THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Essays in Applied Macroeconomics

Łukasz Rachel

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy.

London, April 2019

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 55,000 words.

Statement of conjoint work

I certify that Chapter 3 was co-authored with Lawrence H. Summers, and that I contributed over 50% of the work.

Acknowledgements

I wish to express my gratitude to people that accompanied me through this journey. I cannot find enough words to thank my advisor, Ricardo Reis, whose passion, enthusiasm, knowledge and mentorship have been incredibly inspiring. I am also grateful to the rest of the faculty at the LSE, and in particular to Francesco Caselli, Wouter den Haan, Philippe Aghion, Rachel Ngai, Per Krusell, Ethan Ilzetzki and Silvana Tenreyro for their invaluable support.

I thank the Bank of England for financial and other forms of support. In particular, I am indebted to Katharine Neiss, Fergal Shortall and Simon Price for early encouragement to pursue further study, and Gregory Thwaites for supporting me through the early years of the Ph.D. and sharing his experiences. I am indebted to Niki Anderson and other colleagues, whose patience allowed me to make progress at the latter stages.

During the third year of my studies I visited the Department of Economics at Harvard University. I wish to thank the US-UK Fulbright Commission for this incredible opportunity. I thank Ben Broadbent, Charlie Bean and Fergal Shortall, who supported my application to study in the U.S.. Finally, I express my deepest gratitude to Harvard for their incredible hospitality. I thank Larry Summers, Emmanuel Fahri, Elhanan Helpman and Robert Barro for making me feel welcome at Harvard.

It was a pleasure sharing this journey with friends and colleagues at the LSE. I am grateful to Miguel, Thomas, Andrea, Andy, Nikita, Friedrich, Adrien, Daniel, Nicola, Derek, Kangchul and many others for sharing the few difficult and the many fun moments over the last few years.

I wish to thank my family: Jacek, Gosia, Paweł, Stasia and Halina for the unconditional support over these years and before. You have been an incredible source of inspiration, motivation and strength. I am deeply indebted to all of you.

Finally, my deepest gratitude goes to Ania and Helena. Ania is the person that more than anyone else has shared with me this adventure. Her patience, support and love have been unwavering, and I struggle to express how grateful I am to her. I look forward to thanking Hela – once she is born – for the huge dose of motivation and determination that she has given me over the past few months. I am thrilled at what lies ahead for our family trio.

Abstract

This thesis consists of three chapters, examining the impact of technologies and policies on economic growth and macroeconomic dynamics.

Chapter 1 studies the theory of leisure-enhancing technological change. In the model, two-sided businesses – platforms – invest in leisure-enhancing innovations in order to capture consumers' time, attention and data. I show that the economy transitions endogenously from a balanced growth path with only productivity-enhancing technologies to one with leisure-enhancing innovations. Following the transition, hours worked decline at a constant rate, in line with the trend observed in the data, and growth of the economy declines, which may help account for the recent productivity slowdown. I also analyse the normative properties of the equilibrium.

Chapter 2 studies how technological shifts – such as automation technologies – affect unemployment in general equilibrium. The model economy features endogenous technology choice, incomplete insurance markets and labor market frictions. I show that automation can lead to strong labor market quantities and weak wages. I uncover two novel general equilibrium feedback loops. First, automation shocks raise income inequality – wages of the wealthy increase by more than the wages of the poor – and boost the desire to save: a 'saving for higher wages' effect. Second, with higher unemployment risk, savings rise for precautionary reasons: a 'risk-mitigation effect'. Both mechanisms depress the interest rate and further raise the adoption of capital-intensive technologies.

Chapter 3, co-authored with Larry Summers, concerns the secular trend in long-term real safe interest rates. The main part of the paper uses two general equilibrium models, one capturing life-cycle heterogeneity and another focusing on uninsured idiosyncratic risks and precautionary behaviour, to assess the impact of higher government debt and social security and healthcare spending on the equilibrium. We conclude that government policies have pushed interest rates up by up to 4pp since the 1970s, effectively masking even more subdued private sector interest rates. We also examine the underlying drivers of this weakness using our models. Our results are consistent with the notion that, on its own, the private sector of the developed world may indeed deliver low levels of safe interest rates in the long-run equilibrium.

Contents

List of Tables

List of Figures

Chapter 1

Leisure-enhancing technological change

1.1 Introduction

In models of economic growth, technological change is an abstract generalization of a large and diverse set of innovations undertaken in the real world. What determines the success of these models is whether this generalization captures the essential features of the growth process. In this paper I distinguish between the familiar product- or process-innovations and inventions that are *leisure-enhancing*. I show that making this distinction helps us understand the salient patterns observed in the real world.

The defining difference between the two kinds of technologies is the way they are monetized. Improving a production process or introducing a new or higher-quality product tends to raise the profits of the innovator directly. Instead, leisure-enhancing innovations may sell below the cost of production (indeed, may be given away for free), bringing in little or no revenue. Instead, these innovations are used to capture peoples' time and attention, which are valuable and can ultimately be turned into profit. Some firms – such as platforms in the business of marketing, advertising and brand building – need consumers' time and attention in order to produce. Innovation efforts of these firms are thus focused on creating and improving services that complement leisure, in the hope that consumers recognize their attractiveness and spend more time enjoying them. This form of technological change makes leisure increasingly more attractive. The theoretical model developed in this paper traces out the implications of this phenomenon for the long-run equilibrium.

This paper attempts to make sense of three trends in the data. The first is that leisure-enhancing innovations have been around for a while and are now ubiquitous. The recently-constructed long-run data on the value of free content suggests that the leisure sector has been operating for some time (Figure 1.1).[1] TV channels, websites, social media and free newspapers are all products that many people use every day. Businesses involved in supplying these services are hugely profitable and very large. As of March 2018, the top six of the world's

[1] Measured at the cost of production, the value of free services has been relatively stable at about 1% of GDP over much of the past century in the US (see Figure 1.1). While not negligible, this share in GDP is certainly not large. The theory developed in this paper underscores that leisure-enhancing technological change may be of first-order importance even if the share of the sector is negligible in aggregate.

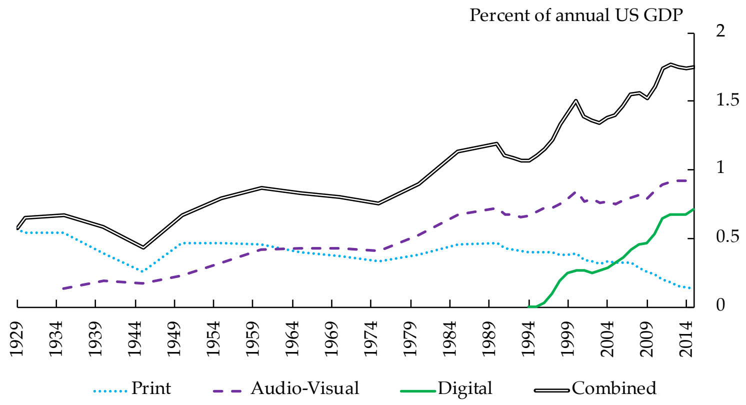

Figure 1.1: Free ad-supported consumer content in the United States

Source: Nakamura, Samuels, and Soloveichik (2017). The figure shows the ratio of free consumer content, measured by the costs of production, to GDP. Thus, for example, it does not capture a welfare measure of the value of Facebook, but only measure the cost of providing it.

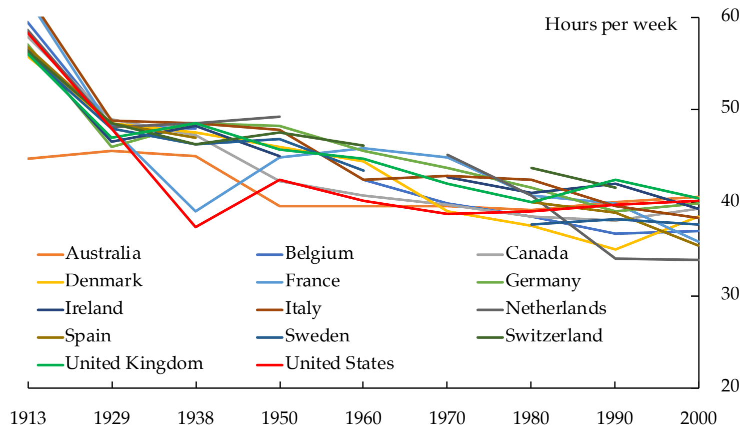

Figure 1.2: Trend growth rates of average hours worked in a range of advanced economies

Source: Huberman and Minns (2007). Note: Data refers to full-time production workers (male and female) in non-agricultural activities.

largest companies by market capitalization – Apple, Alphabet, Microsoft, Amazon, Facebook and Alibaba – are platforms whose business models, at least in part, rely on capturing peoples' time and attention.

The second observation which this work speaks to is the decline in average hours worked and a corresponding rise in leisure hours. Figure 1.2 shows that, following the large falls in average hours worked at the start of the 20th century when the 5-day week became the norm, hours worked continued to decline also in the post-war period, by about 7 hours per week in the latter five decades of the last century.

Furthermore, there is suggestive evidence that changes in time allocation patterns may be linked to technology. Data that tracks how much time people spend on various electronic devices show a very rapid increase in recent years, largely due to widespread smartphone usage (Figure 1.3). This is true across many developed and developing economies.

The third trend that this paper speaks to is the slowdown in productivity growth. Despite

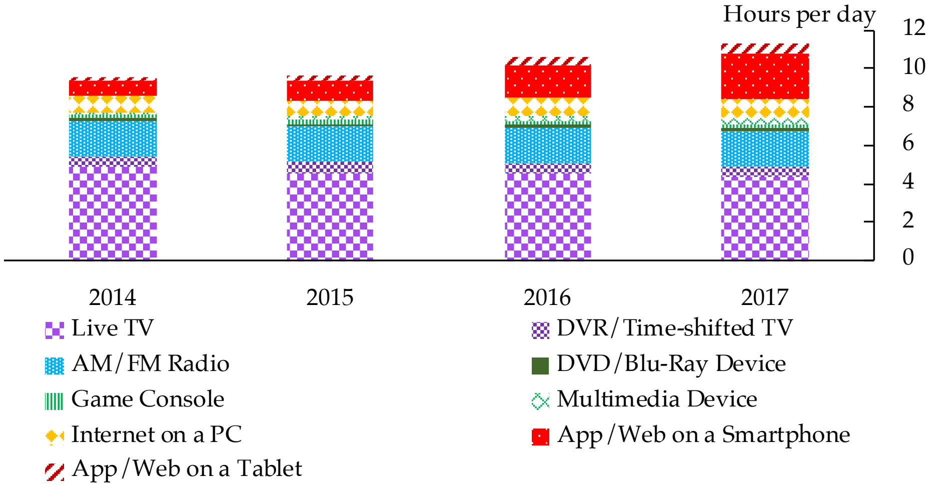

Figure 1.3: Average time spent on media consumption per adult in the US

Source: Nielsen. Note: Figures for representative samples of total US population (whether or not have the technology). Data on TV and internet usage, and the usage of TV-connected devices are based on 248,095 individuals in 2016 and similar sample sizes in previous years. Data on radio are based on a sample of around 400,000 individuals. There are approximately 9,000 smartphone and 1,300 tablet panelists in the U.S. across both iOS and Android smartphone devices.

seemingly rapid technological progress, productivity growth has been surprisingly weak across a wide range of advanced economies in recent decades (Figure 1.4) – what has come to be known as *productivity puzzle*.

