining share of labor in GDP during this period.

The automation of routine jobs generates a pattern of uneven growth reminiscent of that observed in the US over the last forty years. Two features combine to produce this pattern. First, our model generates a decline in real wages at the bottom and middle of the income distribution, which accounts for part of the income stagnation observed at these percentiles. Second, it generates a rising concentration of capital income at the top of the distribution, which accounts for the sharp rise of incomes at the top. In particular, automation generates an increase in tail income inequality (measured by the inverse of the Pareto index) from 0.54 to 0.65, about 75% of the increase estimated by [Piketty, Saez and Zucman \(2018\)](#).

The first contribution of our paper is to the literature on technological change and automation. Like most theoretical papers in this area, we use a task-based framework to model automation ([Zeira, 1998](#); [Acemoglu and Autor, 2011](#); [Acemoglu and Restrepo, 2018](#)). Papers in this literature focus on wage inequality ([Autor, Levy and Murnane, 2003](#); [Autor, Katz and Kearney, 2006](#); [Acemoglu and Autor, 2011](#); [Hémous and Olsen, 2018](#)), or study the effect of automation on aggregates and wages using a representative household framework ([Acemoglu and Restrepo, 2018](#); [Caselli and Manning, 2018](#)). One exception is [Sachs and Kotlikoff \(2012\)](#), who study the possibility of immiserizing growth in an OLG model. We contribute to this literature by moving beyond representative-household models and exploring the implications of automation for inequality of total incomes across households. By doing so, we show that automation might contribute to rising incomes at the top and stagnant or declining wages at the bottom of the income distribution.

Our second contribution is to the literature exploring the determinants of wealth inequality. Several papers explore this question quantitatively in general equilibrium models, including [Krusell and Smith \(1998\)](#); [Castañeda, Díaz-Giménez and Ríos-Rull \(2003\)](#); [Kaymak and Poschke \(2016\)](#); [Hubmer, Krusell and Smith \(2016\)](#); [Straub \(2019\)](#). Relative to this literature, we isolate a new theoretical mechanism: technical change affects income and wealth distribution via returns to wealth. In contrast, other papers studying how technology affects wealth distribution focus on a more mechanical effect through wage dispersion (which is also present in our paper). Our model is deliberately simple and abstracts from labor income risk and a realistic treatment of life-cycle and bequests—important elements in quantitative theories of the wealth distribution. In exchange, we obtain analytical solutions for the steady-state distributions of wages, incomes and wealth.¹⁰ Furthermore, the economy

¹⁰The mechanism generating a Pareto tail in our model is a common feature of random growth processes

aggregates, and solving for its transition dynamics is as easy as in the neoclassical growth model. Thus, our model is closer in spirit to models of the Gorman class (see [Chatterjee, 1994](#); [Caselli and Ventura, 2000](#)) in which policy functions are linear and aggregates do not depend on the wealth distribution.¹¹ However, in contrast to these models, dissipation shocks generate a determinate steady-state wealth distribution and imply that our model does not admit a representative household.

Section 1 lays out our baseline theory of uneven growth. Section 2 presents the extended model and a measurement exercise showing that the relevant return measures have increased in line with our theory. In Section 3 we take the extended model to the data with a calibration of changes in automation across the wage distribution, and in Section 4 we show that this generates a pattern of uneven growth like that observed in the data. Section 5 concludes.

1 Benchmark Model of Uneven Growth

The model is cast in continuous time. For expositional clarity, we outline the model in stationary form. Appendices A and B provide proofs and full derivations. Appendix D describes the full model along the transition path.

1.1 Economic Environment

Households. There is a unit continuum of households that differ in their skills z , with ℓ_z denoting the population share of each skill. Households maximize standard preferences over utility flows from consumption subject to a flow budget constraint and a natural debt limit:

$$\max_{\{c_{z,t}, a_{z,t}\}_{t \geq 0}} \int_0^\infty e^{-\rho t} \frac{c_{z,t}^{1-\sigma}}{1-\sigma} dt \quad (1)$$

subject to: $\dot{a}_{z,t} = w_z + r a_{z,t} - c_{z,t}$, and $a_{z,t} \geq -w_z/r$,

where $a_{z,t}$ denotes assets, $c_{z,t}$ consumption, r the return to wealth and w_z wage income.

Without any modification, Gorman’s aggregation theorem applies (see theorem 5.2 in [Acemoglu, 2009](#)) and this model would admit a representative household. Independently of the production side, the steady state would involve a constant interest rate equal to the discount rate ρ . Moreover, the wealth distribution would be indeterminate (see [Caselli and](#)

(see [Gabaix, 2009](#), for a review). In fact, the process with dissipation shocks in our model is a simple and tractable example of a random growth process, and the idea that it leads to a Pareto distribution has been used before by many authors.

¹¹Because policy functions are linear “macro matters for inequality, but inequality does not matter for macro.” A large recent literature instead argues that models incorporating empirically realistic heterogeneity in household balance sheets and marginal propensities to consume often deliver different aggregate implications than do representative agent models, precisely because aggregates depend on distribution. We view linear policy functions and other abstractions with unrealistic implications as costs worth paying in return for our theory’s analytical tractability.

Ventura, 2000), and so the theory would be ill-suited to study how technology affects income and wealth inequality.

We break Gorman aggregation by assuming that the accumulation of wealth is subject to *dissipation shocks*, which arrive at a constant rate $p > 0$. We operationalize this assumption in the simplest possible way by assuming that after receiving a dissipation shock, households consume all of their wealth and are left with zero assets.¹² Appendix B shows formally that in the presence of dissipation shocks, households solve (1), where t denotes the time elapsed since the last dissipation shock, and households discount the future at a rate $\rho := \varrho + p$, where ϱ is the pure rate of time preference. Our formulation implies that following a dissipation shock we have $a_{z,0} = 0$. In what follows, we will refer to the case with $p = 0$ as the representative-household benchmark.

We view dissipation shocks as a reduced-form way of capturing the risks and hazards involved in raising and maintaining a fortune, both over time and across generations. These risks are a defining feature of wealth accumulation. For instance, recent work documents substantial churn at the top of the wealth distribution (Gomez, 2018; Zheng, 2019). Dissipation shocks provide a natural way of thinking about this churn. As we discuss in more detail in Appendix B.1, there are several possible micro-foundations for dissipation shocks, many of which are key ingredients in existing theories of wealth inequality:

- *perpetual youth*: individuals die with a constant probability p and are not altruistic toward their offspring. If there are no accidental bequests (either because of an annuity market, as in Blanchard, 1985, or because people consume their wealth right before they die), newborns start life with zero assets and only their labor income, mimicking the effect of a dissipation shock.¹³ As shown by Benhabib and Bisin (2007), this setup can be generalized to allow for some bequests (e.g. due to partial altruism or a “joy of giving” bequest motive) so long as individuals do not pass on their entire wealth to their offspring thereby forming a perfect infinitely-lived dynasty.
- *finite lives and stochastic altruism*: finitely-lived individuals are part of a dynasty and pass on their wealth over time. At rate p , the current member of the dynasty stops being altruistic and consumes all of her wealth so that the dynasty is interrupted.
- *population growth*: p captures net increases in population. Being born is isomorphic to receiving a dissipation shock because newborns start life with zero assets.
- *discount rate shocks*: at rate p , households become infinitely impatient and consume all of their wealth. This is a simple and tractable version of the more general process for discount rates assumed in Krusell and Smith (1998).

¹²For households with debt, we assume that dissipation shocks act as a cancellation of their debt. Alternatively, one could assume that households cannot borrow, or that dissipation shocks only hit households with a positive level of assets. These assumptions are not important for our results. As we will show, the steady state of our model involves all households accumulating assets at a positive rate.

¹³When there is an annuity market, the return obtained by households is $r + p$, where the addition of p accounts for the income from annuities. See Appendix B.1.1 for a more detailed discussion.

- *capital income risk*: at rate p , the assets held by a household become obsolete and lose all their value, i.e. a return of minus one hundred percent. This is a simple and tractable version of more general processes for idiosyncratic returns to wealth (e.g. Benhabib, Bisin and Zhu, 2011, 2015).

Technology. Our description of the production process emphasizes the role of tasks, which are either completed by workers or automated. Each skill type z works in a different sector that produces output Y_z . The economy produces a final good Y using these sectoral outputs according to a Cobb-Douglas aggregator

$$Y = A \prod_z Y_z^{\eta_z} \quad \text{with} \quad \sum_z \eta_z = 1.$$

Here, η_z denotes the importance of the sectoral output produced by skill type z in production. The productivity shifter A captures the role of factor-neutral technological improvements.

The production of sectoral output Y_z involves the completion of a unit continuum of tasks u , which are then combined via a Cobb-Douglas aggregator:

$$\ln Y_z = \int_0^1 \ln \mathcal{Y}_z(u) du.$$

These tasks can be produced using capital and skill- z labor as follows:

$$\mathcal{Y}_z(u) = \begin{cases} \psi_z \ell_z(u) + k_z(u) & \text{if } u \in [0, \alpha_z] \\ \psi_z \ell_z(u) & \text{if } u \in (\alpha_z, 1]. \end{cases}$$

The threshold α_z summarizes the possibilities for the automation of tasks performed by workers of skill z . Tasks $u \in [0, \alpha_z]$ are technologically automated and can be produced by capital $k_z(u)$ or labor $\ell_z(u)$. The remaining tasks are not technologically automated and must be produced by labor. Firms in sector z hire in a competitive labor market at a skill-specific wage w_z and they rent capital in a competitive rental market at rental rate $R = r + \delta$ where $\delta \geq 0$ is the depreciation rate. Therefore the unit cost of producing a task with capital is R and that of producing it with labor is w_z/ψ_z .

An increase in α_z captures the development of automation technologies that expand the range of tasks in which capital is now able to substitute workers of type z . For example, workers engaged in white-collar office work devote their time to tasks such as accounting, keeping and locating records, and customer support. Workers engaged in blue-collar work devote their time to tasks such as welding, painting, assembling, machining, and supervision. Over time, technological improvements have allowed the automation of some of these tasks, while others, like customer support or supervision, remain the domain of workers.

Equilibrium. Throughout, we assume competitive input and final good markets and define a steady state equilibrium of the economy as follows.

Definition 1 *The steady state equilibrium is given by aggregate output and capital, factor prices, a set of factor allocations, and consumption and saving policy functions such that:*

- capital and labor, $\{k_z(u), \ell_z(u)\}_{z,u}$, are allocated in a cost minimizing way to produce output Y given factor prices $\{w_z\}_z, R = r + \delta$ and technology.
- policy functions $a_{z,t}, c_{z,t}$ maximize utility given $a_{z,0} = 0, w_z$ and r .
- labor markets clear:

$$\int_0^1 \ell_z(u) du = \ell_z \quad \text{for all } z.$$

- the aggregate capital supplied by households and used by firms equals K :

$$K = \sum_z \int_0^{\alpha_z} k_z(u) du = \sum_z \ell_z \int_0^\infty a_{z,t} p e^{-pt} dt,$$

where $p e^{-pt}$ is the share of households that have accumulated wealth without experiencing a dissipation shock for t periods.

We first study households' consumption and saving decisions, which determine the supply of capital in the economy. We then consider the firms' problem of choosing the cost-minimizing mix of factor inputs, which determines the demand for capital. We are able to derive a closed-form solution for both the supply of and demand for capital, and for the interest rate that equilibrates the capital market.

Lemma 1 (Policy functions) *Suppose that $r > (r - \rho)/\sigma$. The solution to the savings problem is given by policy functions that are linear in effective wealth, $a_{z,t} + w_z/r$:*

$$\dot{a}_{z,t} = \frac{r - \rho}{\sigma} \left(a_{z,t} + \frac{w_z}{r} \right), \quad c_{z,t} = \left(r - \frac{r - \rho}{\sigma} \right) \left(a_{z,t} + \frac{w_z}{r} \right) \quad (2)$$

with $a_{z,0} = 0$.

The lemma shows that households accumulate assets at a rate $(r - \rho)/\sigma$ and consume a constant share of their effective wealth. The rate at which households accumulate assets, and thus the supply of capital, depends positively on the return rate r .

We now turn to the production side of the economy. To simplify the exposition, we make the following assumption for the rest of the paper:

Assumption 1 (Full adoption of available automation technologies)

$$\frac{w_z}{\psi_z} > R \quad \text{for all } z.$$

Assumption 1 ensures that in all tasks for which automation technologies are available, the cost of producing them with labor, w_z/ψ_z , exceeds the cost of producing them with capital, R . As a result, all automation technologies will be adopted.¹⁴

The following lemma characterizes output and factor prices.

Lemma 2 (Output and factor prices) *Suppose Assumption 1 holds. Output is given by*

$$Y(K) := \mathcal{A}K^{\sum_z \alpha_z \eta_z} \prod_z (\psi_z \ell_z)^{(1-\alpha_z)\eta_z}, \quad (3)$$

where \mathcal{A} is a constant that depends on parameters $\{A, \alpha_z, \eta_z\}$. Factor prices are given by

$$w_z = (1 - \alpha_z) \frac{\eta_z}{\ell_z} Y(K), \quad R = \alpha \frac{Y(K)}{K}, \quad (4)$$

where α is the average degree of automation in the economy: $\alpha := \sum_z \alpha_z \eta_z$.

Lemma 2 shows that aggregate output is given by a Cobb-Douglas production function with factor shares linked to the range of tasks performed by each factor. As tasks that were the domain of skill z get automated, this skill loses importance in production ($1 - \alpha_z$ declines) while capital becomes more important (α rises).¹⁵ Changes in the α_z 's thus provide a flexible way of capturing technological improvements that involve an expansion in the range of tasks that capital and machines can do at the expense of workers with particular set of skills.

1.2 Behavior of Aggregates and Wages

Lemma 1 showed how the return to wealth determines the rate at which households accumulate wealth and supply capital, and Lemma 2 characterized how technology shapes the demand for capital. We now characterize the equilibrium of the return to wealth and aggregates. We denote variables in the equilibrium steady state with an asterisk.

Proposition 1 (The return to wealth and the aggregate effects of automation) *The steady state is unique. The return to wealth, r^* , is given by the solution to the equation*

$$\frac{r^* - \rho}{\sigma p + \rho - r^*} = \frac{\alpha}{1 - \alpha} \frac{r^*}{r^* + \delta}. \quad (5)$$

¹⁴Assumption 1 involves endogenous factor prices, so it needs to be verified in equilibrium. Lemma A1 in Appendix A shows that it holds for steady-state factor prices when productivity A is high, so that steady-state wages are high relative to the steady-state rental rate R .

¹⁵Our model abstracts from firms and industries, and so the decline of the labor share following improvements in automation manifests at the aggregate level. In practice, larger and growing firms might be more likely to deploy automation technologies, or firms already operating more automated production processes may expand at the expense of others as automation technologies improve. These dynamics are consistent with the evidence in Autor et al. (2017); Kehrig and Vincent (2018) on the role played by firms in driving the decline in the labor share, and with the theoretical model in Martinez (2019).

Equivalently, denoting by $\alpha_{net}^* := r^*K^*/(r^*K^* + \sum_z w_z^*\ell_z)$ the net share of capital income, the return to wealth satisfies

$$r^* = \rho + \sigma p \alpha_{net}^*. \quad (6)$$

The return to wealth, r^* , the net capital share α_{net}^* , and the capital-output ratio, $(K/Y)^*$, are all increasing functions of α . Moreover, output increases in all α_z .

Equation (5) shows how technology and the parameters governing households' saving behavior determine the equilibrium return to wealth in steady state. We can think of the return to wealth as determined by the supply of capital and the demand for capital. To elaborate this argument, it is convenient to think in terms of capital relative to human wealth

$$k := \frac{K}{\bar{w}/r},$$

where $\bar{w} := \sum_z w_z \ell_z$ is the wage bill and \bar{w}/r is the value of human wealth in the economy.

To derive the supply of capital, we use the characterization in Lemma 1. Integrating equation (2) yields $\dot{K} = \frac{r-\rho}{\sigma} \left(K + \frac{\bar{w}}{r} \right) - pK$, where the term $-pK$ accounts for the wealth loss due to dissipation shocks. In steady state $\dot{K} = 0$ and hence

$$\frac{r-\rho}{\sigma} \left(K + \frac{\bar{w}}{r} \right) = pK.$$

Rearranging gives the (relative) supply of capital for $r \geq \rho$:

$$k^s = \frac{r-\rho}{\sigma p + \rho - r}. \quad (7)$$

This equation is an upward-sloping long-run capital supply curve and we depict it in Figure 1. Intuitively, as the return to wealth rises, households accumulate and supply more capital relative to their human wealth.




Figure 1: Equilibrium determination of the return to wealth.

Figure 1 also plots a capital demand curve which we obtain from the economy's production side: rearranging (4) yields

$$k^d = \frac{\alpha}{1 - \alpha} \frac{r}{r + \delta}.$$

and so (normalized) capital demand is a decreasing function of the return to wealth.

Equation (5) gives the return r^* at which demand equals supply.

The main implication of the supply and demand diagram in Figure 1 is that r^* lies between ρ and $\rho + \sigma p$, which implies that the return to wealth has a premium above ρ . Equation (6) shows that α_{net}^* determines exactly where in this range r^* lies. According to this equation, the premium of r^* above ρ is given by $\sigma \times p \times \alpha_{net}^*$.¹⁶

Figure 1 also illustrates the difference between the equilibrium in our model compared to the representative-household benchmark. With a representative household, the supply of capital is perfectly elastic and the return to wealth is fixed at $r^* = \rho$. In such models, only the quantity of capital adjusts in the long run, ensuring that technology can have at most a short-lived impact on asset returns.

Despite the simplicity of our model, we find the logic behind the upward-sloping supply of capital and the result that the return to wealth has a premium above ρ linked to the net capital share to be quite general. The driving force behind this result is that all the financial wealth and capital in the economy derives from labor income that households save and invest over time. Because everyone at some point in time had nothing but their human wealth, the equilibrium return to wealth r must exceed the discount rate ρ so as to incentivize households to accumulate and supply capital. How high above ρ must the return to wealth be? This depends on the importance of capital relative to human wealth. If the net capital share is high, and hence capital is important relative to human wealth, households must accumulate capital rapidly and this requires that the return to wealth exceeds ρ by a large amount.

The Proposition shows that the average extent of automation in the economy, α , increases the return to wealth and the net capital share. This finding can be seen directly from Figure 2. An increase in α raises the demand for capital relative to labor. This demand shift increases the return to wealth, the ratio of capital to human wealth, k^* , and the net share of capital. The proposition also shows that the capital-output ratio, $(K/Y)^*$, expands as automation increases the relative demand for capital. These findings are intuitive. At a fundamental level, automation makes capital more important than human wealth. A higher return is then required to get households to accumulate and supply the higher level of capital required for production.

Does the higher demand for capital following an increase in automation result primarily

¹⁶ The exact formula for r^* in equation (6) follows from rearranging equation (7) and holds independently of the production structure of the economy. The reason why the return to wealth exceeds ρ is different from the known formula $r = \rho + \sigma g$ which holds along a balanced-growth path where technology grows at a rate $g > 0$. Equation (18) in Section 2 generalizes equation (6) to an environment with sustained growth and shows that the equilibrium return is given by $\rho + \sigma g + p\sigma\alpha_{net}^*$, where α_{net}^* is now the capital share in income net of depreciation and growth.




Figure 2: Effect of automation—an increase in α —on the equilibrium return to wealth, r^* . The green arrows illustrate how the equilibrium changes in our model and in the representative-household benchmark.

in a higher return to capital or in an expansion of the capital-to-output ratio? The answer depends on the capital-supply elasticity, which is inversely linked to p . To illustrate this, consider the special case $\delta = 0$ in which we obtain simple formulas for aggregates:

Example 1 When $\delta = 0$, we have that $\alpha_{net}^* = \alpha$ and

$$r^* = \rho + \sigma p \alpha, \quad (K/Y)^* = \frac{\alpha}{\rho + \sigma p \alpha}. \quad (8)$$

As these equations show, for high values of p , automation leads to a sizable increase in the return to capital and a mild expansion in the capital-to-output ratio.

This is in contrast to what happens in the representative-household benchmark with $p = 0$. As Figure 2 shows, in this case the long-run capital supply is infinitely elastic at $r^* = \rho$. Compared to our model with an upward-sloping long-run capital supply, automation now results in a larger expansion of capital and has no effect on the return to wealth.

Finally, the Proposition shows that output always rises with automation. In particular, the appendix shows that, to a first-order approximation, the steady state effect of automation on output is given by

$$d \ln Y^* = \frac{1}{1 - \alpha} \underbrace{\sum_z \eta_z \ln \left(\frac{w_z}{\psi_z R} \right) d\alpha_z}_{:= d \ln \text{TFP}_\alpha} + \frac{\alpha}{1 - \alpha} d \ln (K/Y)^* > 0. \quad (9)$$

Output increases via two channels. First, by allowing the substitution of labor for capital, automation increases productivity (the first term). The contribution of automation to TFP is a (weighted) sum of the cost-saving at the task-level $w_z/(\psi_z R)$. Second, output increases due to endogenous capital accumulation. As discussed above, for high values of p , this second

force is modest, and automation generates a limited output expansion.

The following proposition summarizes the effects of automation on real wages.

Proposition 2 (The long-run effects of automation on wages)

- An increase in α_z reduces the wage w_z^* relative to other wages w_v^* for $v \neq z$.
- For a given increase in the α'_z 's, there exists a threshold $\bar{p} > 0$ such that, for $p > \bar{p}$, the average wage \bar{w}^* falls; and for $p < \bar{p}$, \bar{w}^* increases.

The effect of automation on relative wages is unambiguous and follows from the fact that $w_z = (1 - \alpha_z) \frac{\eta_z}{\ell_z} Y$ (see also [Hémous and Olsen, 2018](#); [Acemoglu and Restrepo, 2018](#)).

A more novel implication of the proposition is the possibility that automation may lead to stagnant wages for the average worker, which necessarily implies a more pronounced *real decline* in the wages of workers displaced by automation. Whether this is the case or not depends again on p , which determines how inelastic the supply of capital is in steady state.

Two complementary intuitions illustrate the importance of the capital supply in determining the behavior of the wage level. First, following a technological improvement that raises TFP by $d \ln \text{TFP}_\alpha > 0$, the dual version of the Solow residual implies that¹⁷

$$d \ln \text{TFP}_\alpha = (1 - \alpha) d \ln \bar{w} + \alpha d \ln R, \quad R = r + \delta \quad (10)$$

That is, productivity improvements accrue either to workers in the form of higher average wages or to capital owners in the form of a higher return to wealth. While this expression is very general and also holds outside of steady state and in a much broader class of models, consider now the steady state of our economy. In the representative-household benchmark $p = 0$, supply is perfectly elastic and hence $d \ln R = 0$ so that all productivity gains from automation accrue to labor – the inelastic factor.¹⁸ However, as p increases and capital supply becomes more inelastic, an increasingly larger share of these productivity gains accrues to capital in the form of a higher return. When $p > \bar{p}$ so that capital supply is sufficiently inelastic, most of these gains accrue to capital so that average wages actually fall, $d \ln \bar{w} < 0$.

An alternative intuition comes from studying directly the behavior of the average wage. From Lemma 2, we have that $\bar{w} = (1 - \alpha)Y$. From (9) we know that automation results in an

¹⁷This result holds in general whenever aggregate output exhibits constant returns to scale and markets are competitive. Under these assumptions, we have $Y = \bar{w}L + RK$ where we now allow for movements in labor supply L to underline the argument's generality so that \bar{w} denotes the average wage. Differentiating both sides of this identity, we get that, following a technological improvement, we have

$$d \ln Y = d \ln \text{TFP}_\alpha + \alpha d \ln K + (1 - \alpha) d \ln L = (1 - \alpha) d \ln \bar{w} + \alpha d \ln R + \alpha d \ln K + (1 - \alpha) d \ln L,$$

where $\alpha = RK/Y$. The expression in the main text follows by canceling the $d \ln K$ and $d \ln L$ terms. This derivation shows that the results in the proposition extend beyond our model: any type of technological change will raise wages provided that $d \ln R = 0$. See [Jaffe et al. \(2019, ch.18/19\)](#) for a textbook treatment.

¹⁸This part of the proposition is in line with papers that studied the impact of automation in settings with a representative household or an infinitely elastic supply of capital, such as [Simon \(1965\)](#); [Acemoglu and Restrepo \(2018\)](#); [Caselli and Manning \(2018\)](#). These settings correspond to the case $p = 0$ in our model.

output expansion, with the magnitude depending on the productivity increase $d \ln \text{TFP}_\alpha$ and the expansion in capital supply $d \ln(K/Y)^*$. The effect of automation on the average wage is therefore determined by the relative strength of this output expansion and the displacement effect captured by the term $1 - \alpha$. With sufficiently inelastic capital supply (i.e., $p > \bar{p}$), the displacement effect dominates. With sufficiently elastic capital supply (i.e., $p < \bar{p}$), the output expansion dominates.

Transitional dynamics and a numerical example. Appendix D shows that the transition dynamics for the aggregate variables are governed by the stable arm of the system of differential equations

$$\begin{aligned} \dot{C} - p\dot{K} &= \frac{1}{\sigma} (Y'(K) - \delta - \rho) (C - pK) - \mu pK \\ \dot{K} &= Y(K) - \delta K - C \\ \frac{\dot{\mu}}{\mu} &= \mu - Y'(K) + \delta + \frac{1}{\sigma} (Y'(K) - \delta - \rho), \end{aligned}$$

where μ is the (common) marginal propensity to consume out of effective wealth. This is similar to the usual system of equations determining the equilibrium in the neoclassical growth model, but is now coupled with an additional forward-looking equation determining the behavior of μ . Moreover, $C - pK$ —the consumption not triggered by dissipation shocks—follows a standard Euler equation, with an adjustment $-\mu pK$ which accounts for the share μ of dissipated wealth pK that would have been consumed.¹⁹

To illustrate the results from Propositions 1 and 2 and how the transition dynamics in our model deviate from the representative-household benchmark, we present a numerical example that shows the response of economic aggregates following a permanent improvement in automation technologies—an increase in the α_z 's.

For this numerical example, we use a subset of the parameters from the calibration of our extended model in Section 3. In particular, as explained there, we feed into the model a gradual increase in the α_z 's that starts in 1980 and matches the observed increase in the capital share between 1980 and 2014 from 34.5% to 43% and we assume that automated tasks see a cost reduction of 30%. This amount of automation generates a TFP increase of 2.5% (the share of automated tasks times the cost saving gains from automating a task – see (9)), which is a relatively small shock in terms of productivity gains. Finally, we set a value of $p = 4.5\%$, which implies a fairly high capital supply semi-elasticity, $d \ln K/dr$, of 50.

Figure 3 presents the transition dynamics for the labor share (summarizing the technological change fed into the model), output per worker, the net investment rate, the capital-output ratio, the return to wealth, and the average wage per hour. For comparison, we also

¹⁹As for the steady state of the model, we can solve the transition dynamics without knowing the evolution of the wealth distribution. This block-recursive property implies that one can first solve for aggregates and then use the behavior of aggregates to compute the evolution of distributions. This also implies that, in this model, “macro matters for inequality” but “inequality does not matter for macro.”