To elucidate how these trends are related, I develop a simple general equilibrium model with demand for and supply of leisure services and advertising. I start with a static model that helps to illustrate the novel mechanisms in a simple setting. I then add the dynamic choices and explore the implications of leisure-enhancing technologies for the balanced growth equilibrium.

The static model has the following new features. Household leisure is defined as the utility generated from a range of activities, such as watching TV or browsing the web. To 'generate' leisure, households combine their time with *leisure services* (such as TV channels or web content). To isolate the novel economic forces at play, the theory focuses on the services that are *free* (available at zero price) and *non-rival* (equivalently, once produced, there is zero marginal cost of supplying an extra user).[2] Households value variety: the more leisure activities exist, the more attractive leisure is. On the firm side, households' time is an essential input to the production of advertising, which is supplied to the rest of the economy by a marketing platform.[3] Demand for advertising comes from producers of all other goods and services in the economy:

[2]In reality, some leisure services may sell at positive prices but below the cost of production. These products can be thought of as a combination of the standard consumption good and leisure-enhancing products. The underlying economic forces are similar; and the stark split in the theory developed here underscores the novel aspects of incentives driving innovation.

[3]I assume there is only one platform. One way to think about this assumption is as approximating a sector with a high degree of concentration, perhaps due to fixed costs, economies of scale and network effects in production. To the first order, this should be a reasonably good approximation: in particular, Cournot competition with one dominant firm delivers behaviour arbitrarily close to monopoly solution. Still, future research may consider departing from this assumption by allowing for competition within the marketing sector, which may provide an extra impetus towards leisure-enhancing technological progress.

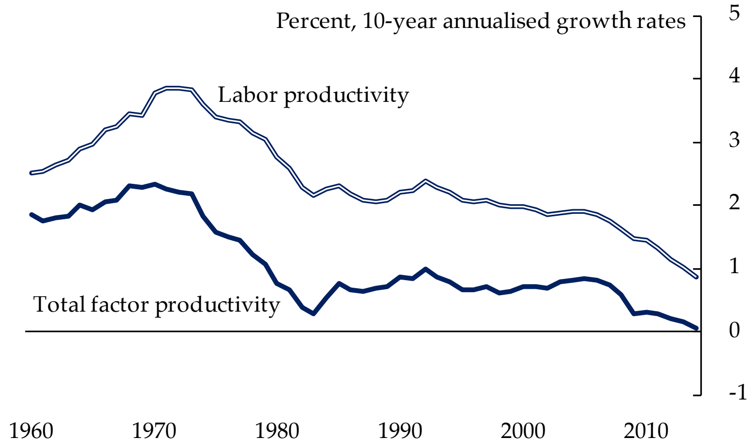

Figure 1.4: Trend growth rates of total factor productivity and labor productivity in advanced economies

Source: Adler et al. (2017). Based on data from Penn World Table 9.0; IMF World Economic Outlook. The figure shows purchasing power parity GDP-weighted average for the largest 20 advanced economies.

each firm can shift its demand curve – sell more at any given price – by advertising its product.

There are two main results that come out of the static setting. First, I show that in the static equilibrium there is positive supply of advertising and of leisure services as long as the economy is large enough. This is driven by the new kind of non-homotheticity: unless leisure is attractive enough, households choose zero leisure. In that sense, the model develops a microfoundation for when and why the leisure sector becomes viable. Second, the model highlights how GDP (as measured by statistical agencies worldwide) misses the value of leisure services when they do emerge. The size of this mismeasurement can be calculated and analyzed in equilibrium.

Embedding the model mechanisms in the dynamic setting brings out three further insights.

First, leisure-enhancing technologies emerge endogenously in equilibrium once the economy is sufficiently developed – a straightforward corollary of the result from the static setup that the economy must be large enough to feature leisure services in equilibrium. Thus the long-run equilibrium is characterized by what may be called a *segmented balanced growth path* (sBGP).

Second, in presence of leisure-enhancing innovations, hours worked decline and leisure hours rise along the sBGP. From a theoretical standpoint, the model provides a novel way to embed falling hours worked in the long-run equilibrium. Empirically, this result matches the trend in time use observed across countries (Aguiar and Hurst, 2007b).

Third, TFP growth declines with the introduction of leisure-enhancing technologies. This is because more leisure translates into a slower growth rate of the pool of resources that are devoted to generating new ideas and knowledge. At the same time, traditionally-measured GDP struggles to capture the value of leisure services, potentially exaggerating the extent of the slowdown.[4]

[4]The present model abstracts from all other forms of leisure and focuses solely on leisure activities that are marketable. In reality, some individuals may also enjoy leisure time that is non-marketable, such as a walk in the park. I certainly hope that is the case. For the purposes of the analysis, in all that follows, leisure is synonymous with marketable leisure.

The final set of results concerns the efficiency of the market equilibrium. The market equilibrium allocation involves two new inefficiencies which I uncover by studying the social planner problem in the static and dynamic settings. In the static equilibrium there is undersupply of leisure technology, as the production of leisure services is governed solely by the prospective profitability of the marketing platform rather than on their true social value. In the dynamic model, there may be oversupply, precisely because of the adverse implications for the long-run growth rate.

Related literature. This study is related to several strands of literature. First and foremost, it builds on and contributes to the extensive literature on endogenous growth (Romer (1990), Aghion and Howitt (1992a), Jones (1995), Kortum (1997), Segerstrom (1998), Acemoglu and Guerrieri (2008), Ngai and Pissarides (2008), Aghion and Howitt (2009), Aghion, Akcigit, and Howitt (2014)). More specifically, it focuses on the ability of endogenous growth models to generate balanced growth equilibrium in which hours worked decline over time, and in this sense it is closely related to the work of Boppart and Krusell (2016). The focus and the underlying economics differ substantially across the two papers, however: while Boppart and Krusell (2016) depart from the balanced growth preferences benchmark to develop the formulation consistent with declining hours and balanced growth (with the income effect stronger than the substitution effect), the theory in the present paper builds on balanced growth preferences and generates these outcomes endogenously in equilibrium.

The present paper draws on and contributes to the literature on the economics of time allocation. Starting from the seminal work of Becker (1965), economists have developed the theory of time allocation that is helpful in the analysis of the complementarities between consumption and leisure (recently summarized in Aguiar and Hurst (2016)). This paper is a first step in the direction of extending this literature to capture a salient features of the modern economy, namely, the rising supply of free and non-rival leisure-enhancing services.

The analysis of empirical evidence on time allocation patterns in recent years presented in the Appendix contributes to the literature that documents the changes in how people spend their time (with seminal contributions from Aguiar and Hurst (2007a) and Ramey and Francis (2009)). The focus on free leisure-enhancing services brings this paper close to the studies by Terranova (2012), Wallsten (2013) and Boik, Greenstein, and Prince (2016), all of whom consider the time and attention allocation problem in the context of the rise of the internet and digital technologies. Unlike these papers, the present study analyses this problem in a fully specified macroeconomic model, which allows for the study of both positive and normative macro implications. The focus on macroeconomic implications of leisure technologies brings the paper close to Aguiar, Bils, Charles, and Hurst (2017) who investigate how video games altered the supply of labor of young men in the United States. I cast the net more broadly, however, studying all leisure-enhancing technologies and highlighting the importance of the endogenous feedback mechanism between the innovations carried out by producers of these technologies and the resulting changes in the allocation of time and attention by consumers.

The present paper draws out implications for the growth rates of productivity and for measurement, contributing to a large and fast-growing literature on the productivity puzzle and on measurement of the modern economy (Aghion, Jones, and Jones (2017)). In particular, the theory proposes a way to reconcile the two sides of the "productivity debate" by relating the rapid leisure-enhancing technological change to the lower rates of TFP growth.[5] The model developed here suggests that in some sense both sides of this debate are correct: rapid technological change in the leisure sector could well cause the growth rates of TFP to decline; thus the perception of rapid progress may well coincide with low observed TFP growth (Figure 1.4). Furthermore, the model allows for a clean analysis of different measures of activity, depending on whether and how one values the leisure services. It contributes to the literature on measurement challenges, particularly those that arise when goods and services are provided at zero prices (Bean (2016), Bridgman (2016), Syverson (2016), Brynjolfsson, Eggers, and Gannamaneni (2017), Byrne, Fernald, and Reinsdorf (2016), Nakamura, Samuels, and Soloveichik (2017), among others).

Finally, this paper builds on several literatures in industrial organization which study two-sided markets and economics of advertising. Classic references on the economics of platforms are Rochet and Tirole (2003) and Anderson and Renault (2006) who study the equilibrium pricing decisions in two-sided markets. More recently, Tirole (2017) provides a non-technical overview of the major forces at play, and highlights the rising importance of the platform business model in the modern economy. There is an extensive literature on the economics of advertising, going back to Marshall (1890) and Chamberlin (1933), and summarized in the IO Handbook Chapter by Bagwell (2007). More recent research has considered the effectiveness of advertising in the context of modern technologies (DellaVigna and Gentzkow (2010), Lewis and Reiley (2014), Lewis and Rao (2015)). The present paper contributes to these literatures by analyzing platforms in a general equilibrium setting and drawing out the implications within a canonical macroeconomic model of competition and growth.

Roadmap. The remainder of this paper is structured as follows. Section 1.2 lays out the static model and derives the main optimality conditions. Section 1.3 shows how these conditions can be adapted to a dynamic endogenous growth setting, defines and characterizes the balanced growth equilibrium. Section 1.4 discusses the welfare properties of the market equilibrium. Section 3.7 concludes and suggests areas for future work.

[5] "Productivity optimists" build their case with anecdotes of groundbreaking technological innovations, speculating about their capacity to change routines and disrupt markets the way societies operate. They are not too concerned about the meagre productivity growth in the data, viewing them either as a transitory phase before the impending technological revolution (Brynjolfsson and McAfee (2014), Brynjolfsson, Rock, and Syverson (2017)) or as an artifact of the inability of the measurement techniques to keep up with the changing nature of the modern economy (Aghion, Bergeaud, et al., 2017). The pessimists view the data as an early sign of persistent weakness, and speculate about economies' ability to generate long-run growth (Gordon (2016), Bloom, Jones, Van Reenen, and Webb (2017)).

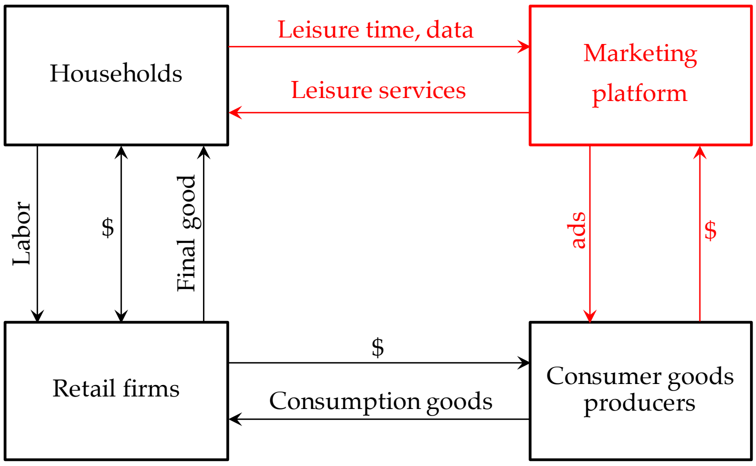

Figure 1.5: The model structure

1.2 The static model

The model builds on the standard framework of monopolistic competition (Dixit and Stiglitz, 1977). Figure 1.5 illustrates the basic structure, with the "leisure sector" highlighted in red. In the reminder of this Section I present a detailed outline of the economy and derive some initial results.