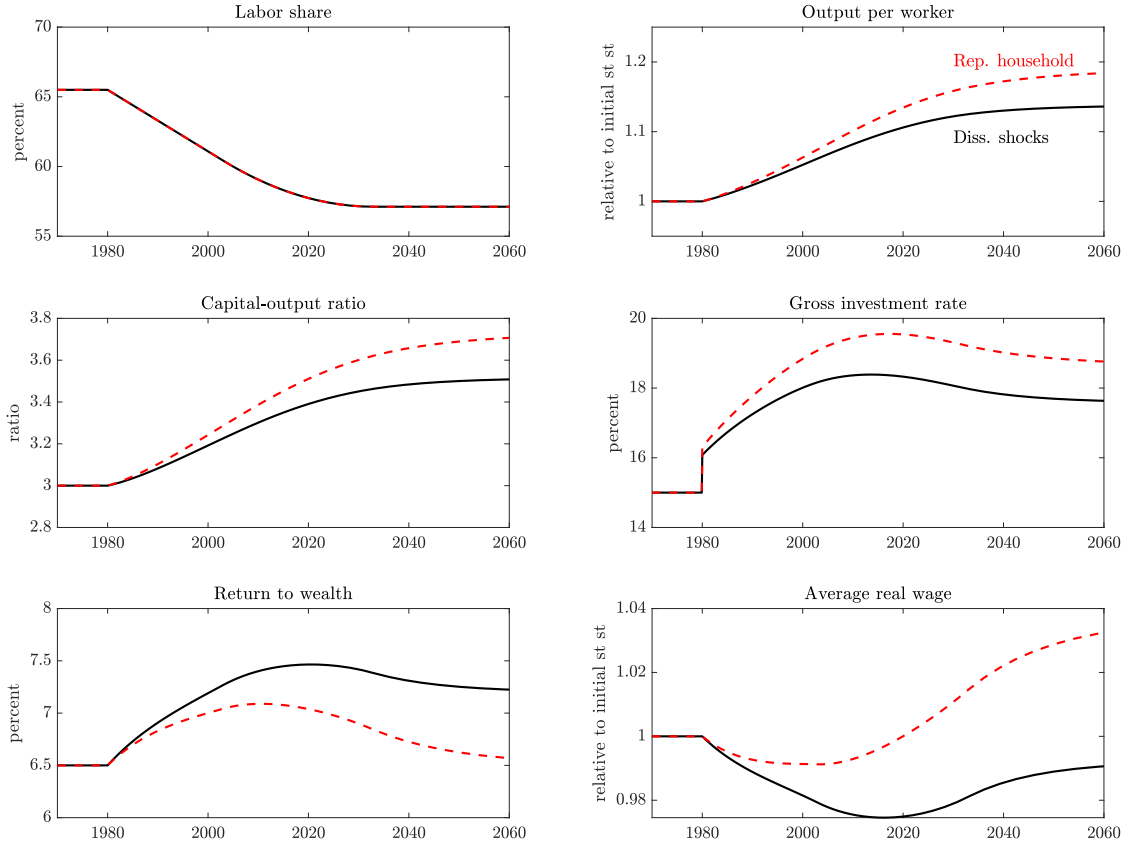


Figure 3: Transitional dynamics for the main aggregates following an improvement in automation technologies in 1980. The solid line presents the transition dynamics in our model. The dashed line presents the transition dynamics in the representative-household benchmark.

plotted the transition dynamics for the representative-household benchmark.

In our model, automation leads to a modest expansion of output of 11%. Despite reducing the labor share, automation generates a modest expansion in the capital-output and investment-output ratios of 15% and permanently increases the return to wealth from 6.5% to 7.2%. Thus, our model with an upward-sloping long-run capital supply provides a partial answer to the question “if the decline in the US labor share was driven by automation, shouldn’t investment and capital have increased?” The answer is “yes, but only slightly, precisely because capital supply is upward-sloping.”²⁰ It does not follow that automation

²⁰To see this in more detail, recall from above that the capital share increased from 0.345 to 0.43 percent between 1980 and 2014, a 25 percent increase. How much would we expect the capital-output ratio to increase? The answer can be seen from manipulating the definition of the capital share $\alpha = RK/Y$:

$$\frac{(K/Y)_{2014}}{(K/Y)_{1980}} = \frac{\alpha_{2014}/\alpha_{1980}}{R_{2014}/R_{1980}} \quad (11)$$

Now consider two scenarios, and assume in both that $R_{1980} = 11.5\%$ as in our calibration. If the rental rate R is unchanged so that $R_{2014}/R_{1980} = 1$, as in the representative household model, the capital-output ratio must increase by 25 percent. But suppose instead that the rental rate R rises by one percentage as in our model. Then $\frac{(K/Y)_{2014}}{(K/Y)_{1980}} = \frac{0.43/0.345}{12.5\%/11.5\%} = 1.15$. Rather than increasing by 25%, the capital-output ratio

predicts massive expansions in investment and capital. While automation leads to an increase in the gross capital share $\alpha = RK/Y$, a substantial fraction of this increase may show up in R rather than K/Y .

Despite the increase in the average output per hour of labor, and in line with Proposition 2, mean wages go down by 3% in 2020 and by 1% in the long run.²¹ This is in contrast to what would happen in the representative-household benchmark, where automation leads to a more pronounced economic expansion propelled by a boom in investment, a *temporary* increase in the return to wealth, and higher average wages in the long run.

The differences between these models underscore the importance of the capital-supply elasticity. Even though the supply of capital in our numerical example is fairly elastic (with a semi-elasticity $d \ln K/dr = 50$), the response of macroeconomic aggregates to automation differs significantly from what one would get in a representative household model.

1.3 Wealth and Income Inequality

We now study the wealth and income distributions. Before proceeding, recall that the steady-state return to wealth satisfies $r^* > \rho$. Therefore, while aggregate capital K is constant in steady state, households accumulate wealth at a positive rate. At the same time, they experience a dissipation shock and lose all their wealth at rate p . The combination of accumulation and dissipation gives rise to a unique, non-degenerate distribution of effective wealth, which in turn determines the income distribution.

Proposition 3 (Automation and the wealth and income distribution) *Denote households' effective wealth by $x_{z,t} := a_{z,t} + w_z^*/r^*$. The stationary distribution of effective wealth for skill type z is Pareto, with a PDF*

$$f_z(x) = \left(\frac{w_z^*}{r^*}\right)^\zeta \zeta x^{-\zeta-1}, \quad \frac{1}{\zeta} := \frac{1}{p} \frac{r^* - \rho}{\sigma} = \alpha_{net}^*.$$

The conditional and unconditional wealth distributions satisfy

$$\Pr(\text{wealth} \geq a|z) = \left(\frac{a + w_z^*/r^*}{w_z^*/r^*}\right)^{-\zeta}, \quad \Pr(\text{wealth} \geq a) = \sum_z \ell_z \left(\frac{a + w_z^*/r^*}{w_z^*/r^*}\right)^{-\zeta};$$

and the conditional and unconditional income distributions satisfy

$$\Pr(\text{income} \geq y|z) = \left(\frac{\max\{y, w_z^*\}}{w_z^*}\right)^{-\zeta}, \quad \Pr(\text{income} \geq y) = \sum_z \ell_z \left(\frac{\max\{y, w_z^*\}}{w_z^*}\right)^{-\zeta}.$$

increases by only 15%. As we discuss in Appendix I.1, this is well within the range observed in the data.

²¹Equation (10) shows why a small increase in the return to wealth of roughly 1 percentage point can have a large effect on wages. Using this equation with $\alpha = 0.345$ —the base point of our calibration—shows that a small increase in R^* of 10% (or 1 percentage point) is enough to ensure that the 2.5% increase in TFP driven by automation results in a decline in real wages, \bar{w} , of roughly 1%.

Moreover, the tail index $1/\zeta = \alpha_{net}^*$ increases with automation, α .

The proposition shows that the distribution of effective wealth for skill type z is Pareto with scale w_z^*/r^* and tail index $1/\zeta = \alpha_{net}^*$. The driving force behind this result is the random growth process governing the accumulation of wealth: starting from their human wealth $x_z(0) = w_z^*/r^*$, households accumulate capital over time at a rate $(r^* - \rho)/\sigma$. The distribution of wealth is stabilized by the dissipation shocks, which arrive at a rate p . Figure 4 depicts the process of accumulation. This random growth process gives rise to a Pareto distribution (see [Wold and Whittle, 1957](#); [Steindl, 1965](#); [Jones, 2015](#)) with tail index

$$\frac{1}{\zeta} = \frac{\text{individual household accumulation rate}}{\text{dissipation rate}} = \frac{(r^* - \rho)/\sigma}{p}.$$

As the formula shows, inequality depends on the ratio of the rate at which households accumulate wealth and the probability with which their wealth dissipates, p . The formula in the proposition follows from the observation that the steady state return to wealth is $r^* = \rho + \sigma p \alpha_{net}^*$, which implies a household accumulation rate of $(r^* - \rho)/\sigma = p \alpha_{net}^*$.




Figure 4: Dynamics of effective wealth accumulation.

The reason why we get a Pareto tail is that, while the average household retains a constant level of wealth due to the dissipation shocks, some fortunate households manage to accumulate wealth exponentially at a rate $(r^* - \rho)/\sigma > 0$, populating the tail of the wealth distribution over time. Higher returns increase the rate at which successful households accumulate wealth but do not change the rate at which the average household accumulates wealth, leading to widening differences in the rate at which fortunate households accumulate wealth and higher levels of wealth inequality.

Although our model delivers this insight in a stylized way, we see dissipation shocks as a simple and tractable way to micro-found a nexus between returns to wealth and inequality that is present in a broader class of more realistic models where the stochastic accumulation of wealth results in fat-tailed distributions. This broader class includes models with stochastic bequest motives (see [Benhabib and Bisin, 2007](#)), models with stochastic returns or discount rates (see [Krusell and Smith, 1998](#); [Benhabib, Bisin and Zhu, 2011, 2015](#)), or the extension

of our model presented in Section 2. While the details differ, in all of these models wealth accumulation at the top of the distribution is governed by a random growth process and this creates a natural nexus between returns and inequality (see Gabaix et al., 2016). In particular, higher returns increase the rate at which successful households accumulate wealth and hence a more fat-tailed wealth distribution. It follows that in this broader class of models, technological changes that lead to a higher return to wealth and more rapid wealth accumulation by successful households will generate a fatter tail of the wealth distribution, as is the case in our model.

The distribution of effective wealth is important in and of itself because it tells us about inequality in consumption and welfare (a corollary of Lemma 1). But our model also allows us to characterize the conditional and unconditional distributions of wealth and income, which is what we typically measure in the data. The remaining formulas in Proposition 3 provide an exact characterization of the wealth and income distributions.

The formulas in Proposition 3 show that technology affects both distributions via wages—which determine the scale parameters—, but more novel via the net capital share—which determines the return to wealth and the tail index of the income distribution, $1/\zeta$. These formulas illustrate why automation might contribute to the pattern of uneven growth described in the introduction. By depressing the real wages of displaced workers, automation will reduce incomes at the bottom of the distribution, especially for displaced households with few assets. By raising the return to wealth, automation will generate rising income inequality at the top driven by the rapid capital accumulation of some fortunate households. Both effects depend crucially on the fact that in our model, the long-run supply of capital is less than perfectly elastic, which implies that automation is more likely to reduce real wages for some workers and increase returns for households at the top of the income distribution.

Composition of top income earners. The formulas in Proposition 3 can also be used to study the composition of top income earners and quantify their share of national income. As we show in Appendix C, for small enough q , the share of national income held by the top q (i.e., the top 0.1%) is

$$S(q) = \Lambda q^{1-\alpha_{net}^*}, \tag{12}$$

where Λ is a constant that depends on the wage distribution. When there is no heterogeneity in wages, $\Lambda = 1$ and we obtain the usual formula for the top q percent share in a Pareto distribution (see Jones, 2015). One implication of this formula is that technology might affect $S(q)$ through wages (via the Λ term) but this would cause a proportional increase in $S(q)$ for all $q < \bar{q}$. Instead, by raising α_{net}^* , technology will increase the share of income held by higher percentiles disproportionately.

Moreover, the probability that someone with a wage w_z reaches the top q (i.e., the top

0.1%) of income earners is

$$\Pr(\text{skill} = z | \text{top } q) = \frac{\ell_z w_z^{1/\alpha_{net}^*}}{\sum_v \ell_v w_v^{1/\alpha_{net}^*}},$$

which implies that the share of labor income among the top q is

$$\frac{\mathbb{E}[\text{labor income} | \text{top } q]}{\mathbb{E}[\text{income} | \text{top } q]} = (1 - \alpha_{net}^*) q^{\alpha_{net}^*} \frac{\sum_z \ell_z w_z^{1+1/\alpha_{net}^*}}{\left(\sum_z \ell_z w_z^{1/\alpha_{net}^*}\right)^{1+\alpha_{net}^*}}.$$

As this formula shows, as we move up the income distribution, households derive more of their income from capital ownership. This reflects the fact that the tail of the income distribution is increasingly made of successful investors for whom labor income represents a small part of their earnings.

As shown by these formulas, the effect of technology on the composition of top incomes is ambiguous. Technologies that increase the relative wage of high wage earners, like the automation of middle and low-skill jobs and skill-biased technical change more generally, will increase the share of high skill workers and the importance of labor income among top earners. On the other hand, by raising the net capital share, automation also makes relative wages less important in determining the composition of top income earners, and reduces the importance of labor income at the top of the distribution. Intuitively, many more households that managed to accumulate assets for long periods at the higher rate brought by automation will be among top earners independently of their wage.

Other mechanisms. Starting with [Meade \(1964\)](#) and more recently with [Piketty \(2014\)](#), several authors have emphasized that a rise in the net capital share will generate income inequality via a compositional effect. The argument is that because capital is more unequally distributed than labor income, a rise in the relative importance of capital would contribute to inequality. This argument differs from ours, since we emphasize how technology might lead to a more concentrated ownership of wealth, with major implications for income inequality.

We can use our model to illustrate the differences between compositional effects and our mechanism. As above, denote by $S(q)$ the top q percent income share. Also denote by $\tilde{S}_k(q)$ and $\tilde{S}_\ell(q)$ the share of aggregate capital income and labor income earned by the top q percent of the distribution of *total* income. The appendix shows that

$$S(q) = \alpha_{net} \times \tilde{S}_k(q) + (1 - \alpha_{net}) \times \tilde{S}_\ell(q).$$

This simple formula is precisely the one derived by [Meade \(1964, p.34\)](#) and we can use it to

decompose changes over time in $S(q)$ as

$$dS(q) = \underbrace{(\tilde{S}_k(q) - \tilde{S}_\ell(q)) \times d\alpha_{net}}_{\text{Compositional effect at } q} + \underbrace{\alpha_{net} \times d\tilde{S}_k(q) + (1 - \alpha_{net}) \times d\tilde{S}_\ell(q)}_{\text{Changes in within factor distribution at } q}. \quad (13)$$

This decomposition highlights the difference between compositional effects and our mechanism. In our mechanism, automation increases the concentration of capital ownership, leading to a rise in income inequality via changes in the within factor term.

The decomposition also shows that composition effects will be small in practice. For example, for the top 1% in the US in 1980 we had $\tilde{S}_\ell(q) = 8.3\%$ and $\tilde{S}_k(q) = 15.3\%$, so that an increase in the net capital share of one percentage point would yield an increase in the top 1% income share of only $(\tilde{S}_k(q) - \tilde{S}_\ell(q)) \times d\alpha_{net} = 0.07 \times 0.01 = 0.07$ percentage points.²² As we will see, these compositional effects are small when compared to the effects in our model and in the data.

2 Extended Model: Multiple Assets and Growth

2.1 An Extended Version of the Model

We now present an extended version of our model that clarifies how automation and other concurrent changes in capital markets affect the returns of different assets. This extension also identifies the key returns that are linked to rising wealth inequality in our theory, and explains how these returns can be measured and estimated.

We modify the baseline model in four ways:

1. We assume that there are two assets, risky capital (or equity) and safe bonds, and that only a fraction of the population can invest in capital, but are constrained to do so because of risk or financial frictions.
2. We include markups in the production sector, which allows for the possibility that part of the decline in the labor share might be due to markups, and that the return to capital may include some of the profits generated by markups.
3. We model the taxation of capital, which will also affect capital markets.
4. We allow for factor-neutral technological change causing sustained growth at a rate g .

All proofs and derivations are in Appendix E.

²²In 1980, the top 1% earned 10% of US income. 35% of the income of that group was capital income and 65% was labor income (as opposed to a 20% and 80% breakdown in the aggregate) (see [Piketty and Saez, 2003](#)). Using these numbers, we have $\tilde{S}_k(q) = \frac{\alpha_{net}(q)}{\alpha_{net}} S(q) = 15.3\%$ and $\tilde{S}_\ell(q) = \frac{1 - \alpha_{net}(q)}{1 - \alpha_{net}} S(q) = 8.3\%$.

Households and Investors. Households maximize their expected lifetime utility. Because we will want to separately analyze changes in risk aversion and intertemporal substitution, we work with Duffie-Epstein-Zin preferences

$$v_0 = \mathbb{E}_0 \int_0^\infty f(c_t, v_t) dt \quad \text{with} \quad f(c, v) = \frac{\rho(1-\vartheta)v}{1-\sigma} \left(\left(\frac{c}{((1-\vartheta)v)^{1/(1-\vartheta)}} \right)^{1-\sigma} - 1 \right),$$

where $\vartheta > 0$ measures risk aversion and $\sigma > 0$ is the inverse of the IES. As in the baseline model, households face dissipation shocks at rate p so that the discount rate ρ equals $\varrho + p$.

A fraction $\chi \in (0, 1)$ of households can invest in both capital and safe bonds. We refer to these households as investors, and assume that investors come with equal probability from all skill groups. The remaining $1 - \chi$ households can only invest in bonds.

The bond pays a safe return r_B . The return to capital instead features idiosyncratic risk: the after-tax return over a time interval of length dt equals $dR_t = r_K dt + \nu dW_t$, where r_K is the average after-tax return per dollar of capital, ν is the standard deviation of returns, and W_t is an idiosyncratic standard Brownian motion. This return risk might arise from shocks to the productivity of businesses or might also capture idiosyncratic capital gains.

The budget constraint of an investor with wage w_z is therefore given by

$$da_{z,t} + db_{z,t} = (r_K a_{z,t} + r_B b_{z,t} + w_z + T - c_{z,t})dt + a_{z,t} \nu dW_t.$$

Investors also face a borrowing constraint:

$$-b_{z,t} \leq \theta a_{z,t} + \frac{w_z + T}{r_B - g},$$

where $\theta \in (0, 1]$ parameterizes the extent to which investors can pledge their capital.²³ Investors thus solve a portfolio allocation problem similar to the one in Merton (1969). The remaining non-investor households only have access to the bond market and their budget constraint is $\dot{b}_{z,t} = r_B b_{z,t} + w_z + T - c_{z,t}$.

Taxation. Capital income from bond holdings and equity is taxed at a rate τ and the revenue is returned to households via a common lump-sum transfer T . To simplify the exposition, we assume that idiosyncratic capital returns—the term $a_{z,t} \nu dW_t$ in investors' budget constraint—are not taxed.²⁴

Technology and Market Structure. Firms operate the same technology as in our baseline model but now charge a constant markup $\varphi \geq 1$ with all profits accruing to equity owners. In addition, to generate balanced growth, we assume that ψ_z grows at a constant

²³The usual formulation used in the literature is $-b_{z,t} \leq \theta a_{z,t}$. Relative to this, our formulation assumes that human wealth is pledgeable, which makes the model more tractable.

²⁴If these were taxed, the after-tax income risk would be $(1 - \tau)a_{z,t} \nu dW_t$, and taxes would reduce capital income risk. Instead, here we are only interested in the effect of taxes on after-tax returns r_K and r_B .

rate of $g > 0$ for all skill groups z . Finally, we assume throughout that Assumption 1 holds, so that output is given by (3).

Balanced-growth Equilibrium (BGE). In a balanced-growth equilibrium, output, capital, bond holdings, and wages grow at a rate $g > 0$, and the returns r_K and r_B are constant, with $r_B > g$ so that human wealth $(w_z + T)/(r_B - g)$ is finite. Moreover, in the main text, we will focus on characterizing the BGE of a closed economy in which bonds are in zero net supply. We assume that $\rho + (\sigma - 1)g > 0$, which is a sufficient condition to ensure that the equilibrium exists and features finite wealth.²⁵

Define the return on investors' wealth as

$$r_W := \underbrace{\kappa r_K + (1 - \kappa)r_B}_{\text{portfolio return}} + \underbrace{\frac{1}{2}(\sigma - 1)\vartheta\nu^2\kappa^2}_{\text{risk adjustment}} \quad (14)$$

where $\kappa := \min\left\{\frac{1}{1-\theta}, \frac{r_K - r_B}{\vartheta\nu^2}\right\}$ denotes the share of investors' effective wealth invested in risky capital. The formula for the return on investors' wealth takes the return on their portfolio, $\kappa r_K + (1 - \kappa)r_B$, and adjusts it by the extra term $\frac{1}{2}(\sigma - 1)\vartheta\nu^2\kappa^2$, which accounts for the effect of capital income risk on investors' saving decisions.²⁶ Appendix E shows that r_W is the relevant return governing the rate at which investors accumulate wealth. In particular, the effective wealth of investors follows a random growth process of the form

$$dx_{z,t} = \frac{r_W - \rho}{\sigma} x_{z,t} dt + \kappa\nu x_{z,t} dW_t,$$

where, in addition, $x_{z,t}$ is reset with probability p . Non-investors solve the same problem as in the baseline model so that Lemma 1 applies and their effective wealth evolves as

$$\dot{x}_{z,t} = \frac{r_B - \rho}{\sigma} x_{z,t},$$

and resets with probability p .

In what follows, define the *after-tax share of capital inclusive of profits* in GDP

$$\tilde{\alpha} := (1 - \tau) \left(\frac{\varphi - 1}{\varphi} + \frac{\alpha}{\varphi} \right), \quad (15)$$

²⁵Alternatively, we can assume that bond markets are open, so that r_B is fixed in international markets at a level \bar{r}_B , with $g < \bar{r}_B < \rho + \sigma(p + g)$. This case is analyzed in the Appendix and delivers similar insights.

²⁶In the presence of capital income risk, investors recognize that a given investment will bring them a lower future consumption equivalent. This generates an income effect (the lower value of future consumption pushes current consumption down and savings up) and a substitution effect (the lower consumption equivalent pushes current consumption up and savings down). For $\sigma = 1$ these effects cancel, but for $\sigma > 1$, as in most calibrations of the (inverse of the) intertemporal elasticity of substitution, the income effect dominates, and capital income risk leads to higher savings, which shows up in our formula as if investors faced a higher return. The observation that capital income risk has an ambiguous effect on savings that depends on the (inverse of the) intertemporal elasticity of substitution σ is in line with Obstfeld (1994) and Angeletos (2007).

which increases in α , φ and the keep rate $1 - \tau$. As we will see, all these primitives have the same effects on returns and wealth inequality operating through $\tilde{\alpha}$.

Proposition 4 *There exists a threshold constraint-tightness $\bar{\theta} \in (0, 1)$ such that:*

- *If investors are risk neutral and $\theta \geq \bar{\theta}$, there is a unique BGE. In this equilibrium, $r_W^* = r_K^* = r_B^* = r^*$, where*

$$\frac{r^* - \rho - \sigma g}{\sigma(p + g) + \rho - r^*} = \frac{\tilde{\alpha}}{1 - \tilde{\alpha}} \frac{r^* - g}{r^* + \delta}. \quad (16)$$

- *If investors are risk averse or if $\theta < \bar{\theta}$, there is a BGE. In any BGE,*

$$r_W^*, r_K^* > r_B^*, \rho + \sigma g.$$

- *In any BGE, the distributions of wealth and income have Pareto tails with tail index*

$$\frac{1}{\zeta} := \frac{r_W^* - \rho - \sigma g - \frac{\sigma \kappa^{*2} \nu^2}{2} + \sqrt{\left(r_W^* - \rho - \sigma g - \frac{\sigma \kappa^{*2} \nu^2}{2}\right)^2 + 2\sigma^2 \kappa^{*2} \nu^2 p}}{2p\sigma}. \quad (17)$$

The first part of the Proposition shows that, when investors are risk neutral and their borrowing constraint does not bind, the return of capital and bonds is equalized. The extended model then collapses to a one-asset model which is similar to our baseline model except that it features profits, capital taxes and growth. The common return that equalizes supply and demand is given by the solution to (16), which slightly generalizes equation (5) in the baseline model. This return now satisfies a generalized version of equation (6):

$$r^* = \rho + \sigma g + p\sigma \alpha_{net}^*(\tilde{\alpha}), \quad \alpha_{net}^*(\tilde{\alpha}) := \frac{\tilde{\alpha} - \delta K/Y - gK/Y}{1 - \delta K/Y - gK/Y}, \quad (18)$$

with the only difference that now, $\alpha_{net}^*(\tilde{\alpha})$ is the capital share net of depreciation, capital taxes, and growth (but including all profits).

The second part of the proposition shows that, if investors are risk averse or if financial frictions bind, the equilibrium risky and safe returns instead differ and the model no longer collapses to a one-asset model. In such an equilibrium, the return to capital and investors' wealth exceeds the return on bonds, and also $\rho + \sigma g$, which is the level of returns in a BGE of the representative-household benchmark.

The third part of the proposition characterizes the tail properties of the wealth and income distributions. When they are not hit by a dissipation shock, investors accumulate wealth at a rate $(r_W^* - \rho)/\sigma$ which exceeds the rate at which the average household accumulates wealth and the rate at which wages grow, g – hence the dependence of the tail index in (17) on $r_W^* - \rho - \sigma g$. In fact, because returns are stochastic, some investors obtain even larger returns in a sustained way. As a result, the tail index of the wealth and income

distributions depends not only on $r_W^* - \rho - \sigma g$, but also on κ —determining their exposure to risk and therefore the possibility that some of them will obtain higher-than average returns. This second term explains why relative to the formula for the tail index in Proposition 3, the extended model with additional capital income risk generates higher wealth inequality.²⁷

Our extended model also provides a more realistic description of the process by which investors reach and leave the top of the wealth distribution. Because of the additional capital income risk, some investors will quickly reach the top after obtaining a series of high returns and rapidly growing their fortunes. The top of the wealth distribution will therefore be populated by investors who have not experienced dissipation shocks in a long period—being able to grow their wealth over time and presumably across generations—and also by newcomers who grew their fortunes quickly. By the same token, there is now much more churn among this group, with investors dropping out of the top after a dissipation shock or obtaining a series of low returns.²⁸

In what follows, define the *return premium* as the gap between the return to investors’ wealth r_W^* and $\rho + \sigma g$. Using (14), this return premium can be written as

$$r_W^* - \rho - \sigma g := \underbrace{\kappa^* r_K^* + (1 - \kappa^*) r_B^*}_{\text{portfolio return}} + \underbrace{\frac{1}{2}(\sigma - 1)\vartheta\nu^2\kappa^{*2}}_{\text{risk adjustment}} - \underbrace{(\rho + \sigma g)}_{\text{usual BGP return}} . \quad (19)$$

As we discussed above and can be seen from (17), the return premium $r_W^* - \rho - \sigma g$ is the key object determining top inequality in our model.

The expression for the return premium (19) and its link to top inequality in equation (17) are the key results of the extended model. The two expressions identify which returns are linked to rising wealth inequality in our theory and answer the question: how does inequality respond to a change in equilibrium safe returns r_B^* or returns to capital r_K^* ? The answer depends on investors’ portfolio shares κ^* . In particular, note that if investors are primarily invested in business capital (κ^* close to one), inequality will be largely unaffected by movements in the safe return r_B^* . In fact, (19) shows that if investors are leveraged, $\kappa^* > 1$, it is a *reduction* in safe rates that increases inequality. Instead, if investors at the top of the wealth distribution are leveraged or hold few safe assets in their portfolios, the return to business capital r_K^* is tightly connected to inequality.

In a way, our formula (19) generalizes Piketty’s (2014) “ $r - g$ ” which is the special case of the return premium (19) with logarithmic preferences $\sigma = 1$ and zero discounting

²⁷In particular, when $\nu = 0$, the formula in (17) boils down to $(r_W^* - \rho - \sigma g)/(\sigma p)$, which is the natural extension of the formula in our baseline model to a context with growth. When $\nu > 0$, this extra source of capital income risk generates a fatter tail, as some investors are lucky to obtain high returns consistently.

²⁸Recent work documents the importance of churn at the top of the wealth distribution, as well as the contribution of newcomers to the rising share of wealth held at the top (Gomez, 2018; Zheng, 2019). Random growth processes offer a natural way of thinking about this churn. For example, consider the group of investors at the top of the wealth distribution at time t_0 . In both our baseline model and in the extended version, the wealth of these investors grows at an average rate of g between t_0 and t , for any $t > t_0$. On the other hand, the wealth of investors who reach the top between t_0 and t grows at a rate higher than g .