1.2.1 Household's time allocation problem

The economy is populated by two kinds of households: workers and capitalists. Capitalists do not work; they own all firms in the economy and finance hand-to-mouth consumption with the flow of corporate profits.[6]

In a static setting, there is a unit-measure of working households. Households have separable utility over consumption and leisure. Specifically, I work with a canonical form of balanced growth preferences, as in King, Plosser, and Rebelo (1988) and (2002), with log utility from consumption and a constant-elasticity power utility of leisure (see also Boppart and Krusell (2016)). I allow for the possibility that there are network effects in leisure, external to individual choices, which is a plausible feature of some of the modern leisure services such as social media. The instantaneous utility function is:

$$u(c, l) = log(c) + \frac{l^{1-\xi}}{1-\xi} \cdot f(\bar{H}_L) \tag{1.1}$$

where l is leisure and $f(\bar{H}_L)$ is an increasing function that takes as argument the average hours spent on leisure, denoted by $\bar{H}_L$. I further assume that $0 \leqslant \xi < 1$, which implies that utility is weakly concave in leisure ($\xi \geqslant 0$) but that the curvature is not too extreme ($\xi < 1$). This ensures that leisure-enhancing innovations raise leisure time.[7]

[6] This removes the profit flow from the working households budget constraint which allows for a closed-form solution to the model. I denote aggregate consumption of capitalists C.

[7] As highlighted in Section 2, the evidence suggests that leisure-enhancing technologies indeed raises the total time spent on leisure.

In a workhorse macro model of labor supply, leisure and leisure time are one and the same thing. Here I instead view leisure as utility that comes from spending time on a range of leisure activities of measure M:

$$l = \left(\int_0^M \underbrace{\left[\min\{h(j), m(j)\}\right]}_{activity(j)}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}}. \tag{1.2}$$

In this expression activities are indexed by j, and $h(j)$ is the time spent on activity j and $m(j)$ denotes leisure services required for this activity. Parameter ν is the elasticity of substitution across different activities, and is assumed to be greater than unity: $\nu > 1$.

Several assumptions are embodied in this formulation. First, the assumption that the elasticity ν is greater than 1 implies that different activities are imperfect substitutes: there is *love for variety* in preferences, in the sense that, all else equal, more varieties yield higher leisure utility. Second, to generate leisure, each activity requires both time and leisure services, and only those. Thus the formulation assumes away any potential complementarities with consumption.[8] Finally, there is no substitutability between time and free services: enjoying a favorite TV show or surfing the net necessarily requires free time in fixed proportion. This assumption is a natural in the context of services that are available free of charge. Allowing for some substitutability between leisure services and leisure time would make the problem ill-defined.

The household chooses how to allocate their time across activities, leading to the result summarized in the following Lemma:

Lemma 1 *Given (1.2), individuals' optimal allocation of time across available leisure varieties implies that the leisure production function is linear in total leisure time:*

$$l = H_L M^{\frac{1}{\nu-1}}. \tag{1.3}$$

The simple proof is in the Appendix. The key step is to note that, given that services $m(j)$ are free of charge and are non-rival, household is never constrained by $m(j)$.

Equation (1.3) says that for a given amount of leisure time, leisure utility increases in M. This love for variety effect vanishes as $\nu \to \infty$: if different varieties of leisure are perfectly substitutable, only the total time spent on leisure matters and the number of activities become irrelevant.

The problem of choosing the optimal labor supply can be written in the following way:

$$\max_{c, H_L} \log(c) + \frac{l^{1-\xi}}{1-\xi} \cdot f(\bar{H}_L)$$

[8] Conceptually, there are several kinds of inputs that may feature in leisure production: (i) time; (ii) non-durable consumption expenditure (e.g. monthly broadband fee); (iii) durable consumption expenditure (e.g. TV or a laptop); (iv) free, non-excludable and non-rival services (e.g. radio station, web content, free apps). To simplify the analysis and to focus on the key and novel aspects of this problem, in what follows I maintain only the first and the last of these four groups of inputs. Extending the framework so that it encompasses (ii) and (iii) should not alter the key conclusions of this paper.

subject to:

$$l = H_L M^{\frac{1}{\nu-1}},\ c = w \cdot H_W,\ H_W + H_L = 1,\ H_W \geqslant 0,\ H_L \geqslant 0,$$

where H_W denotes time spent working, and time endowment is normalized to 1.[9]

Focusing on the interior solution, the first order condition then implies:

$$\frac{1}{H_W} = H_L^{-\xi} M^{\frac{1-\xi}{\nu-1}} f(\bar{H}_L). \tag{1.4}$$

Because the time endowment constraint implies that $H_W = 1 - H_L$, and in equilibrium $H_L = \bar{H}_L$, equation (1.4) pins down the optimal choice of leisure hours H_L^* , which is an increasing function of M. It is then straightforward to find a solution for all other variables of interest, such as labor supply H_W^* and consumption c^*.

Implications for balanced growth. This characterization of the static equilibrium carries implications for the dynamic model I study in Section (1.3). The key one is that, without further restrictions, optimality condition (1.4) is *not* consistent with balanced growth in a dynamic setting if M changes over time. To see this, suppose that M grows at some constant net rate γ_M. It is easy to see that the optimal labor supply changes with M. For balanced growth we usually require that hours worked are constant; but in principle there exists another possibility, which is that hours worked decline at a constant rate: $\gamma_H \leqslant 0$.[10] On such a path, hours decline asymptotically towards zero, and leisure hours increase at a *non-constant* rate.[11] This is inconsistent with condition (1.4), however, because this condition implies that under the hypothesis of constant growth rates of M and H_W, the growth rate of leisure hours is constant too.

Additional assumptions on the functional form of the utility function need to be imposed to ensure balanced growth. These requirements are contained in the following Proposition:

Proposition 2 *Consider a dynamic environment in which household preferences are described by the instantaneous utility function in (1.1), in which the number of leisure varieties M grows at a constant rate in equilibrium. A balanced growth equilibrium with hours worked decreasing at a constant rate is consistent with these preferences under either of the two conditions:*

Condition 1. $\xi = 0$; that is, utility is linear in leisure, and $f(\bar{H}_L) = 1$, that is, there is no network effect; or

Condition 2. The network effect takes the form of a power function with exponent equal to ξ: $f(\bar{H}_L) = \bar{H}_L^{\xi}$.

[9]Clearly, this formulation focuses only on marketable leisure, abstracting from other uses of time, such as home production, non-marketable leisure, or sleeping.

[10]Note that this is the only other possibility: for example, leisure hours rising at a constant rate would imply zero hours worked at finite horizon, which would yield zero consumption, an outcome clearly inconsistent with utility maximization.

[11]To see this mathematically, note that if $x(t) + y(t) = constant$, taking logs of both sides and differentiating the time constraint with respect to time yields: $\frac{1}{x+y}\left(\frac{dx(t)}{dt} + \frac{dy(t)}{dt}\right) = 0$. The expression on the left hand side can be rewritten as $\frac{x}{x+y}\frac{dx(t)}{dt}\frac{1}{x} + \frac{y}{x+y}\frac{dy(t)}{dt}\frac{1}{y} = \frac{x}{x+y}\gamma_x + \frac{y}{x+y}\gamma_y$ where γ denotes net growth rate of a variable, which itself is, in general, a function of time. Setting this equal to zero, the two growth rates satisfy $\gamma_y = -\frac{x}{y}\gamma_x$. This proves that if γ_x is constant, γ_y is not.

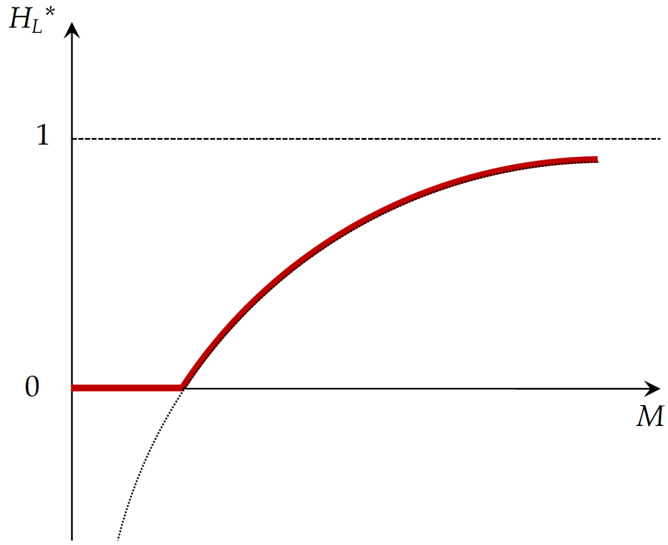

Figure 1.6: Optimal choice of leisure hours as a function of M

Proof. Both of these conditions eliminate the dependence of the right hand side of (1.4) from H_L, implying that $H_W = M^{\frac{1-\xi}{\nu-1}}$. Thus if M grows at a constant rate, hours worked can grow at a constant rate with no further restrictions on the growth rate of leisure hours. ■ In the remainder of the paper I assume that Condition 1 of Proposition 2 holds.[12] Under these assumptions, the optimal choice of leisure hours is given by:

$$H_L^* = \max\{0, 1 - M^{\frac{1}{1-\nu}}\}. \tag{1.5}$$

The key result is that, for M sufficiently large, optimal leisure hours depend positively on the measure of available leisure options. The generic shape of the leisure supply function is depicted in Figure (1.6). Because households like variety, more abundant leisure options make leisure more attractive relative to work, and households optimally choose to work less and play more. There is, however, the possibility of the corner solution: for small M, the optimal choice of leisure hours is zero.

1.2.2 Demand for advertising

The competitive retail sector produces final output by combining labour with a continuum of consumer products $x(i)$ indexed from 0 to A. The key is that the share of each product in the mix is affected by the relative advertising spending of each producer i. In particular, the final good is produced according to the following constant returns to scale aggregate production function:

[12] Thus, to simplify the notation and the algebra, in the remainder of the paper I assume that $\xi = 0$, so that v is a linear function and marginal utility of leisure is constant and equal to 1. All of the economics of the model carries through under the alternative assumption (2) in Proposition (2) and with $0 < \xi < 1$. In particular, the possibility of corner solution arises irrespectively of the curvature of v. This is because with u logarithmic and so as long as v a power function with ξ strictly smaller than 1, marginal utility of consumption goes to infinity faster than marginal utility of leisure as both consumption and leisure get arbitrarily close to zero.

$$Y = \int_0^A \left(\left(\frac{b(i)}{\bar{B}} \right)^{\chi \cdot \mathcal{O}} x(i) \right)^{\alpha} L_Y^{1-\alpha} di \tag{1.6}$$

where $x(i)$ are the individual consumer varieties, and L_Y is labour employed in the retail sector. The term $\left(\frac{b(i)}{\bar{B}}\right)^{\chi \cdot \mathcal{O}}$ guides how advertising enters the model. Here, $b(i)$ is firm i's ad spending and $\bar{B}$ is the average marketing expenditure across all producers: $\bar{B} \equiv \frac{1}{A}\int_0^A b(i)di$. The fraction $\frac{b(i)}{\bar{B}}$ measures the relative advantage of firm i due to its marketing efforts as compared to its competitors.[13] Parameter $\chi \geqslant 0$ measures firms' perceptions of the sensitivity of demand for their product to advertising efforts: the value of zero corresponds to the situation where individual firms view advertising as useless, and the model collapses to the standard monopolistic competition setting found in the literature. Finally, $\mathcal{O}$ is a binary $\{0, 1\}$ variable which denotes the choice of the marketing platform whether to operate or not (discussed below).