$\rho = 0$ (so that investors reinvest all their wealth). However, as we show next, our theory also recognizes that the return premium is an endogenous object and that its long-run level depends on secular forces affecting capital markets such as automation.

Proposition 5 *Suppose that investors are risk averse and/or $\theta < \bar{\theta}$. There exists a threshold $\bar{\alpha} \in (0, 1]$ such that, for $\tilde{\alpha} < \bar{\alpha}$, the balanced-growth equilibrium in the closed economy is unique, and following an increase in $\tilde{\alpha}$, we have that:*

- *The return to wealth r_W^* , the return premium $r_W^* - \rho - \sigma g$, and the return spread $r_K^* - r_B^*$ all strictly increase, and the portfolio share of capital κ^* weakly increases;*
- *Top tail inequality $1/\zeta^*$ in (17) strictly increases.*

The key takeaway from this proposition is that automation results in a higher return to wealth r_W^* , a higher return premium $r_W^* - \rho - \sigma g$, and higher wealth inequality.²⁹ Part of the increase in tail inequality is due to the effect described in our baseline model, namely that investors' return to wealth and return premium increase and that they therefore accumulate wealth at a faster rate. However, our extended model also features an additional effect: because the return spread $r_K^* - r_B^*$ increases, unconstrained investors expose themselves more to risky capital. On the other hand, automation has an ambiguous effect on the bond rate r_B^* (or no effects in the open-economy version of our extended model). Thus, changes in the return premium $r_W^* - \rho - \sigma g$ and the spread $r_K^* - r_B^*$ are more informative for understanding wealth inequality than changes in the bond rate r_B^* .

2.2 Other Drivers of an Increasing Net Capital Share

An immediate implication of our extended model is that automation is not the only secular change that has the potential to explain rising wealth inequality. Instead higher markups (φ) or lower capital taxation (τ) have similar effects. In particular both higher markups and lower capital taxes also increase the *after-tax share of capital inclusive of profits* $\tilde{\alpha}$ defined in equation (15) and Proposition 5 showed that it is this object that determines asset returns and top wealth inequality. Intuitively, all of automation, higher markups and lower capital taxes increase the demand for capital and hence the return to wealth, the return premium, and the spread between the return to capital and safe assets.

While our paper's main focus is on understanding the consequences of automation on income and wealth inequality, we view our framework's ability to speak to these alternative trends as one of its key advantages. In our quantitative exercise below and the discussion that follows, we study the effects of automation on inequality. But readers should keep in

²⁹The result for r_W^* holds even for large values of $\tilde{\alpha} > \bar{\alpha}$ provided that the equilibrium is unique. If there were multiple equilibria, the result would still apply to the equilibria with the lowest and the largest levels for r_W^* . On the other hand, as discussed in detail in the Appendix, $r_K^* - r_B^*$ might fall for very high values of $\tilde{\alpha}$ (above 0.85 in our calibrations with high levels of risk aversion).

mind that the model-implied effects of automation on capital markets, returns, and wealth inequality at the top of the distribution are identical to those generated by an increase in markups or a reduction in capital taxes causing the same reduction in the after-tax share of capital inclusive of profits. That being said, while these other secular trends affecting capital markets have similar implications as automation for returns and wealth inequality, they do have different implications for wage inequality and aggregates such as investment.

2.3 Has the Return Premium Risen?

We now turn to a measurement exercise that shows that the return premium has risen since 1980, at the same time as we have seen a large secular decline in safe rates. We here highlight the main takeaways and relegate a more detailed discussion to Online Appendix G.

As a first step, Figure 5 plots estimates of returns on bonds—our proxy for r_B . The red line with round markers depicts the long-run return to safe assets, which we measure using Holston, Laubach and Williams’ (2017) estimate of the real natural interest rate for the United States. The safe rate has been on a secular decline since the 1970s, falling by about 1.5 percentage points in 1960–2007 and then experiencing a sharp and persistent decline of 1.5 percentage points after the 2007 Great recession. A significant proportion of this decline can be explained by demographic factors (captured by a lower ρ in our model) and slower trend productivity growth (see Rachel and Summers, 2019; Auclert et al., 2019); while developments in international markets, like the “savings glut,” could have also contributed to the decline.




Figure 5: Estimates of the return to US business capital, r_K , and the safe rate, r_B . Notes—see Appendix G for data sources and measurement details.

More importantly, as our theoretical discussion has highlighted, capital income of U.S. households, the return to wealth and the return premium may be largely decoupled from safe asset returns. This is a particular concern when investors at the top of the wealth distribution are leveraged or hold few safe assets in their portfolios.

Figure 5 therefore plots other return series, namely measures of returns to business capital—our proxy for r_K —in the US economy using national accounts data. Our theoretical discussion suggests that these return measures are more relevant for wealth inequality.³⁰ The first series is from [Gomme, Ravikumar and Rupert \(2011\)](#), who measure the return to U.S. business capital from the “firm side.” Their approach defines capital income Y_k as net output minus wage payments, benefits and taxes and computes the return to capital as $r_K := Y_k/K + p_K$ where K is the stock of capital measured at replacement cost and p_K captures capital gains arising from changes in the price of capital over time.³¹ We plot a version of their measure that abstracts from variation in capital gains over time and sets capital gains to their average over the period of analysis. Although this measure reflects the returns to investing in U.S. businesses, it does not tell us whether it is U.S. investors and households who are reaping the higher returns. Figure 5 therefore also plots a measure of returns to business capital from the “household side.” We use the same formula for returns but measure the components differently. In particular, we define capital income Y_k as dividends or profits from business ownership obtained by U.S. households net of personal income taxes (from NIPA); K by the market value of equity held by households (from the Integrated Macroeconomic Accounts—IMAs); and capital gains from revaluations (also from the IMAs). Using this approach, we constructed measures of returns to capital in the corporate and noncorporate sectors. As before, we start with a baseline that abstracts from time variation in capital gains.

In stark contrast with what we found for the safe rate, all measures of returns to business capital exhibit an increasing trend between 1970 and 2017, rising by 1 to 2 percentage points.³² This upward trend in returns is driven by the income component of returns, and not by higher capital gains over time. In particular, Appendix Figure A7 shows that, relative to the income component, the revaluation component is highly volatile, with large fluctuations at high frequencies but no visible trend.

To illustrate the implications of these divergent return trends for wealth inequality, we provide several estimates of the return to investors’ wealth r_W and the return premium $r_W - \rho - \sigma g$. Because some of their components in equations (14) and (19) are hard to measure empirically, we proxy them as follows.

We proxy the return r_W by imposing two key simplifications. First, we simply ignore the risk-adjustment term in equation (14) which would require us to take a stance on the risk

³⁰These calculations exclude housing and focus only on productive or business capital. As shown by [Gomme, Ravikumar and Rupert \(2011\)](#), the resulting patterns are broadly similar including housing.

³¹This approach to measuring returns is also used by [Poterba \(1998\)](#); [Mulligan \(2002\)](#); [Karabarounis and Neiman \(2018\)](#) and the BLS multifactor productivity program.

³²The divergence between safe rates and returns computed from income flows is at the heart of the literature on “factorless income.” The basic idea is that using risk-free interest rates to impute net capital income as $r_B \times K$ (or gross capital income via a user cost formula) leads to a large and volatile residual category of “factorless income” (see [Karabarounis and Neiman, 2018](#)). In our extended model, factorless income is equal to $(r_K - r_B) \times K$, and rises in importance with automation, financial frictions, and risk compensations. This aligns with the argument in [Karabarounis and Neiman \(2018\)](#) that factorless income could reflect a mismeasurement of returns (what they refer to as “case R”).

preferences of investors.³³ Second, we exclude all other assets besides capital and bonds, such as housing. We thus proxy r_W as

$$\hat{r}_W = \text{share equity in net assets} \times r_K + \text{share bonds in net assets} \times r_B,$$

where r_K is the mean return to business capital obtained by U.S. households (a weighted average of the two lines from IMA/NIPAs in Figure 5), r_B is given by Holston, Laubach and Williams (2017) estimate of the real natural interest rate for the United States, and the portfolio shares will be described momentarily. Next, to proxy the return premium $r_W - \rho - \sigma g$ we not only replace r_W by its proxy \hat{r}_W (as just discussed) but also ignore the discounting term ρ . This latter simplification is arguably conservative because the effective discount rate has arguably declined over time, e.g. due to demographic trends, so that it would have pushed up the return premium.

To compute our estimates of the return to investors' wealth and the return premium, we use several different portfolio share measures. In particular, some measures use aggregate data on all households and others use data for the top 1% of the income distribution so as to capture the idea that it is the returns of those households that are most relevant for the mechanism in our theory. When focusing on all households, we have data on portfolio shares from the IMAs, the DINAs, and also from the Distributional Financial Accounts (or DFAs), which are available since 1989 and are maintained by the Federal Reserve. For the top 1%, we have data on portfolio shares from the DINAs and the DFAs. Finally note that in all calculations, we compute "bonds" as fixed income assets net of debt and thus the share of bonds can be negative if households are leveraged.

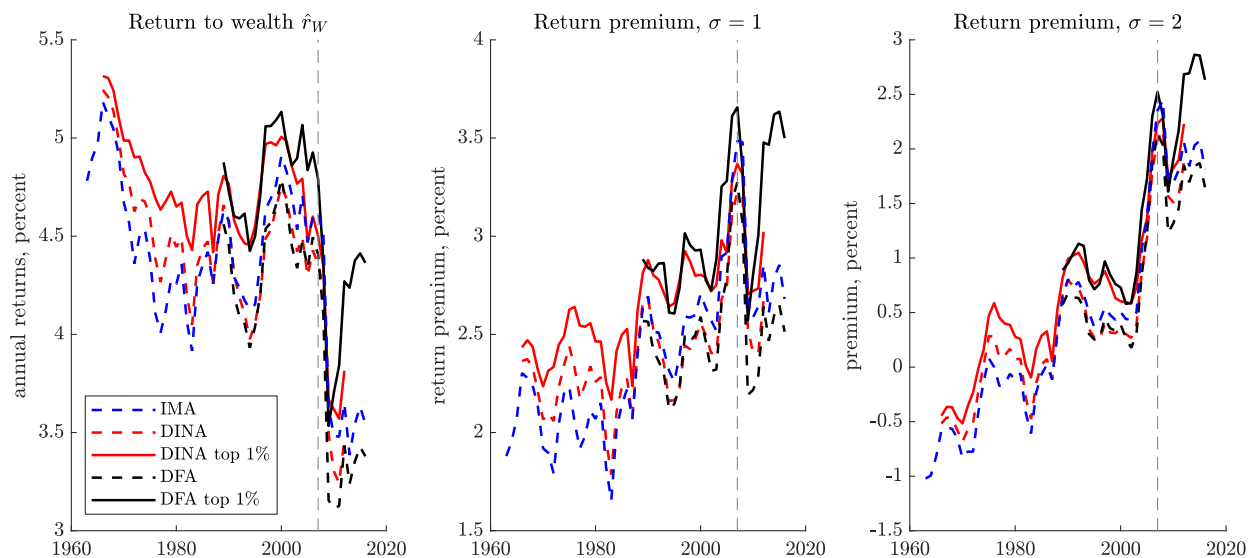


Figure 6: Estimated returns to wealth (left panel) and the return premium for $\sigma = 1$ (middle panel) and $\sigma = 2$ (right panel). *Notes*—see Appendix G for details.

³³We conjecture that taking this risk-adjustment term into account would contribute to a larger increase in the return to investors' wealth r_W . We therefore view ignoring it as conservative.

The left panel in Figure 6 presents estimates for portfolio returns based on the observed portfolio composition of all households and the top 1% income earners. The figure shows that in 1960–2007, portfolio returns followed a U-shape pattern, dropping from 5 to 4 percentage points per year between 1960 and 1980, and then rising again from 4 to 4.7 percentage points between 1980 and 2007. As we show in the Appendix, the early decline in portfolio returns during 1960–1980 was largely due to changes in the growth rate of the economy but also due to a rising share of pension wealth in households’ portfolio over time. The upward trend between 1980 and 2007 was interrupted by the Great Recession of 2008, when there was a large drop in the share of equity in net assets (presumably driven by a loss of business value) and a sizable decline in the safe rate r_B .

The middle and right panels turn to our proxy of the return premium which our theory identifies as the key object determining top tail inequality. The panel in the middle uses an inverse intertemporal elasticity of $\sigma = 1$, which we view as conservative. The right panel uses our baseline value of $\sigma = 2$. We measure g using a 10-year moving average of the CBO’s estimation of the potential growth rate of the economy. In both cases, we find that the return premium before 1980 was roughly stable, suggesting that the deceleration of growth was one of the key driving forces behind the U-shape pattern behind portfolio returns documented in the left panel. More importantly, we find that the return premium rose by 1–2 percentage points since 1980, with a large but temporary drop during the Great Recession. As a whole, we view these estimates as supporting the idea that returns to wealth and in particular the return premium rose between 1980 and 2007 by 0.7–2 percentage points. This growing return premium is thus a potential driver of the observed rise in wealth and income inequality.

3 Calibration and Implications for Aggregates

3.1 Calibration

To study the effects of automation, we feed changes in α_z into the model and explore the consequences of this technological change for aggregates and inequality using the extended version of the model. We will focus on changes in automation between 1980 and 2014, a period with a marked shift in technology towards automation (Acemoglu and Restrepo, 2019), especially of routine tasks both in manufacturing and services (Autor, Levy and Murnane, 2003; Acemoglu and Autor, 2011).³⁴ Because we are not feeding any other drivers of inequality into the model besides automation, the results here must be interpreted as isolating the effects of automation on the economy while holding all other potential drivers

³⁴Our focus on this period does not imply that there was no automation before then. As discussed in Acemoglu and Restrepo (2019), before 1980 jobs were automated in some specific industries and tasks, but automation was counteracted by other technological improvements that raised labor shares in other industries or introduced new labor-intensive roles for labor in production. As a result, the aggregate labor share—the key object determining how technology affects wealth inequality—remained stable during this period. Technological change might have contributed to rising *wage* inequality before 1980, but because the labor share remained constant, our mechanism did not contribute to rising *wealth* inequality back then.

of distributional and macroeconomic trends at their 1980 levels.

To bring the model to the data, we interpret z as indexing the group of workers in a given percentile of the wage distribution, so that we have 100 skill groups. The main ingredient in our calibration is a measure of how automated the tasks being performed by workers in each percentile of the wage distribution have been over time, $\alpha_{z,t}$. We think of changes in $\alpha_{z,t}$ as been driven by the automation of routine tasks. To operationalize this view, we assume that all routine tasks have been automated at the same rate over time starting in 1980, that other tasks did not experience additional automation during this period, and that each skill group had a fixed relative exposure to routine tasks. Moreover, to benchmark the size of this shock, we assume that the rate at which routine tasks were automated explains the overall decline in the labor share between 1980 and 2014, though we also present results under different scenarios below.³⁵ Finally, we normalize $\alpha_{z,1980}$ to be equal across all z , which requires $\alpha_{z,1980} = \alpha_{1980} = 34.5\%$ —the gross capital share in this year. As we show in Appendix F, these assumptions imply that

$$\frac{1}{1 - \alpha_{z,2014}} - \frac{1}{1 - \alpha_{z,1980}} = \omega_z^R \left(\frac{1}{1 - \alpha_{2014}} - \frac{1}{1 - \alpha_{1980}} \right), \quad (20)$$

which relates the change in α_z between 1980 and 2014 to the increase in the capital share during this period (from $\alpha_{1980} = 34.5\%$ to $\alpha_{2014} = 42.8\%$, respectively, in the BLS series for the non-financial corporate sector) interacted with the exposure of workers in each percentile to routine jobs, ω_z^R . This exposure measure is in turn defined as the share of labor income derived by workers in percentile z from routine jobs relative to the labor income derived by all workers from routine jobs.³⁶ This procedure results in the 1980 and 2014 values for $\alpha_{z,t}$ plotted in the left panel of Figure 7, which shows greater exposure to the automation of routine jobs among workers at the middle and bottom of the wage distribution. The average change in $\alpha_{z,t}$ (weighted by η_z) is of 8.4 percentage points, which roughly matches the observed decline in the labor share during our period of analysis (8.3 p.p decline).³⁷

³⁵We view this stark assumption as providing a useful benchmark. In particular, our calibration asks: if the observed decline in the labor share were explained by the automation of routine jobs, what would the distributional and aggregate consequences be? We recognize that there are other forces behind the decline in the labor share, including rising markups (Autor et al., 2017; De Loecker and Eeckhout, 2017). The extended model in Section 2 shows that all these potential drivers of the decline in the labor share have similar implications as automation for wealth inequality at the top.

³⁶We measure ω_z^R using the 2000 Census—a point in the middle of the period we study. In our model, the composition and specialization patterns of a skill group are assumed invariant. However, in the data, the composition of workers in a given wage percentile might change over time, as the relative ranking of groups of workers with different characteristics changes. In our baseline calibration, we used the 2000 values for ω_z^R as describing the level of specialization of different groups in routine jobs. We also experimented with measuring ω_z^R using the 1980 Census and obtained similar results. The reason is that ω_z^R is highly correlated over time (the correlation between the 1980 and 2000 measures is of 0.9714).

³⁷The average change in $\alpha_{z,t}$ in the model is slightly larger than the one observed in the data. The reason is that over time, we also have changes in η_z that are not fed to the model. In particular, the observed changes in η_z imply that other technological changes resulted in a reallocation of value added from highly automated skills to less automated ones at the top of the wage distribution.




Figure 7: Calibrated α_z by wage percentile in 1980 and the new steady state (left panel).

Finally, we feed a smooth increase in $\alpha_{z,t}$ from its value in 1980 to its final value in 2014, which results in the smooth change in the aggregate labor share shown in Figure 3 (recall that the labor share is $1 - \alpha_t = \sum \eta_z(1 - \alpha_{z,t})$ in our model).

For the remaining labor parameters, we calibrate η_z to match the wage distribution in 1980 (obtained from the 1980 Census). The η_z 's might have changed over time as a result of other forms of skill biased technical change not modeled here, but we do not explore this possibility. We pick ψ_z to ensure that human labor is 30% more costly than capital in automated tasks. This number is in line with studies exploring the cost-saving gains from industrial robots (see [Acemoglu and Restrepo, 2020](#)). As already discussed, this implies that automating a task reduces its cost by 30%, and that, to a first-order approximation, the productivity gains from the calibrated increase in $\alpha_{z,t}$ between 1980 and 2014 is of only 2.4%—or 10% of the total increase in TFP during this period ([Fernald, 2014](#)).

Turning to investors, we focus on a calibration with risk averse investors with $\gamma = 2$. We also present results for a low risk aversion calibration in the appendix, with $\gamma = 0.1$. We assume that the extent of capital income risk is given by $\nu = 6\%$, and we calibrate the share of investors in the economy to $\chi = 8.1\%$ to match the observed tail index of the income distribution in 1980.³⁸ Finally, we set $\theta = 0.5$, which implies that investors can hold up to two times their effective wealth in risky capital.

The remaining parameters are chosen to match aggregates in 1980 and available estimates of the elasticity of capital supply. We take a capital-output ratio of $K/Y = 3$, which implies a rental rate of capital $R = 11.5\%$. We take a depreciation rate of 5% so that the net capital

³⁸The share of income held by the top 0.1% (inclusive of capital gains) rose from 4% in 1980 to a peak of 12% before the Great Recession ([Jones, 2015](#); [Piketty and Saez, 2003](#)), which implies a tail index of 0.54 in 1980 and 0.7 before the Great Recession. In this model, the tail index of the wealth distribution and the income distribution coincides, and so by matching the tail of the income distribution we cannot match the thicker tail for wealth inequality.

Table 1: List of calibrated parameter values

DESCRIPTION		VALUE	TARGET / SOURCE
PREFERENCES			
σ	Inverse IES	2	Standard calibration
p	Dissipation rate (p.a.)	4.5%	Target capital-supply elasticity $d \ln K/dr \approx 50$
γ	Risk aversion	2	High risk aversion
χ	Share investors	5.8%	Match tail index of income in 1980
ρ	Discount rate (p.a.)	1.31%	Target $r = 6.5\%$
ν	Capital risk	6%	Assumed
θ	Borrowing constraint	0.5	Assumed
TECHNOLOGY			
g	Growth rate of ψ_z	1.5%	Standard calibration
δ	Depreciation rate (p.a.)	5%	Standard calibration
A	Hicks-neutral productivity term	0.14	Y/L in 1980
η_z	Skill demand shifters in 1980	vector	Wage levels in 1980 Census/ACS
$\psi_{z,1980}$	Productivity of labor relative to capital	vector	Automation reduces costs by 30% ($= \frac{w_z}{\psi_z R}$)
AUTOMATION SHOCK			
α_{1980}	Capital share in 1980	0.345	BLS labor share in 1980
α_{2014}	Capital share in 2014	0.428	BLS labor share in 2014
ω_z^R	Routine jobs share in each pctl in 2000	vector	Acemoglu-Autor + 2000 Census/ACS

Notes: The table provides the parameters used in our baseline calibration of the model. For details, see the main text and Appendix F.

share equals 23% and the return to wealth equals $r = 6.5\%$. We pick ρ to ensure a 6.5% return to wealth and pick p to target a long-run elasticity of capital supply $d \ln K/dr$ of about 50. We view this choice of p as conservative, in the sense that much of the evidence suggests a more inelastic supply of capital (see Appendix F). Table 1 summarizes the parameters used in our exercise.

3.2 Implications for Aggregates and Returns

With this calibration in hand, we now explore the effects of changes in the α_z 's on macroeconomic aggregates and asset returns. In the next section, we will then turn to the distributional effects of these changes.

Figure 3 already plotted the transition dynamics of our baseline model following a gradual increase in automation. In what follows, we will therefore focus on steady-state comparisons. The results are summarized in Figure 8.

The horizontal axis plots $\alpha = \sum_z \alpha_z \eta_z$, the average extent of automation in the economy. The different panels plot various steady-state aggregates and returns. In line with the discussion in Section 2, automation leads to a sharp rise in the return to capital; while the return on bonds and safe assets increases slightly. As a result, we see a widening gap between r_K and r_B . Because some of the productivity gains from automation accrue to investors, mean wages are stagnant and median wages decline, and we see a modest expansion in output and the capital-output ratio. Finally, despite the large distributional implications

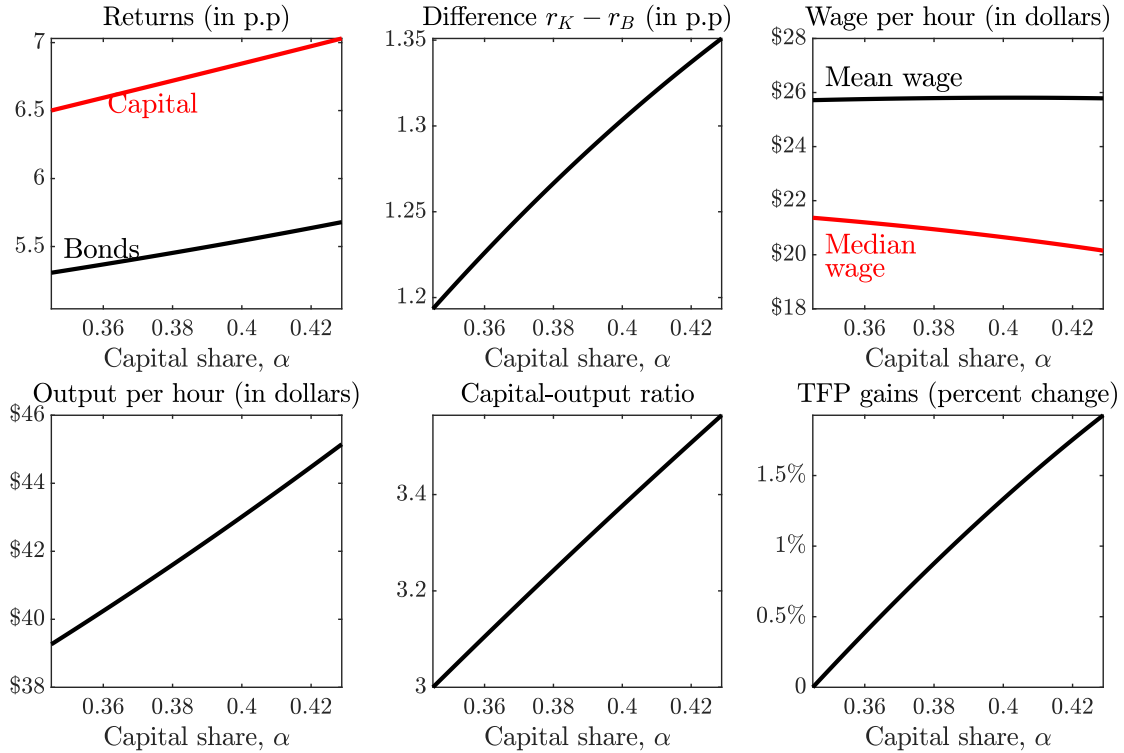


Figure 8: Aggregate effects of changes in $\alpha = \sum_z \alpha_z \eta_z$ in steady state.

that we will document, a rise in the capital share driven by automation generates an increase in TFP of 1.9%, which is smaller than what our first-order approximation suggests but of a similar order of magnitude. In Appendix I.1 we further show that the modest increase in the capital-output ratio implied by our model is in line with that observed in the data.

4 Uneven Growth: Inequality in Model and Data

We now turn to our paper’s main focus, namely to understand inequality trends observed in the data, in particular the pattern of highly uneven growth of different income percentiles since the 1980s. Section 4.1 briefly describes our model’s implications for changes in wages and wage inequality. Section 4.2 then presents our model’s implications for inequality of overall income, including not only wages but also capital income, and Section 4.3 then confronts these model implications with the analogous data. Section 4.4 briefly discusses a discrepancy between model and data regarding the speed at which these changes occur at the top of the income distribution.

4.1 Wage Inequality

As we already saw in Figure 8, average wages in our model are roughly stagnant when the capital share increases. However, the constant average wage masks substantial heterogeneity.

This can be seen in Figure 9 which plots the change in steady-state wages by wage percentile. The real wage of workers below the 80th percentile of the wage distribution declines over time, but the most pronounced effects are for workers at the 25th percentile of the wage distribution, whose real wages fall by 10%. In contrast, the real wages of workers at the 95th percentile of the wage distribution rise by 7%.³⁹ For comparison, the figure also plots on a different vertical axis the observed change in wages by percentile between 1980 and 2012–2016 (using data from the US Census and the American Community Survey (ACS)).




Figure 9: Predicted change in wages by wage percentile (left axis) and observed change in wages by wage percentile (right axis). *Notes*— Observed wage changes computed using the 1980 Census and 2012–2016 ACS. See Appendix F for details.

4.2 Uneven Growth in the Model

We now turn to our model’s implications for income inequality. Figure 10 presents the change in total income for different percentiles of the income distribution across steady states. The figure reveals a pattern of uneven growth. Below the 50th quantile, households experienced declines in total income of 5–10%. Between the 50th and 80th percentile, households experienced stagnant incomes. This is in contrast to the top income quantiles, which experienced an increase in income ranging from 20% (for households in the top 1%–0.5%) to 55% (for the top 0.1%). Although other technologies not modeled here also shifted incomes, the figure shows that automation is capable of generating substantial income gains at the very top of the distribution and stagnant or declining incomes at the bottom.

Both wage and wealth inequality combine to produce the pattern observed in Figure 10. The blue-shaded area plots the contribution of changes in labor income. The fall in real wages for households at the bottom of the wage distribution contributed to declining incomes

³⁹Because our model’s wage distribution is very skewed (just like in the data), the rising wages for the top 25 percent of the wage distribution in Figure 9 are enough to offset the declining wages for the bottom 75 percent of the distribution so that the average wage in Figure 8 is roughly constant.




Figure 10: Predicted change in income by income percentile decomposed into the contribution of capital and labor income. The dotted red line plots the results from the representative-household benchmark. *Notes*—see Appendix E.3 for a detailed explanation of the figure’s construction.

for households at the middle and bottom of the income distribution. The red-shaded area plots the contribution of changes in capital income. This is uniformly positive, as everyone benefits from a higher return to wealth. But the benefits from a higher return to wealth are highly dispersed. People at the bottom of the distribution have few assets and only invest in bonds and safe assets, and so do not benefit as much from an increase in the return to capital. In contrast, a higher return to capital allows investors at the top of the income distribution to accumulate large swaths of wealth and earn a high capital income.