The implication of equation (1.6) is that only by advertising more than its competitors can a firm boost demand for its product: marketing is all about *relative* advantage. This is consistent with wealth of empirical evidence, summarized in Bagwell (2007): while advertising may have significant impact on any individual firm's sales, the effect disappears once the unit of observation is expanded to a sector or the macroeconomy.

The present model is silent on the channels through which advertising operates, and assumes there is no direct utility impact of advertising. The literature highlights three reasons why advertising may work: providing information, persuasion, and complementarity with consumption. In this paper I take it as given that advertising efforts shift relative demand, but do not take a stance on the channels through which this operates. In the same vein, I assume that ads do not affect household's welfare directly: one possible interpretation is that the positive information effects broadly offset the negative utility impact associated with exposure to ads. None of the positive economics of what follows is affected by this assumption.

Differentiated goods are produced by a continuum of monopolistically competitive firms indexed by i. For simplicity, I assume that each firm has access to a technology that converts one unit of final output into one unit of consumer good: $x(i) = X$, where X is the amount of final good used in the production of good i. Each firm faces a demand curve for its product, and each has the option to spend resources $b(i)$ on advertising, a service provided by the marketing firm at a given price p_B. Any individual producer i takes the average level of advertising $\bar{B}$ in the economy as given.

[13]The retail sector anticipates any shifts in relative demand due to firms' marketing efforts and demands more of the varieties that boasted higher advertising relative to competitors. The fact that it is the retail sector that "demands more" of a firm's product as a result of its advertising activity is inconsequential to the analysis. The model would be unchanged under an equivalent formulation where consumers were choosing the products directly, and their relative taste for specific varieties was affected by ad expenditures.

1.2.3 The marketing platform

The marketing platform simultaneously operates along two margins: it supplies advertising to businesses, and produces varieties of leisure services. The latter is necessary to operationalize the former, because peoples' time and attention are required inputs in the production of advertising.

This structure of the business poses interesting problems for optimal pricing and production strategies, a subject of the literature in industrial organization. Classic papers by Rochet and Tirole (2003) and Armstrong (2006) and the recent book by Tirole (2017) conclude that the equilibrium pricing structures depend on the elasticities of demand on both sides of the market as well as on externalities between two groups of customers. If one side of the market benefits substantially more than the other from the interaction, then that side will be charged more, including the possibility that the other side may not pay anything.

In practice, many leisure services are indeed provided free of charge to the consumers (Figures 1.1 and 3.13 provide some indication). This equilibrium pricing configuration appears to be prevalent in the leisure sector, perhaps because the interaction between consumers of leisure services (people) and consumers of advertising (firms) directly benefits the latter. The consumer side ends up being the "loss leader", while the business side is where the revenue is generated. While the full analysis of these pricing strategies is beyond the scope of this paper, armed with the logic of economic forces identified in the literature, I assume that the platform chooses to charge a positive price only to the business side. Thus the problem of the platform is to choose whether or not to operate, and to maximize profits by optimally choosing the price of advertising p_B. This choice will then imply the optimal supply of advertising and of leisure services.

The platform is characterized by two technologies and two constraints. The constraints are the household's leisure supply schedule and the firms' demand for advertising.[14] I make the following assumptions with regards to the technologies:

1. Production function for the number of varieties is linear in final good used by the platform, X_B:
$$M = X_B; \tag{1.7}$$
2. Household leisure time is the only input required to produce marketing, and the production function takes the following functional form if the platform chooses to enter ($\mathcal{O} = 1$):
$$B = \begin{cases} 0 & if \quad H_L = 0 \\ \frac{1}{1-H_L} & if \quad H_L > 0 \end{cases}, \tag{1.8}$$
and $B = 0$ otherwise. Here B is total marketing sector output.

These are strong assumptions and they merit further discussion. Assumption (1) is driven by the desire for parsimony: the linear one-to-one technology is perhaps the simplest possible

[14] The former was derived above, the latter can be easily obtained by solving firm i optimization problem: $p_B = \alpha^{\frac{2}{1-\alpha}} \chi \frac{1}{b(i)} \left(\frac{b(i)}{B}\right)^{\frac{\alpha}{1-\alpha}\chi} L_Y$ (see Appendix for full derivation).

formulation. Note that production function for M reflects the fact that services $m(j)$ are non-rival. Production of an extra variety of leisure - increasing M - requires real effort and resources; but once a specific variety j has been produced, it provides unlimited amount of services $m(j)$ within a period. All M varieties are perishable, in that M is modeled as a flow and not a stock. These simplifying assumptions stand for a more general idea that the marketing firm spends resources on improving leisure options of the consumer. Future work could usefully explore different ways of modeling M and the process of technological improvements of the leisure options.[15]

Formulation (2) consists of a set of assumptions, with varying degrees of importance. The most fundamental one is that H_L is an input into the production for marketing, which makes leisure time a valuable commodity in this economy. To provide advertising, the platform must attract "eyeballs".

Second, it assumes that there are *increasing returns to scale* in the production of advertising content. This is a necessary feature for the model to be consistent with balanced growth, where one of the factors of production is bounded (in this case, $H_L \leqslant 1$). This ensures that the marketing platform is able to produce arbitrarily large amounts of advertising services B as the economy grows indefinitely. One justification for a convex production function is that, at higher levels of H_L, the marketing firm is better able to learn about consumers' tastes and preferences, and thus to provide more targeted advertising. Arguably, such information acquisition appears to play an increasingly important role in the real world, and increasing returns is a simple way to capture this force. This assumption is not without controversy, however. The marginal effectiveness of advertising may decline as the cognitive and memory limits start to bite, suggesting that decreasing returns may be a more plausible alternative. Without convexity, there will be a limit on the size of the marketing sector, and at some point along the growth path advertising prices will rise rapidly as demand outstrips supply.

Furthermore, some of the technical constraints faced by firms in this sector suggest that decreasing returns may be a better assumption. For example, the success of machine learning techniques used for targeting consumers' behavior depends on the accuracy of their predictions. But the prediction accuracy increases with a square-root of the number of observations (Varian (2017)), which is a hard, statistics-driven constraint on the growth potential of leisure platforms.

All in all, the assumption of increasing returns imposed here should be interpreted as one that is necessary for the economy to admit a balanced growth path. The sheer size and profitability of the largest platforms suggests that the assumption of increasing returns is a useful one to describe the past experience. The validity of this assumption going forward is an interesting area for future work.

Equation 1.8 makes it possible to solve the problem in closed form. Combining (1.7) and (1.8) yields a quasi-production function that consolidates the platform technologies:

[15] One possibility is that the leisure time itself is an important part of the production process of leisure varieties – particularly for the digital ones, such as social media. See Arrieta Ibarra et al. (2018) for discussion along those lines.

$$B = X_B^{\frac{1}{\nu-1}}. \tag{1.9}$$

Equation (1.9) reveals that the set of assumptions discussed above implies that the production function of the platform is a simple, smooth and concave neoclassical production function with diminishing marginal productivity in the final-good input. This implication facilitates a closed-form solution to the model.

1.2.4 Equilibrium in the static model

1.2.4.1 Definition of market equilibrium

Definition 3 *The market equilibrium in this economy is a set of aggregate quantities* $\{c, C, Y, B, \bar{B}, M, X, X_B, \Pi_B, H_W, H_L, L_Y\}$*, individual-variety quantities* $\{m(j), h(j), x(i), b(i), \Pi(i)\}_{j,i=0}^{j=M,i=A}$*, and prices* $\{w, p_B, q(i)\}_{i=0}^{A}$ *such that:*

- Working households maximize utility subject to the budget constraint and the time endowment constraint, taking wages w and the number of leisure options M as given;
- The retail sector maximizes profit, taking wages w, prices of the consumer sector goods $\{q(i)\}_{i=1}^{A}$, and advertising spending $\{b(i)\}_{i=1}^{A}$ as given;
- Each producer i chooses quantity $x(i)$, price $q(i)$ and the advertising expenditure $b(i)$ to maximize profits, subject to demand curve for their product and taking the average advertising spending in the economy $\bar{B}$ as given;
- The marketing platform sets the price of advertising services p_B, supplies quantities $\{b(i)\}_{i=1}^{A}$, and produces M varieties of free services to maximize profits, taking as given the average level of ad spending $\bar{B}$.
- Capitalists consume all profit income;
- Goods market clears: $Y = c + C + X + X_B$;
- Labour market clears: $L_Y = H_W$.

1.2.4.2 Characterization of the equilibrium

I solve the model under the assumption that the platform finds it worthwhile to be active: $\mathcal{O} = 1$. I then derive the conditions under which this is indeed the platform's optimal choice.

The retail sector. The maximization problem of the retailer is:

$$\max_{x(i),L_Y} \int_0^A \left(\left(\frac{b(i)}{\bar{B}} \right)^{\chi} x(i) \right)^{\alpha} L_Y^{1-\alpha} di - \int_0^A x(i)q(i)di - wL_Y \tag{1.10}$$

The first order conditions yields the standard labour demand equation and the inverse demand curve for the product of firm i:

$$\alpha\left(\frac{b(i)}{\bar{B}}\right)^{\alpha\chi} x(i)^{\alpha-1}L_Y^{1-\alpha} = q(i) \tag{1.11}$$

where $q(i)$ is the price of variety i of the consumer good. Equation (1.11) shows that higher advertising expenditures of producer i, holding advertising of all other firms constant, shifts the demand curve out (for any given price $q(i)$, firm i can sell a higher quantity of its product).

Consumer sector. Intermediate producers maximize profit, subject to the linear technology (recall that $x(i) = X$) and demand curve in (1.11). Each firm can also buy advertising at a given price p_B. In effect, the problem of each producer is:

$$\max_{x(i),b(i)} \alpha\left(\frac{b(i)}{\bar{B}}\right)^{\alpha\chi} x(i)^{\alpha}L_Y^{1-\alpha} - x(i) - p_B b(i). \tag{1.12}$$

The first order conditions to this problem yield the standard result that, for each producer, price equals mark-up over marginal cost ($q(i) = q = \frac{1}{\alpha}$). They also give the inverse demand curve for marketing:

$$p_B = \alpha^{\frac{2}{1-\alpha}}\chi\frac{1}{b(i)}\left(\frac{b(i)}{\bar{B}}\right)^{\frac{\alpha}{1-\alpha}\chi} L_Y. \tag{1.13}$$

Marketing platform. The platform chooses the quantity of advertising supplied to maximize profits, subject to the demand curve for marketing services (1.13) and to the consolidated production function (1.9). This maximization problem can be written as follows:

$$\underset{b(i)}{max} \int_0^A \alpha^{\frac{2}{1-\alpha}}\chi\left(\frac{b(i)}{\bar{B}}\right)^{\frac{\alpha}{1-\alpha}\chi} L_Y di - \left(\int_0^A b(i)di\right)^{\nu-1}$$

Optimality and the symmetry of the problem imply that, in equilibrium:

$$B = (\kappa A L_Y)^{\frac{1}{\nu-1}} \tag{1.14}$$

where $\kappa = \frac{\alpha}{1-\alpha}\chi^2\alpha^{\frac{2}{1-\alpha}}(\frac{1}{\nu-1})$ is a constant. Next, combining (1.14) and (1.9) gives the expression for the inputs used by the platform:

$$X_B = \kappa A L_Y. \tag{1.15}$$

Finally, note the labour market clearing implies that:

$$L_Y = H_W = X_B^{\frac{1}{1-\nu}} \tag{1.16}$$

where the last equality follows from household optimization under the assumption that marketing sector is active. Combining (1.15) and (1.16) gives the expressions for equilibrium levels of final output used in production of advertising X_B and total advertising output B as function of the exogenous parameters of the model:

$$X_B = (\kappa A)^{\frac{\nu-1}{\nu}}, \tag{1.17}$$

$$B = (\kappa A)^{\frac{1}{\nu}}. \tag{1.18}$$

Unsurprisingly, the equilibrium size of the marketing platform depends positively on the measure of varieties: the more goods varieties there are, the higher the demand for advertising. Marketing firm is also larger for higher χ and lower ν. A higher χ means that individual firms believe advertising is a powerful way through which one may increase profits, and a lower ν means that the link between the leisure-technology and leisure supply is stronger.