Interestingly, although automation is skill-biased and raises wages at the top of the wage distribution (see Figure 9), the contribution of wages at the top of the overall income distribution is small or even negative. As discussed in section 1.3, there are two effects at play here. On the one hand, the skill-biased nature of automation raises the wages of households at the top of the distribution more than at other percentiles. However, the permanent increase in returns also means that the top of the income distribution becomes increasingly populated by low-wage investors with very high capital incomes. This shift in the composition of top earners dominates at the top of the income distribution and generates the observed low or negative contribution of labor income.⁴⁰

This finding can be contrasted with the representative-household benchmark (see Ap-

⁴⁰Formally, the expected wage of households at the top of the income distribution is $\sum_z \ell_z w_z^{1+1/\alpha_{net}^*} / \sum_z \ell_z w_z^{1/\alpha_{net}^*}$. Skill-biased changes in wages raise this expected wage, but increases in α_{net}^* reduce it. See the related discussion of Proposition A1. Note that as we keep moving up the tail, the contribution of labor income converges to zero since households’ own mostly capital income.

pendix D for details). In this model the steady-state wealth distribution is indeterminate, but one can still trace its evolution starting from a given initial distribution, and so we assume the same initial distribution of wages and wealth as in the extended model with dissipation shocks. The dashed line in Figure 10 shows the change in income for different percentiles of the income distribution in this model. The representative household model does not generate uneven growth: there is a fairly uniform increase in income between 10% and 18% for all quantiles. This reflects two differences. First, there is no wage stagnation in the representative household model, since the infinitely elastic supply of capital implies that all productivity gains from automation accrue to workers. Second, the temporary increase in returns to wealth benefits all households equally, as they are all able to scale their effective wealth by the same amount. This is in contrast to our model with dissipation shocks, where only a few households benefit from the higher return to capital.

How strong is the new mechanism in our model linking changes in the capital share due to automation, returns, and the resulting increase in income inequality? To address this question, Figure 11 plots the behavior of the return premium $r_W^* - \rho - \sigma g$ and the resulting tail index of the income distribution as functions of the capital share in our calibrated model.




Figure 11: Steady-state inequality effects of changes in $\alpha = \sum_z \alpha_z \eta_z$. Appendix Figure A6 reproduces the same figure for an alternative calibration with low risk aversion.

The figure shows that the return premium in our calibrated model rises by 1.2 percentage points from 4 to 5.2 percentage points per year. Despite this relatively modest rise in the return premium, the tail index of the income distribution rises from 0.54 (the targeted level in 1980—see footnote 38) to 0.65, which is roughly 70% of the observed increase in the data. A large part of this increase in tail inequality is due to the effect described in our baseline model, namely that investors’ return to wealth r_W increases and that they therefore accumulate wealth at a faster rate. However, in our extended model with idiosyncratic capital income risk, there is also an additional effect: when the return spread $r_K - r_B$ increases (see Figure 8), investors increase the share of equity in their portfolios and become more exposed to this risky capital. The right panel of the figure shows that the higher return to wealth (the

drift effect) generates an increase in the tail index from 0.54 to 0.63; while greater exposure to risk explains the remaining increase from 0.63 to 0.65.⁴¹

Of course, this steady-state comparison must be taken with a grain of salt, since as we discuss in Section 4.4, the transition dynamics in our model are slow.

4.3 Uneven Growth in the Data

The main implication of our model is that technological change involving the automation of tasks performed by low and middle-wage workers generates the pattern of uneven growth documented in Figure 10: stagnant or decreasing labor incomes at the bottom and rising incomes at the top, with an increasing role for capital income at the top of the distribution.

We now explore whether the evolution of capital and labor income in the US economy has followed a similar pattern of uneven growth. To this end, we use two data sources: the NBER IRS public use sample from 1980 to 2012 and the synthetic micro-files from the Distributional National Accounts (DINAs) of Piketty, Saez and Zucman (2018) for the same period.⁴² The IRS dataset is based on administrative tax records and yields reliable information on incomes for the very top of the distribution. However, it only records *fiscal income* and therefore omits pensions and other sources of income that are tax-exempt. Likewise, fiscal income misses income flows such as corporate taxes, retained earnings, or housing services for home owners, which are part of national income but do not show up in households' tax filings. The DINAs account for this "missing income" by imputing tax-exempt sources of income and measuring corporate income and other sources of capital income using a "capitalization approach." For example, this approach allocates aggregate measures of equity from the IMAs to households according to positive corporate income reported to the IRS, and then imputes corporate income based on these stock measures. Given this, the main advantage of the DINAs are the more comprehensive income concept being used that also adds up to national income but this comes up at the cost of strong assumptions and numerous imputations. Conversely, the main advantage of the IRS data is the transparency and simplicity of the data construction but this comes up at the cost of substantial amounts of missing income and, in particular, missing capital income.

Figure 12 provides our analysis of uneven growth in the United States. Panel (a) uses

⁴¹This second channel is only operational in the calibration of our model with sufficiently high risk aversion and a sufficiently loose financial constraint so that investors are unconstrained, a point we explore in our low-risk aversion calibration in Appendix F.4. Nevertheless, even in our alternative calibration that shuts down this second channel, our model predicts a sizable increase in tail inequality.

⁴²The NBER IRS public use sample is available through the NBER. See <https://users.nber.org/~taxsim/gdb/> and <https://www.nber.org/taxsim-notes.html>. This IRS data is also used in Piketty and Saez (2003). The synthetic micro-files for the Distributional National Accounts can be accessed via Gabriel Zucman's website <http://gabriel-zucman.eu/usdina/>. The DINAs merge semi-public and private IRS data with national accounts and the US Census/ ACS to construct consistent measures of wealth and income across the income and wealth distribution. Their construction is detailed in Piketty, Saez and Zucman (2018). We focus on the period 1980–2012 for which the underlying IRS data used in the DINAs are available, and present additional decompositions for 1980–2007 in the Appendix.

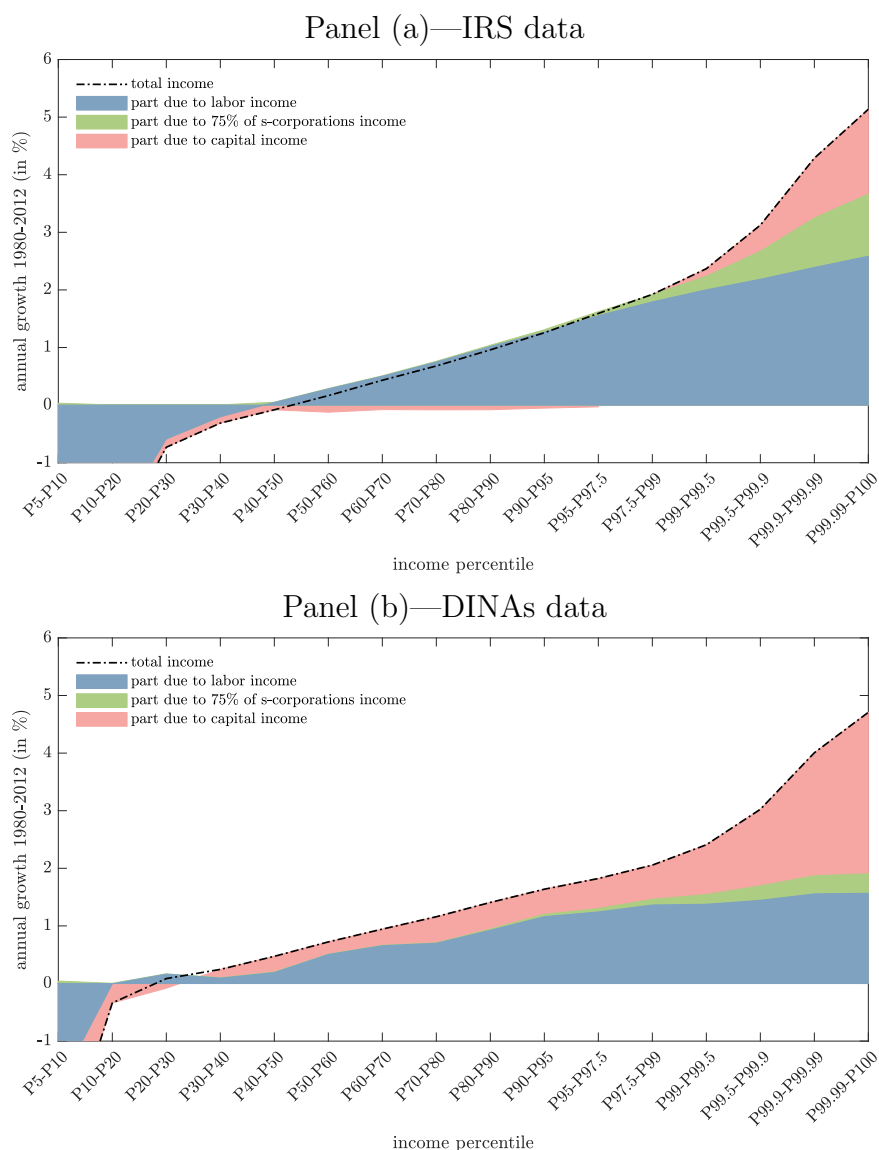


Figure 12: Income growth by percentile between 1980–2012 and decomposition into capital and labor income. *Notes*— see Appendix H for details.

the IRS data and panel (b) the DINAs. The dash-dotted black lines labeled “total income” plot the average annual income growth between 1980 and 2012 for different percentiles of the income distribution. The income percentiles corresponding to the lower half of the distribution have stagnated or declined between 1980 and 2012. In contrast, top income percentiles have grown rapidly. For example, in both data sets, the top 0.1 percentile increased at a yearly rate of 5 percent. The hockey stick shape of income growth at different percentiles in the figure is qualitatively similar to that predicted by our model (see Figure 10).⁴³

⁴³The change in total incomes in panels (a) and (b) is qualitatively similar but there are some quantitative differences in terms of both the average growth and the relative growth across different income percentiles. The blue line in panel (b) labeled “total income” is the same as that in Figure II(a) of [Piketty, Saez and Zucman \(2018\)](#) and they discuss this discrepancy. In particular, they note that their series yields “more growth for the bottom 90% since 1980 than suggested by the fiscal data studied by [Piketty and Saez \(2003\)](#)” which is essentially our IRS series.

What is the contribution of capital and labor income to the pattern of uneven growth in Figure 12? Decomposing income growth into these components is difficult because a substantial share of income is derived from self-employment or ownership of passthrough entities, and these cannot be easily classified as labor or capital income. As a starting point, we follow [Piketty, Saez and Zucman \(2018\)](#) who propose measuring capital income as the sum of corporate income (including the income of both C- and S-corporations) plus 30% of the mixed income from non-corporate passthrough businesses (including self-proprietorships and partnerships). However, [Smith et al. \(2019\)](#) criticize this approach and argue that (in addition to the 70% of non-corporate passthrough income already classified as wage income by [Piketty, Saez and Zucman](#)) also about 75% of S-corporation income should be counted as labor income rather than capital income. The argument is that these 75% of S-corporation income are, in fact, returns to the human capital of owners and are only reported as corporate income to minimize tax obligations. To address this criticism, we also present an alternative measure of capital income that relabels 75% of S-corporation income as labor income.

The shaded areas in both panels of Figure 12 decompose the growth in total income into parts due labor income, capital income, and the ambiguous S-corporation income.⁴⁴ The blue area is the contribution of what both [Piketty, Saez and Zucman \(2018\)](#) and [Smith et al. \(2019\)](#) classify as labor income and the red area is the contribution of what they agree is capital income. The green area in the middle corresponds to the contribution of the 75% of S-corporations income that [Piketty, Saez and Zucman \(2018\)](#) attribute to capital and [Smith et al. \(2019\)](#) attribute to the human capital of owners. The sum of the green and red areas thus provides the contribution of capital to income growth under the classification favored by [Piketty, Saez and Zucman \(2018\)](#), while the red area provides the contribution of capital to income growth under the classification favored by [Smith et al. \(2019\)](#).⁴⁵

In both datasets, capital income is unimportant for income growth at the bottom of the distribution. As in our model, the stagnant incomes at the bottom of the distribution are driven by declining or stagnant real wages. However, further up in the distribution and especially within the top 1 percent, capital income becomes more important. The extent to which capital income accounts for the observed uneven growth and at which percentiles it starts to be important differs between the two datasets. The IRS data in panel (a) suggest that up to the 90th percentile, essentially all growth in total income is accounted for by labor income. Capital income becomes important within the top 1% and especially at the top 0.1% where it accounts for around 60% of the cumulative growth from 1980-2012, or 30% if one focuses only on the red area. In contrast, the DINA data in panel (b) suggest that capital income plays an important role already around the median of the income

⁴⁴Denoting by $y_t(q)$ the q th income percentile at time t and by $y_{\ell,t}(q)$ and $y_{k,t}(q)$ that percentile's labor and capital income, we decompose the annualized growth rate from a base year $t = 0$, $\frac{1}{T} \frac{y_T(q) - y_0(q)}{y_0(q)}$, into a part due to labor income $\frac{1}{T} \frac{y_{\ell,T}(q) - y_{\ell,0}(q)}{y_0(q)}$ and a part due to capital income $\frac{1}{T} \frac{y_{k,T}(q) - y_{k,0}(q)}{y_0(q)}$.

⁴⁵In these figures, we measure S-corporation income using the reported *fiscal income*. The appendix shows an alternative strategy using DINAs where we compute corporate income from S-corporations using the capitalization method. Both approaches yield very similar quantitative results.

distribution, accounts for around half of cumulative growth of the top 1% and is twice as important as labor income for the rise in top 0.01% incomes. When using the DINA data, the relabeling of S-corporations income does not alter these conclusions. These discrepancies across datasets are to be expected, given the differences in coverage between the IRS and DINAs. In particular, the share of capital in fiscal income using the IRS data is of less than 10% whereas this share rises to 25–30% in the DINAs (a level that matches the net capital share in national income). Thus, the IRS data is likely to understate the contribution of capital to income growth, and relabeling part of S-corporations income as labor income is more consequential when using the IRS data.

As a whole, the empirical patterns in both data sets are consistent with those generated by our model. Our analysis suggests that the exact contribution of capital to income growth at the very top of the distribution depends on the dataset and assumptions used, but ranges from 30% (in the IRS data) to 70% (in the DINAs). Moreover, both sources of data suggest that incomes at the middle and bottom have stagnated due to slow labor-income growth.

The main difference between the patterns in these figures and that in our model is that, in the data, wage income plays a more prominent and positive role at the top of the income distribution. This suggests that there are forms of skill-biased technical change other than automation affecting relative wages at the top which are not included in our model (see the discussion of Figure 10).

Our discussion focused on inequality trends observed in the U.S. since the 1980s. But it is natural to ask whether our model’s prediction that rises in the net capital share are accompanied by large increases in top income inequality receives support from other countries or time periods. Appendix I.2 shows that this link is visible across countries and during historical periods of rapid mechanization, such as the industrial revolution in Britain.

4.4 Speed of Top Inequality Dynamics

Figure 10 plotted the model-implied changes in inequality between two steady states. Figure 12, in contrast, showed the analogous changes between two calendar years, 1980 and 2012, and in particular showed that the data features not only very large but also very fast changes in top inequality. It is therefore natural to ask whether our model can replicate such fast transitions. The answer is “no”: while our theory can account for large changes in inequality between steady states, it cannot generate rapid transition dynamics of top income inequality.

That theories like ours cannot generate fast transition dynamics is a known result: [Gabaix et al. \(2016\)](#) have shown that standard theories of the Pareto tails of the income and wealth distributions, which build on a random growth mechanism, generate transition dynamics that are too slow relative to those observed in the data. Our theory is exactly a special case of such a theory – see Proposition 3 – and is therefore subject to the same criticism.

The good news is that we know how to “fix” random growth theories to deliver fast transition dynamics like those observed in the data. [Gabaix et al. \(2016\)](#) show that what is needed

are particular deviations from Gibrat’s law, what they call “type- and scale-dependence.” For the case of wealth dynamics, heterogeneous and persistent rates of return to wealth are one candidate for generating such type- and scale-dependence and seem to be a prevalent feature of the data (Fagereng et al., 2016). Future work should build more quantitatively serious theories of the general-equilibrium interaction between technology and income and wealth distribution that feature these model elements.⁴⁶

5 Conclusion

In this paper, we developed a tractable framework to study the effects of technology on income inequality. Our theory allowed us to go beyond wages and to explore how technology affects wealth inequality and overall income inequality. We used our framework to study the effects of automation and identified a new channel through which technology affects inequality. The benefits of new technologies accrue not only to high-skilled labor but also to owners of capital in the form of higher capital incomes and returns.

There are two fruitful avenues for future work. First, one could use our tractable framework to study the distributional consequences of other types of technical change, changes in market structure and markups, and government policies, like the taxation of capital or estates, or redistributive policies. For example, it could be worthwhile to integrate our model with a theory of international trade so as to examine the effects of globalization on income and wealth distribution, and not just that of wages as is common in the trade literature. Similarly, our analytically tractable theory featuring a less-than-perfectly-elastic capital supply and non-degenerate wealth distribution may serve as a useful laboratory for exploring the optimal taxation of capital income and wealth.

Second, one could devise more elaborate quantitative models to study the effect of technologies on inequality. As explained in the introduction, these more elaborate models should retain the two key features underscored by our analysis: an upward-sloping supply of capital and a return inequality nexus. Our model in Section 2 is a first step in that direction. More elaborate versions of our model could include realistic life-cycle structures, a careful treatment of intergenerational transfers, and additional sources of heterogeneity in portfolio and return rates. Successful quantitative extensions should also include some form of scale-dependence and type-dependence to account for the rapid rise of inequality in the data.

References

- ACEMOGLU, D. (2009): *Modern Economic Growth*. Princeton University Press, New York.
ACEMOGLU, D., AND D. AUTOR (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings* vol. 4 of *Handbook of Labor Economics*, chap. 12, pp. 1043–1171. Elsevier.

⁴⁶Alternatively, theories with changing asset prices—another salient feature of the data that we do not model—are promising for generating fast wealth inequality dynamics.

- ACEMOGLU, D., AND P. RESTREPO (2018): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108(6), 1488–1542.
- (2019): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33(2), 3–30.
- (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128(6), 2188–2244.
- ACHDOU, Y., J. HAN, J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2017): “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach,” NBER Working Papers 23732, National Bureau of Economic Research, Inc.
- ALLEN, R. C. (2009): “Engels’ pause: Technical change, capital accumulation, and inequality in the british industrial revolution,” *Explorations in Economic History*, 46(4), 418–435.
- ANGELETOS, G.-M. (2007): “Uninsured idiosyncratic investment risk and aggregate saving,” *Review of Economic Dynamics*, 10(1), 1 – 30.
- AUCLERT, A., H. MALMBERG, F. MARTENET, AND M. ROGNLIE (2019): “Demographics, Wealth, and Interest Rates in the 21st Century,” Discussion paper, Stanford.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. V. REENEN (2017): “The Fall of the Labor Share and the Rise of Superstar Firms,” NBER Working Papers 23396, National Bureau of Economic Research, Inc.
- AUTOR, D. H., AND D. DORN (2013): “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market,” *American Economic Review*, 103(5), 1553–1597.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96(2), 189–194.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118(4), 1279–1333.
- BENGTSSON, E., AND D. WALDENSTRÖM (2018): “Capital Shares and Income Inequality: Evidence from the Long Run,” *The Journal of Economic History*, 78(03), 712–743.
- BENHABIB, J., AND A. BISIN (2007): “The distribution of wealth: Intergenerational transmission and redistributive policies,” Discussion paper, New York University.
- (2018): “Skewed Wealth Distributions: Theory and Empirics,” *Journal of Economic Literature*, 56(4), 1261–91.
- BENHABIB, J., A. BISIN, AND S. ZHU (2011): “The Distribution of Wealth and Fiscal Policy in Economies With Finitely Lived Agents,” *Econometrica*, 79(1), 123–157.
- (2015): “The wealth distribution in Bewley economies with capital income risk,” *Journal of Economic Theory*, 159(PA), 489–515.
- BLANCHARD, O. J. (1985): “Debt, Deficits, and Finite Horizons,” *Journal of Political Economy*, 93(2), 223–247.
- CABALLERO, R. J., E. FARHI, AND P.-O. GOURINCHAS (2017): “Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares,” *American Economic Review*, 107(5), 614–620.
- CASELLI, F., AND A. MANNING (2018): “Robot arithmetic: new technology and wages,” LSE Research Online Documents on Economics 87371, London School of Economics and Political Science, LSE Library.
- CASELLI, F., AND J. VENTURA (2000): “A Representative Consumer Theory of Distribution,” *American Economic Review*, 90(4), 909–926.
- CASTAÑEDA, A., J. DÍAZ-GIMÉNEZ, AND J.-V. RÍOS-RULL (2003): “Accounting for the U.S. Earnings and Wealth Inequality,” *Journal of Political Economy*, 111(4), 818–857.
- CENSUS BUREAU, U. (2015): “Real Household Income at Selected Percentiles: 1967 to 2014,” Discussion paper, U.S. Census Bureau, Current Population Survey, 1968 to 2015 Annual Social

and Economic Supplements.

- CHATTERJEE, S. (1994): “Transitional dynamics and the distribution of wealth in a neoclassical growth model,” *Journal of Public Economics*, 54(1), 97–119.
- DE LOECKER, J., AND J. ECKHOUT (2017): “The Rise of Market Power and the Macroeconomic Implications,” NBER Working Papers 23687, National Bureau of Economic Research, Inc.
- FAGERENG, A., L. GUIISO, D. MALACRINO, AND L. PISTAFERRI (2016): “Heterogeneity and Persistence in Returns to Wealth,” NBER Working Papers 22822, National Bureau of Economic Research, Inc.
- FARHI, E., AND F. GOURIO (2018): “Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia,” NBER Working Papers 25282, National Bureau of Economic Research, Inc.
- FERNALD, J. (2014): “A quarterly, utilization-adjusted series on total factor productivity,” Federal Reserve Bank of San Francisco.
- GABAIX, X. (2009): “Power Laws in Economics and Finance,” *Annual Review of Economics*, 1(1), 255–293.
- GABAIX, X., J.-M. LASRY, P.-L. LIONS, AND B. MOLL (2016): “The Dynamics of Inequality,” *Econometrica*, 84, 2071–2111.
- GOMEZ, M. (2018): “Displacement and the Rise in Top Wealth Inequality,” Discussion paper, Columbia University.
- GOMME, P., B. RAVIKUMAR, AND P. RUPERT (2011): “The Return to Capital and the Business Cycle,” *Review of Economic Dynamics*, 14(2), 262–278.
- GOURIO, F., AND T. KLIER (2015): “Recent trends in capital accumulation and implications for investment,” Chicago Fed Letters 344, Chicago Fed.
- HÉMOUS, D., AND M. OLSEN (2018): “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality,” Working paper, University of Zurich.
- HOLSTON, K., T. LAUBACH, AND J. C. WILLIAMS (2017): “Measuring the natural rate of interest: International trends and determinants,” *Journal of International Economics*, 108, S59–S75.
- HUBMER, J., P. KRUSELL, AND A. SMITH (2016): “The Historical Evolution of the Wealth Distribution: A Quantitative-Theoretic Investigation,” NBER Working Papers 23011, National Bureau of Economic Research, Inc.
- JAFFE, S., R. MINTON, C. MULLIGAN, AND K. MURPHY (2019): *Chicago Price Theory*. Princeton University Press.
- JONES, C. I. (2015): “Pareto and Piketty: The Macroeconomics of Top Income and Wealth Inequality,” *Journal of Economic Perspectives*, 29(1), 29–46.
- KARABARBOUNIS, L., AND B. NEIMAN (2018): “Accounting for Factorless Income,” NBER Working Papers 24404, National Bureau of Economic Research, Inc.
- KATZ, L., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963–1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107(1), 35–78.
- KAYMAK, B., AND M. POSCHKE (2016): “The evolution of wealth inequality over half a century: The role of taxes, transfers and technology,” *Journal of Monetary Economics*, 77(C), 1–25.
- KEHRIG, M., AND N. VINCENT (2018): “The Micro-Level Anatomy of the Labor Share Decline,” NBER Working Papers 25275, National Bureau of Economic Research, Inc.
- KRUSELL, P., L. E. OHANIAN, J.-V. RÃOS-RULL, AND G. L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–1053.
- KRUSELL, P., AND A. A. SMITH, JR. (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- LINDERT, P. (2000): “Three Centuries of Inequality in Britain and America,” in *Handbook of Income Distribution*, ed. by A. Atkinson, and F. Bourguignon, vol. 1, chap. 03, pp. 167–216. Elsevier, 1 edn.

- MARTINEZ, J. (2019): “Automation, Growth and Factor Shares,” Discussion paper, London Business School.
- MEADE, J. E. (1964): *Efficiency, Equality and the Ownership of Property*. George Allen & Unwin. Available at <http://www.princeton.edu/~moll/meade.pdf>.
- MERTON, R. C. (1969): “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case,” *The Review of Economics and Statistics*, 51(3), 247–57.
- MULLIGAN, C. B. (2002): “Capital, Interest, and Aggregate Intertemporal Substitution,” NBER Working Papers 9373, National Bureau of Economic Research, Inc.
- OBSTFELD, M. (1994): “Risk-taking, global diversification, and growth,” *The American Economic Review*, pp. 1310–1329.
- PIKETTY, T. (2014): *Capital in the Twenty-First Century*. Harvard University Press, Cambridge.
- (2015): “About Capital in the Twenty-First Century,” *American Economic Review*, 105(5), 48–53.
- PIKETTY, T., AND E. SAEZ (2003): “Income Inequality in the United States, 1913-1998,” *The Quarterly Journal of Economics*, 118(1), 1–41.
- PIKETTY, T., E. SAEZ, AND G. ZUCMAN (2018): “Distributional National Accounts: Methods and Estimates for the United States,” *The Quarterly Journal of Economics*, 133(2), 553–609.
- POTERBA, J. M. (1998): “The rate of return to corporate capital and factor shares: new estimates using revised national income accounts and capital stock data,” *Carnegie-Rochester Conference Series on Public Policy*, 48(1), 211–246.
- RACHEL, L., AND L. H. SUMMERS (2019): “On Secular Stagnation in the Industrialized World,” *Brookings Papers on Economic Activity*, Spring, 1–76.
- REIS, R. (2020): “The Constraints on Public Debt when $r < g$ but $g < m$,” Working paper, London School of Economics.
- SACHS, J. D., AND L. J. KOTLIKOFF (2012): “Smart Machines and Long-Term Misery,” Working Paper 18629, National Bureau of Economic Research.
- SARGENT, T., N. WANG, AND J. YANG (2020): “Earnings Growth and the Wealth Distribution,” Working paper, NYU.
- SIMON, H. A. (1965): *The shape of automation for men and management*, vol. 13. Harper & Row, New York.
- SMITH, M., D. YAGAN, O. M. ZIDAR, AND E. ZWICK (2019): “Capitalists in the Twenty-First Century,” NBER Working Papers 25442, National Bureau of Economic Research, Inc.
- STACHURSKI, J., AND A. A. TODA (2018): “An Impossibility Theorem for Wealth in Heterogeneous-agent Models with Limited Heterogeneity,” Papers 1807.08404, arXiv.org.
- STEINDL, J. (1965): *Random Processes and the Growth of Firms: a Study of the Pareto Law*. Griffin, London.
- STRAUB, L. (2019): “Consumption, Savings, and the Distribution of Permanent Income,” Discussion paper, Harvard University.
- WOLD, H. O. A., AND P. WHITTLE (1957): “A Model Explaining the Pareto Distribution of Wealth,” *Econometrica*, 25(4), 591–595.
- ZEIRA, J. (1998): “Workers, Machines, and Economic Growth,” *The Quarterly Journal of Economics*, 113(4), 1091–1117.
- ZHENG, G. (2019): “Wealth Shares in the Long Run,” Discussion paper, UCLA Anderson School of Management.