This structure implies the following Lemma:

Lemma 4 *The platform has a minimum scale of operation. It is active ($\mathcal{O} = 1$) if and only if $M^* = X_B^* \geqslant 1$, where M^** and *X_B^* are the optimal choices of the monopolist.*

The Lemma follows from the production functions for M and B in (1.7) and (1.8), respectively, and from the household's leisure supply (1.5): since the household allocates no proportion of their time to leisure if $M \leqslant 1$, it follows that the platform can produce positive levels of B only when $M \geqslant 1$.

Given (1.9), it follows that the sufficient conditions for the platform to be active are that: (i) its equilibrium output is greater than one ($B^* > 1$) and (ii) that its profits are non-negative. The latter is simply a restriction on the parameters of the model which I assume is satisfied.[16] The following Proposition characterizes the former:

Proposition 5 *The marketing sector is active only in an economy that is large enough. In particular, the condition for the emergence of the marketing sector is given by:*

$$A > \frac{1}{\kappa}. \tag{1.19}$$

If the economy is small, and the above condition is not met, the marketing sector remains inactive, with $\mathcal{O} = 0$. In that case the structure of equilibrium is identical to a standard model of monopolistic competition. If (1.19) holds, the marketing sector is active and the scale of its operations in equilibrium is given by (1.18).

Proof. The result follows immediately from equation (1.17) and Lemma 4. The result that the economy collapses to a standard model when the marketing firm is inactive follows immediately from setting $\mathcal{O}$ equal to zero in (1.6) and noting that for $M = 0$ households' labour supply is inelastic. ■

Intuitively, when the number of product varieties is small, there is little profit to be made from supplying marketing services. Relative to these limited profit opportunities, the marketing platform finds it too costly to produce enough of the leisure varieties M to get the households to spend any time on leisure.

[16]Marketing platform profits are $\Pi_B = (\alpha^{\frac{2}{1-\alpha}}\chi - \kappa)AL_Y$, which is positive if and only if $\frac{\alpha}{1-\alpha}\chi\frac{1}{\nu-1} < 1$. I assume that this second condition is met by reasonable parameter values. This is not a strong requirement: it holds for all reasonable values of the parameters (in particular, it could only fail if α was extremely close to 1).

1.2.4.3 GDP in the model and in the data

Is GDP as currently measured in the data a useful metric for activity in the modern economy? To begin thinking about this issue, is useful to review how the activity of the marketing sector is measured in the data at the moment. Bean (2016) concisely describes the current practice:

> Most of the web's popular destinations, such as Google, Facebook, and YouTube, rely on advertising to generate income. [...] Digital products and services are effectively paid for by the advertisers. As such, the 2008 UN System of National Accounts (SNA) treats them as an intermediate input in the advertising industry. Therefore, the advertisement expenditure adds to the value added of the industries supplying advertisement space, and at the same time detracts from the value added of the advertising industries. Consequently the value of digital products nuanced through selling advertising space will be accounted for in aggregate GDP only to the extent that it also translates into higher consumption of the goods and services being advertised.

In our model, advertising is a beggar-thy-competitor, zero-sum activity: there is no increase in consumption that results from it in equilibrium. Consequently, it is clear that GDP, the way it is measured in the data, would not capture any of the value generated by the marketing sector.

To be more specific, consider the true measure of GDP in the context of the present model, denoted by GDP_{model}. It is a sum of gross output in each sector: the retail sector (Y), consumer sector ($\int_0^A x(i)di$), advertising firm (B and $\int_0^M m(j)dj$), less the value of the inputs used in these production processes.[17] The key challenge is to measure the value of leisure services. One way to do so is to rely on current wages that reflect the opportunity cost of time (Goolsbee and Klenow (2006), Brynjolfsson and Oh (2012)). This leads to a valuation method where the value of these services is set to equal the time spent on them, evaluated at the current wage rate. In light of the equilibrium conditions above, the value of these services is thus:

$$w\cdot\int_0^M m(i)di = w\cdot\int_0^M h(i)di = wH_L = w\cdot\left(1-\left(\frac{1}{\kappa A}\right)^{\frac{1}{\nu}}\right) = (1-\alpha)\alpha^{\frac{2\alpha}{1-\alpha}}A\left(1-\left(\frac{1}{\kappa A}\right)^{\frac{1}{\nu}}\right).$$

Taking this methodology of valuing leisure services as given, the equilibrium GDP in this economy can be expressed a function of the primitives of the model:

$$GDP_{model} = \underbrace{\kappa^{-\frac{1}{\nu}}A^{\frac{\nu-1}{\nu}}\left[\alpha^{\frac{2}{1-\alpha}}\left(\frac{1}{\alpha^2}-1\right)-\kappa\right]}_{GDP_{measured}} + \underbrace{(1-\alpha)\alpha^{\frac{2\alpha}{1-\alpha}}A\left[1-\left(\frac{1}{\kappa A}\right)^{\frac{1}{\nu}}\right]}_{Value\,of\,leisure\,services}.$$

True GDP is a sum of two components: the measure of activity of the "traditional" sector, which is properly captured by statistics agencies, and the value of leisure services. For any set

[17] Derivations of gross outputs of each sector and the verification of the goods market clearing condition are in the Appendix.

of parameters, the model allows for a precise analysis of the degree of mismeasurement. This is particularly interesting in dynamic model, to which I now turn.[18]

1.3 Dynamic model

There are three key features of the dynamic model relative to the static setup above. First, there is the R&D sector, which employs researchers to invent new varieties of consumption goods. The size of the economy – measured as the number of varieties A – is no longer a parameter; instead, it is endogenously determined by the optimal choices of the firms in the R&D sector. Second, households now solve an intertemporal consumption-saving problem. Third, there is population growth, and agents can be employed in the final goods sector as well as in R&D, which alters the labor market clearing condition.

1.3.1 The R&D sector

Time is continuous. New varieties of differentiated goods are invented by the R&D sector employing researchers. I follow Jones (1995) and propose a general formulation of the ideas production function, with the possibility that there are externalities from the stock of ideas over time (represented by $\phi < 1$), and that there may be diminishing returns to employing researchers (represented by $0 < \lambda \leqslant 1$):

$$\dot{A}(t) = \zeta A(t)^{\phi} L_A(t)^{\lambda}. \tag{1.20}$$

Given this technology, R&D producers choose how much labour to employ in order to maximize profits:

$$\max_{L_A} \dot{A}(t) P_A(t) - w(t) L_A(t)$$

where P_A is the price of a patent which gives the holder of the patent a perpetual right to produce that particular variety. The first order condition of this problem is:

$$P_A(t) = \frac{w(t)}{\zeta A(t)^{\phi} L_A(t)^{\lambda - 1}}. \tag{1.21}$$

A consumer firm is willing to purchase a patent as long as the return from doing so exceeds the required rate of return on investment, r, meaning that the following no arbitrage condition

[18] An alternative way of incorporating the zero services into the measured GDP is to consider the production account, and instead of valuing the services at the opportunity cost of peoples' time, focus on the costs of actually providing these leisure services to individuals. This is the approach taken Nakamura, Samuels, and Soloveichik (2017). Naturally, the production approach suggests the contribution of the free services an order of magnitude smaller than the calculation of the opportunity cost of time – the latter measure is closer to capturing the consumer surplus from those services. This is particularly important in the recent digital era, when the cost of producing content is low and its reach is very widespread. In that sense, the production approach is likely to be a lower bound for the mismeasurement effects. Nakamura, Samuels, and Soloveichik (2017) find that free media and information services amount to around 1.8% of US GDP in level terms, and that the correcting for the rapid growth of the online free services over the past 20 years would raise the US GDP growth rate by around 0.1pp each year, not a negligible amount given the adopted methodology.

holds in equilibrium:

$$\frac{\Pi(t)}{P_A(t)} + \frac{\dot{P_A}(t)}{P_A(t)} = r(t) \tag{1.22}$$

This condition has a standard interpretation: the flow return of purchasing a patent (the dividend plus the capital gain) must equal the rate of interest.[19]

1.3.2 Decentralized equilibrium

To save on notation, I drop time subscripts when defining the decentralized equilibrium.

Definition 6 *In light of the previous equations, a dynamic market equilibrium is a set of paths of: aggregate quantities:* $\{c, C, X, X_B, H_W, H_L, L_A, L_Y, A, M, B\}_{t=0}^{\infty}$*; quantities of individual varieties:* $\{x(i), b(i), m(j), h(j)\}_{i \in A, j \in M; t=0}^{\infty}$*; prices of each intermediate input and of marketing services:* $\{q(i), p_B\}_{i \in A_t, t=0}^{\infty}$*; interest rates and wages:* $\{w, r\}_{t=0}^{\infty}$*; marketing platform activity indicator:* $\{\mathcal{O}\}_{t=0}^{\infty}$*; such that:*

- Working households maximize utility subject to the law of motion for per-capita assets, the leisure production technology, and the no-Ponzi game condition, taking as given all prices and the number of available leisure options;
- Capitalist households consume firms' profits every period;
- Retail sector firms maximize profit taking prices and advertising spending of each variety as given;
- Intermediate goods producers maximize profits subject to the technology of production and demand curve for their product, taking as given the average level of advertising in the economy;
- The marketing platform maximizes profits subject to the marketing production technology, leisure varieties production technology, households' optimal time allocation decisions, and the demand curve for marketing services. It becomes active only if it can be profitable.
- The R&D producers maximize profits subject to the ideas production function and the free entry condition.
- Labour market and goods market clear.

I now characterize the equilibrium allocation.

---

[19]Note that the price of a patent grows at the same rate as population: $\frac{\dot{P_A}}{P_A} = n$. To see this, note that the price of a patent will reflect the stream of profits generated by variety i, discounted by the effective discount rate, which will be the market interest rate r minus growth rate of population n:

$$P_A = \int_t^{\infty} e^{-(r-n)s} \Pi ds = \frac{\Pi}{r-n}$$

Combining this with the no-arbitrage condition above yields the result.

1.3.3 Household intertemporal problem

The standard household problem yields the Euler Equation and the intratemporal optimality condition (the Appendix contains the full derivation):

$$\frac{\dot{c(t)}}{c(t)} = r(t) - \rho; \tag{1.23}$$

$$\frac{w(t)}{c(t)} = M(t)^{\frac{1}{\nu-1}}. \tag{1.24}$$

1.3.4 Condition for the marketing sector to be active

To determine whether or not the marketing platform is active, it is useful at this point to derive a dynamic equivalent of the condition in Proposition 5. The condition differs from the one in the static setting because the labour market clearing condition is different: workers are now employed both in the retail and in the R&D sectors and population is growing over time. Labour employed by the retail sector is now:

$$L_Y(t) = H_W(t)(1 - s(t))N(t) \tag{1.25}$$

where s is the share of the workforce employed in the R&D sector (which will be asymptotically constant in the balanced growth equilibrium). Equation (1.25) is simply an accounting identity which states that total labour supply in the final goods sector is equal to average hours worked times the number of workers employed in the retail sector.

Equations (1.15), (1.16) and (1.25) imply that in the dynamic setting the optimal choice of X_B is given by:

$$X_B(t) = (\kappa(1 - s(t))A(t)N(t))^{\frac{\nu-1}{\nu}}. \tag{1.26}$$

We thus have the following Proposition:

Proposition 7 *In a dynamic economy, the platform is active and there is leisure-enhancing technological change if and only if*

$$A(t) \cdot (1 - s(t)) \cdot N(t) > \frac{1}{\kappa}. \tag{1.27}$$

Proof. The Proposition follows from combining equation (1.26) with the fact that, for the marketing sector to be active, we must have $X_B > 1$. ■

Proposition 3 describes a watershed moment for this economy. On the balanced growth path, the left hand side of condition (1.27) is a growing quantity, while the right side is a constant. Along a BGP, it is just a matter of time when (1.27) is satisfied. When that happens, the marketing sector emerges endogenously. The economic intuition is similar to that in the static model: the marketing platform becomes a viable business proposition only when the economy reaches a certain size. Only then the demand for advertising and the revenue generated through the business-to-business side is sufficient to compensate for the costs of creating the required supply of leisure services.

1.3.5 Segmented Balanced Growth Path

Given the result in Proposition 7, the balanced growth equilibrium in this economy is defined as follows:

Definition 8 *A* segmented balanced growth path *(sBGP) is an equilibrium trajectory along which: (i) when (1.27) is not satisfied, per capita consumption, output and the measure of varieties A all grow at a constant rate; (ii) as $t \to \infty$, per capita consumption, output, A and M grow at possibly distinct but constant rates, and average hours H_W decrease at a constant rate.*

The segmented Balanced Growth Path is then characterized by the following Proposition:

Proposition 9 ***Characterization of the sBGP.***

There exists $\hat{t}$ such that (1.27) holds $\forall t \geqslant \hat{t}$.

For $t \leqslant \hat{t}$, average hours worked are constant $(H_W = 1)$, and per capita consumption, output, wages and knowledge all grow at the same constant rate given by:

$$g_c = g_y = g_w = g_A = \frac{\lambda n}{1-\phi}. \tag{1.28}$$

For $t \geqslant \hat{t}$, the marketing platform is active and the economy transitions to the other segment of the sBGP. As $t \to \infty$, average hours worked decline at a constant rate $\gamma_H = -\frac{1}{\nu}(\gamma + n)$ and the growth rate of A is lower than g_A. Specifically, it is given by:

$$\gamma = \frac{n\lambda\left(\frac{\nu-1}{\nu}\right)}{1-\phi+\left(\frac{\lambda}{\nu}\right)}. \tag{1.29}$$

Proof. The first part of the Proposition follows from the fact that the left hand side of (1.27) is a growing quantity (as there is positive population growth) while the right hand side is a constant.

Consider the section of the BGP where the condition (1.19) does not hold. This implies that the platform is inactive: $\mathcal{O} = 0$. It is then straightforward to see that the economy is identical to the model economy studied in Jones (1995). In that economy, the measure of consumer goods varieties A, wages, output and consumption per capita all grow at the same constant rate given by (1.28). The proof is available in Jones (1995) and is omitted for brevity.

To prove the final part of the proposition, note first that the solution to the household's intratemporal optimization problem implies that average hours worked are a power function of the number of leisure varieties (recall equation (1.5)). Thus the growth rate of average hours must be proportional to the growth rate of varieties of leisure if the two variables grow at constant rates, so that:[20]

$$\gamma_H = \frac{1}{1-\nu}\gamma_M \tag{1.30}$$

[20]Note that, for variables X and Y, if $Y = X^a$ then gross growth rates satisfy $\zeta_Y = \zeta_X^a$ and net growth rates satisfy $\gamma_Y = a\gamma_X$.

Next, equation (1.7) implies that the M and X_B grow at the same rate in equilibrium. Equation (1.26) implies that this rate is given by:

$$\gamma_{XB} = \gamma_M = \frac{\nu - 1}{\nu}(\gamma + n) \tag{1.31}$$

where γ is the net growth rate of A. Combining (1.31) and (1.32) gives the expression for the growth rate of average hours worked as a function of γ:

$$\gamma_H = -\frac{1}{\nu}(\gamma + n) \tag{1.32}$$

Next, note that γ is, by definition, given by:

$$\gamma \equiv \frac{\dot{A}(t)}{A(t)} = \frac{\zeta L_A(t)^\lambda}{A(t)^{1-\phi}} \tag{1.33}$$

On the segmented balanced growth path, the left-hand side of this equation is asymptotically constant by definition. That means that the numerator and the denominator of the fraction on the right-hand side must grow at the same rate. More formally, using the fact that $L_A = H_W \cdot N \cdot s$, differentiating (1.33) with respect to time gives and solving for γ yields:

$$\gamma = \frac{\lambda(\gamma_H + n)}{1 - \phi} \tag{1.34}$$

Together, the two relationships, (1.32) and (1.34), are two equations in two endogenous variables that readily give the expression for γ in equilibrium. The fact that $\gamma < g_A$ follows immediately from the comparison of the two expressions.[21,22] ■

Discussion. Proposition 9 contains several key results which I now discuss.

First, it tells us that the economy transitions endogenously from one segment of the BGP to another, when condition (1.27) is first satisfied. Fundamentally, this is driven by the juxtaposition of the household leisure supply, which is zero for low values of M, and the fact that the platform requires some of household's time to produce. The theory highlights the importance of market size (and the associated strength of demand for advertising) as the key force behind the emergence of the marketing sector.

To be sure, in reality there are other reasons, absent from the parsimonious model presented above, for why leisure-enhancing technological change emerges in equilibrium. For example, the invention of the washing machine and other appliances dramatically cut the time required to complete the necessary household duties, which meant that households had more time to allocate to other activities, including leisure. This meant that the potential market for leisure services expanded, which could have directed innovations towards the leisure sector, as in the

[21] The associated rate of decline of average hours worked is given by $\gamma_H = -\frac{1}{\nu}\left[\frac{\lambda n\left(\frac{\nu-1}{\nu}\right)}{1-\phi+\left(\frac{\lambda}{\nu}\right)} + n\right]$ which goes to zero as $\nu \to \infty$.

[22] Finally, it remains to check that utility is bounded on the sBGP. This is indeed the case for a discount rate that is large enough.

directed technical change literature (Acemoglu (2002)). More recently, other technological breakthroughs - notably the invention of the internet and the subsequent hardware advances - opened up new possibilities to reach people's consciousness throughout their day (recall Figure 3.15), amass large quantities of data, and target advertising with extraordinary precision. The theory presented here, while falling short of modeling these forces explicitly, is not contesting their importance. Instead, it is complementing these well-known explanations with a new one that might have previously been overlooked.

Second, the balanced growth equilibrium features declining hours worked. This is a unique feature of the theory developed here: the model matches (at least qualitatively) the key trend observed in the data without the need to hardwire the decline in hours into household preferences, and shows that despite this trend, growth can be balanced, in the sense that all variables of interest grow at constant rates. In a standard model with endogenous labor supply, falling hours worked on the balanced growth path can be ruled as violating household intratemporal condition in finite time. Here, instead, this very intratemporal choice is the key force driving the downward trend. Furthermore, the model stresses the importance of innovation that is targeted specifically at capturing people's time, a mechanism that is arguably at the heart of the modern economy but is new to macroeconomic models.

The third result is that TFP growth slows down as the leisure sector emerges. What is the intuition behind the decline in TFP and does it make sense? In the class of semi-endogenous growth models (to which the present model belongs), such as Jones (1995), Kortum (1997) and Segerstrom (1998) models, the long-run growth rate is determined by the growth rate of the total pool of resources that can be used to generate ideas. Bloom, Jones, Reenen, and Webb (2017) highlight that this is the relevant framework to think about growth, because of the key observation that research productivity (defined as $\frac{\frac{\dot{A}}{A}}{L_A}$) is falling across the board, be it in the aggregate economy or at the industry- and firm-level. Exponential growth, if and when it is observed, is a result of a combination of declining research productivity and increasing pool of resources devoted to research. Leisure-enhancing technological change acts to limit the growth of this pool (by generating decline in hours worked). This lowers the resulting rate of economic growth. This abstract reasoning may correspond to some of the phenomena discussed in the real world. For example, technologies and content available today may act to divert peoples' time and attention away from creative thinking. Some suggest that more distracted minds may lower workers' productivity (e.g. Terranova (2012), Nixon (2017)), while the nascent experimental evidence shows that leisure-technologies occupy limited cognitive resources, significantly decreasing cognitive performance (Ward, Duke, Gneezy, and Bos, 2017). The present paper can be seen as formalization of these ideas, and a way to explore the linkages between those micro-observations and macroeconomic performance.

Several other implications follow from Proposition 9. Growth of consumption per capita declines by more than the growth of TFP. On the sBGP, it must be that consumption grows only as fast as earnings, so that: $\gamma_w + \gamma_H = \gamma_c$. Because labor is paid its marginal product in the retail sector, wages satisfy $w = (1-\alpha)\frac{Y}{L_Y} = \alpha^{\frac{2\alpha}{1-\alpha}} A$ and grow at the same rate as TFP: $\gamma_w = \gamma$.

Given the rate of decline of average hours worked in Proposition (9), the growth of per-capita consumption is given by:

$$\gamma_c = \frac{\lambda n \left(\frac{\nu-1}{\nu}\right)^2}{1-\phi+\left(\frac{\lambda}{\nu}\right)} - \frac{n}{\nu}$$

which is lower than in the economy with no leisure-enhancing technological change: not only is each worker working fewer hours, but they are also less productive.[23]

The Euler Equation on the sBGP gives the interest rate as $t \to \infty$:

$$r = \rho + \gamma_c. \tag{1.35}$$

Equation (1.35) indicates that the decline in the growth rate of consumption growth is associated with the decline in equilibrium real interest rate - a prediction that is consistent with the trend observed in the data (Rachel and Smith, 2015) and a subject of the debate over the secular stagnation hypothesis (Summers, 2015a). Intuitively, facing lower steady state growth rate of income, households ought to revise down their consumption growth. For that to happen, the interest rate must fall to make the current consumption more attractive relative to future consumption.

The Appendix contains the derivation of the equilibrium share of the labour force employed in the R&D sector, s. This share is determined by the relative profitability of R&D, which in turn is tightly linked to the profitability of the consumer sector (recall that the price of a patent is a discounted sum of consumer firm's profits). The share is thus lower when marketing is active, as a chunk of would-be profits of the consumer sector is captured by the marketing firm.[24]

1.3.6 An illustrative parametrization

To further explore the workings of the model and the quantitative properties of the balanced growth equilibrium and the transition path, it is useful to consider an illustrative parametrization (Table 1). The parametrization corresponds to annual frequency, with the discount rate of 2% and population growth of 1% per annum. I follow Jones (1995) and set parameters of the R&D production function – λ and ϕ – to hit the long-run per-capita growth rate of the economy (with inactive advertising) of 2%.

Parameter ν is the key one in this calibration. It pins down the strength of the link between

[23] It is possible that in balanced growth equilibrium consumption per capita declines over time. This curious result highlights the fact that, when households choose how much leisure to enjoy, they do not take into account the impact of that decision on the pace of technological progress in equilibrium. The condition for positive consumption growth in the presence of leisure enhancing technology is $\nu > 2 + \frac{1-\phi}{\lambda}$, that is, the strength of the love-for-variety effect needs not be too strong. In the calibration below, the right-side of this condition is equal to around 2.5, so that with $\nu > 2.5$, there will be positive asymptotic growth in consumption per capita along the sBGP.

[24] There is also a second effect which works through the lower steady state growth rate: with lower growth, future profits are less, again making R&D less profitable. The causal effect runs solely from the growth rate to the proportion of workers in R&D and not vice-versa: the share of workers in R&D only lowers the *level* of output, but not its asymptotic sBGP growth rate.

leisure enhancing technological change and the resulting shifts in time allocation. A higher ν makes this link weaker: advancements in leisure technology do not shift households' decisions by much. Conversely, the closer ν is to unity, the more sensitive households' time allocation decisions become. Calibrating this parameter is difficult, and future work may fruitfully consider the strength of the link between technology and time allocation empirically. Through the lens of the theory developed in this paper, parameter ν is an elasticity of substitution across the leisure varieties,[25] so one way to gauge a sense of the plausible magnitudes is to consider the estimates of different elasticities from the literature. For example, Goolsbee and Klenow (2006) estimate the elasticity between internet vs. everything else of about 1.5, highlighting that such elasticies can be quite low when the composite goods are very broadly defined. Elsewhere, Broda and Weinstein (2006) study the welfare gains from increased variety as a result of the rising trade penetration in the US economy. In the process, the authors estimate thousands of elasticities of substitution between similar products imported from different countries. For example, they establish that the elasticity of substitution across cars imported from various places is around 3, while the corresponding elasticity for apparel and textiles is about 6. Within products classified as differentiated, the median elasticity is around 2 and the mean is about 5. To err on the side of caution, the parametrization that follows assumes that $\nu = 5$.

Parameter χ corresponds to the perceived effectiveness of advertising: the value of 0.05 means that each producer believes that spending twice as much on advertising as their competitors would increase their sales by 2.5%.[26] The empirical literature on the effects of advertising highlights that estimation is plagued with difficulties, not least the endogeneity of the advertising campaigns and the residual volatility that is very large compared to any potential treatment effect (DellaVigna and Gentzkow (2010), Lewis and Reiley (2014), Lewis and Rao (2015)). Some of these studies report the increases in sales that follow a large advertising campaign of about 5%, putting the value of 0.05 broadly in the right ballpark.

Finally, α is a share of consumer goods in the final good production, with $1-\alpha$ corresponding to the share of labour in the retails sector. I set α to $\frac{1}{3}$, following Jones (1995).

Under this parametrization, Figure 1.7 compares the constant growth rates the two sections of the segmented balanced growth path: the initial section (for $t < \hat{t}$) and then the asymptotic one (as $t \to \infty$). In the initial section of the BGP the marketing sector is inactive and the growth rate of GDP, given by equation (1.28), is equal to 3%, with growth of A at 2% and population growth of 1% per year.[27] The second set of bars in each panel in the Figure illustrates the

[25] It is useful to draw parallels with the growth literature of the early 1990s. The present paper, highlighting the "love for variety" effect, incorporates the horizontal innovation mechanism popularized by Romer (1990). It is possible, however, to recast the model in the spirit of Grossman and Helpman (1991) and Aghion and Howitt (1992b), highlighting the quality improvements brought about by innovation. Indeed, it is possible to mix the two approaches in a more general model. The basic economics and the implications of the leisure-enhancing technologies will likely remain unchanged, as they depend on M growing over time and do not rely on any specific interpretation of what is behind these increases. Nonetheless, these alternative modeling schemes might perhaps highlight some new channels and interactions. Consequently, this is a useful avenue that may be pursued in future research.

[26] To see this, note that equilibrium quantity is $x(i) = [\alpha^2 \left(\frac{b(i)}{B}\right)^{\alpha\chi} L_Y^{1-\alpha}]^{\frac{1}{1-\alpha}}$, thus the elasticity is $\frac{\alpha\chi}{1-\alpha}$, which is 0.025 for $\alpha = \frac{1}{3}$.

[27] Note that the variables that describe the behaviour of the marketing sector - such as M, B or Π_B - are equal

Table 1.1: Illustrative parametrisation

ρ	Household discount rate	0.02	$r \approx 4\%$
n	Population growth	0.01	AEs data
α	Share of consumer goods in Y	0.33	Jones (1995)
λ	Diminishing returns to labor in R&D	0.65	Jones (1995)
ϕ	Increasing returns in ideas production	0.65	Jones (1995)
ν	Elasticity of substitution across leisure varieties	5	see text
χ	Perceived effectiveness of advertising	0.05	see text

asymptotic growth rates following the transition. The growth of TFP is about 1pp lower in this new section of the BGP, illustrating the magnitudes of the effects described in Proposition 9. This result is quite remarkable: given the cautious calibration, the effects are economically very large. The model easily explains 1pp drop in total factor productivity growth once the economy adjusts, which has tremendous implications for economic policy across the board. The growth rate of consumption per capita predicted by the model is lower still, reflecting both the decline in productivity growth and the downward trend in hours worked.

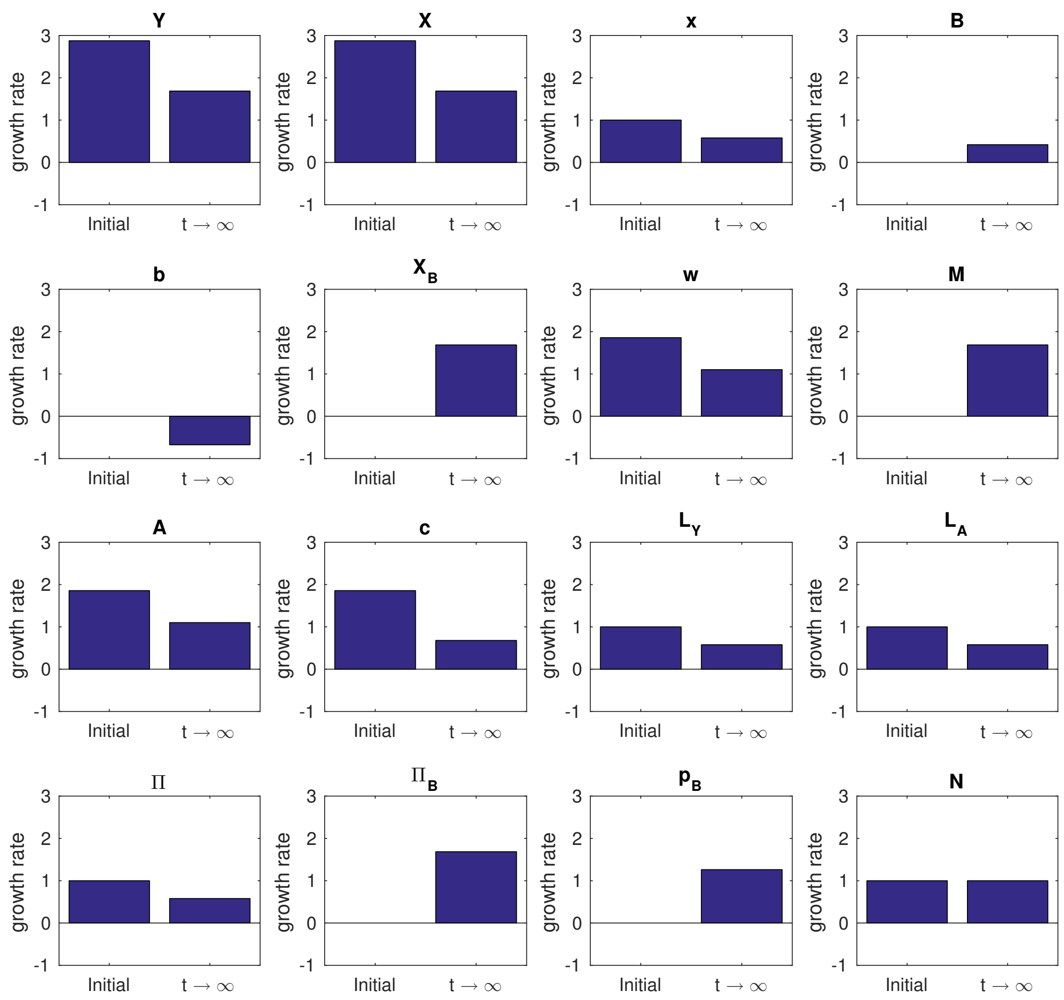

Figure 1.7: The two segments of the balanced growth path: annual growth rates before and after leisure-enhancing technological progress

to zero on the initial section of the sBGP, and their growth rates are not well defined.

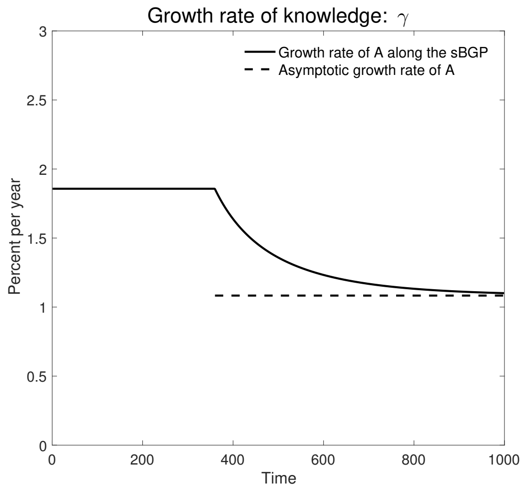

Figure 1.8: Growth rate of knowledge in the sBGP

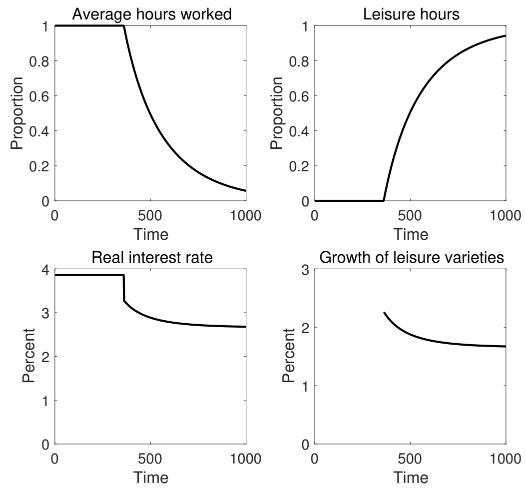

Figure 1.9: Some key variables on the sBGP

The economy does not jump from one section of the BGP to another; instead, there are transitional dynamics (Figures 1.8 and 1.9).[28,29] Figure 1.8 shows that the growth rate of knowledge declines gradually.[30] Remarkably, the decline in TFP occurs exactly at the point when the rate of technological progress is booming as a result of the emergence of the leisure sector (the bottom-right panel of Figure 1.9 shows the growth rate of leisure technology). This result may help to reconcile the productivity paradox, that the weak observed productivity growth occurs at the time when technological change appears to be very rapid.

The two upper panels of Figure 1.9 show the time allocation choice along the sBGP. Hours worked decline at a constant rate of about -0.4% per annum, which matches the average rate of decline of weekly hours worked illustrated in Figure 1.2. Thus the model's key prediction is in line with the trends observed in the data.

The bottom-left panel of 1.9 shows that the equilibrium real interest rate falls sharply on impact, and then adjusts further to settle at a level that is about 1.3pp lower. This economically significant decline in the real rate of interest is matches the decline in real interest rates across advanced economies that we have seen over the past 20 years (Rachel and Smith, 2015), arguably the period over which the leisure enhancing technological change has been particularly rapid.

Of course, the parsimonious, qualitative approach developed here may be too simplistic to give definitive answers on the question of how much of the observed slowdown in productivity

[28]To see this, consider the fact that the growth rate γ depends on the stock of ideas A, which the model predicts the economy will adjust gradually to the new steady state (see equation 1.33). This is true not just for γ, but also for all other endogenous variables - not least the share of workers in R&D and average hours worked.

[29]The logic of the numerical algorithm used to compute the transitional dynamics is the following: for an initial guess for the time path of the number of people employed in R&D, $\{L_A^0\}_{t=0}^T$, compute the growth rate of ideas γ^0 and a host of other variables consistent with that guess. Then recompute L_A^1 consistent with this guess equilibrium, and compare with the initial guess. Iterate on this procedure until L_A^1 and L_A^0 are numerically close enough. This procedure converges robustly.

[30]The transitional dynamics are fairly slow - this is driven by the (relatively high) value of ϕ, and is a well-known result in this class of models, discussed at length in Jones (1995).

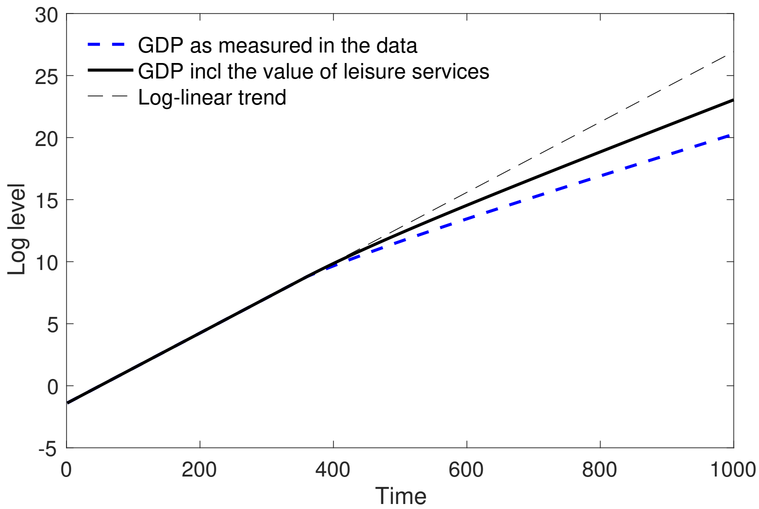

Figure 1.10: Level of two measures of GDP on the sBGP

and hours worked can be accounted for by leisure-technologies. It is nonetheless interesting to speculate how the model's predictions match with the historical developments. Section 3.1 has documented that leisure-enhancing technologies have been around since at least the 1920s, and became increasingly important in the 21st century. Two basic interpretations of the findings above are thus possible. One is that the "watershed moment" came when the leisure sector first appeared, perhaps a century ago, when free newspapers and magazines became important and when the radio first became popular (early 1920s). The other focuses on the setup of the problem, and in particular on the feature that the *work-leisure margin is key*. In the model, improved leisure options crowd out working time, including time spent creatively on generating new ideas. According to this interpretation, the watershed moment corresponds more closely to the inventions of the internet and mobile technologies, which arguably made the competition between leisure and work much more direct and quantitatively important. In other words, improving TV services may have crowded out other leisure activities and housework, while the smartphone may be eating into work and creative hours, dragging on productivity growth as a result. Future work may usefully enrich the framework to provide more concrete quantitative answers to these issues.

1.3.7 Measurement

As highlighted in Section 1.2, the theory can also contribute to the debate about best ways to measure the modern economy. Figure 1.10 compares two measures of GDP along the balanced growth equilibrium. The dashed blue line shows GDP as currently measured in the data - that is, treating the zero price leisure services as an intermediate input into the advertising industry. The second measure, in the solid black line, illustrates the path of GDP incorporating the value of leisure services by multiplying the time spent on these services by the current wage rate. The conclusion is that the growth rate of the economy goes down with the arrival of leisure enhancing technologies, but the measured data do not capture the value of leisure services and thus may exaggerate the extent of the slowdown.

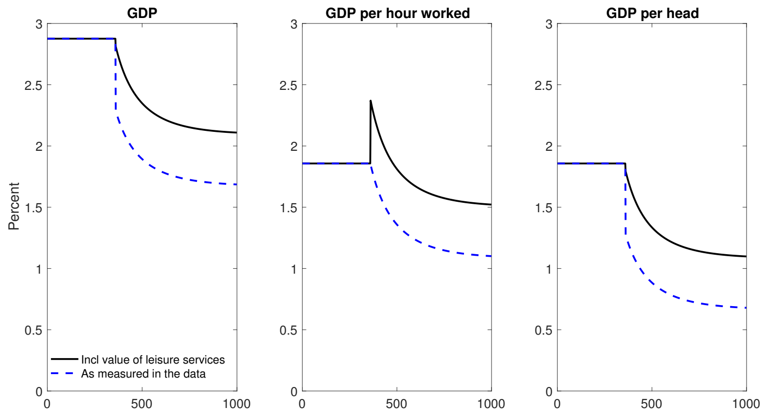

Figure 1.11: Growth rates of GDP and productivity

Figure 1.11 plots growth rates of the level of GDP and of hourly and per-head productivity along the sBGP. Each subplot contains two metrics, analogous to those in Figure 1.10. Along the transition, a gap opens up between the two measures: the economy produces leisure varieties that are not captured in measured statistics. In fact, if one tries to include the value of leisure services in the statistics, and assuming that hours worked are measured correctly, the growth rate of GDP *per hour* actually shoots up. The intuition behind this result is straightforward: the measure $\frac{GDP_{measured}+LeisureServices}{Hours\ Worked}$ rises "on impact" as the denominator declines and the numerator is little changed. In the long-run, however, the decline in the growth of TFP dominates, and we see the decline across all three growth rates in Figure 1.11.

Overall, then, the model provides a novel yet nuanced contribution to the productivity debate highlighted in the Introduction. The rapid leisure-enhancing technological change may be one reason why productivity has slowed down. The model also predicts that the drag on productivity growth may be a permanent feature of the modern economy; a striking finding with numerous policy implications.[31]

1.4 Welfare

Leisure-enhancing technological change introduces novel welfare considerations to the market equilibrium.[32] There are two new inefficiencies. The *static inefficiency* arises because the supply of leisure services is only driven by the profit-maximizing platform. This and the fact that these services are supplied at zero-prices means that the equilibrium level of supply may be suboptimally low. The *dynamic inefficiency* results from the adverse effects of leisure-sector on productivity in the rest of the economy, as documented in Section 1.3, and suggests that supply is inefficiently high. Intuitively, agents do not accurately internalize the impact of today's time allocation decisions on future productivity.

[31]Of course, modern technologies benefit the economy through multiple channels that are completely absent from the model, which may well ultimately prevail and lead to higher productivity and prosperity.

[32]In addition to the inefficiencies familiar from the literature on optimal growth: the presence of monopolistic competition and externalities to R&D.

1.4.1 Static inefficiency

In the static setting, the planner's problem is to maximize household's utility subject to the resource constraint, time endowment constraint, available technologies and non-negativity constraints. The problem can be written as follows:

$$\max_{c,X,X_B,H_L} \log(c) + l \tag{1.36}$$

subject to:

$$Y = c + X_B + \int_0^A x(i)di$$

$$Y = \int_0^A x(i)^\alpha H_W^{(1-\alpha)} di$$

$$M = X_B$$

$$l = H_L M^{\frac{1}{\nu-1}}$$

$$H_W + H_L = 1$$

$$0 \leqslant H_L \leqslant 1,\ 0 \leqslant H_W \leqslant 1$$

$$X_B, X, M \geqslant 0.$$

From the setup of this problem, the following Proposition follows immediately:

Proposition 10 *A direct implication of the planner's problem (1.36) is that B=0 in the optimal allocation.*

Advertising per se is socially useless in this model, so the optimal solution is to set $\mathcal{O} = 0$, that is, to get rid of the marketing sector altogether. In the optimal allocation leisure services are provided directly, and not as a by-product of marketing activity.

Full characterization of the optimal allocation requires the use of a computer.[33] I proceed by solving this maximization problem numerically with parameter values specified in Table 1. Figure 1.12 illustrates the objective of the planner as a function of leisure hours for different values of A. The black line traces out the optimal choice of leisure hours. When A is small, the optimal allocation is at the corner; for A large enough, the choice of leisure is increasing in A. Thus the supply of leisure services in the socially optimal allocation is qualitatively similar to the market allocation. The difference is qualitative: recall that the condition for the marketing sector to emerge in the market equilibrium is $A > \frac{1}{\kappa}$, which, for the current parametrization, translates into $A > 88,000$. For comparison, the planner provides positive supply of leisure services when $A > 100$. So the key takeaway from the static setting is that there is under-provision of leisure services in market equilibrium.[34]

[33]The Appendix outlines the steps necessary to reduce the dimensionality of the problem above.

[34]The high threshold in the market equilibrium relative to the planner solution is in large part driven by the parameter χ - the perceived effectiveness of advertising. This parameter drives the willingness to pay for advertising by the consumer firms, and thus is the key determinant of the profitability of the platform. The low value of 0.05

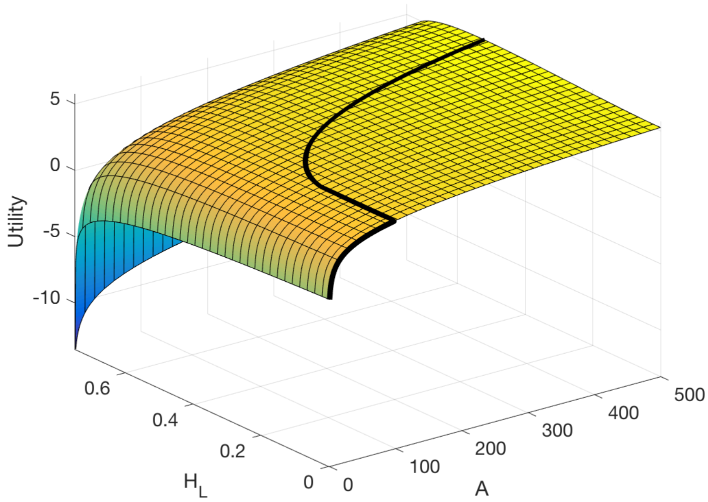

Figure 1.12: Planner's objective as a function of H_L for different values of A, and the optimal allocation

The economics underlying the static inefficiency is not new: a similar mechanism was identified by Spence and Owen (1977), who analyze the television market. They highlight that the marginal cost of supplying a TV program to an additional consumer is zero, and thus efficiency calls for a zero price. Under the ad-supported business model, the pricing is indeed efficient. The authors note, however, that *"the program may not be supplied unless revenues cover the cost of producing it, a cost that is independent of the number of viewers. The revenue under advertiser support comes from advertisers who pay a price of roughly two cents per viewer per hour of prime time. The issue with respect to program selection then is whether two cents is a reasonable estimate of the average value of the program to the viewers of it. If it is not, then revenues may understate the social value of the program, and some programs with a potential positive surplus may not be profitable."* This is very similar to the inefficiency identified above: the leisure sector may not be active if the revenues from providing advertising do not compensate for the cost of supplying leisure services.

1.4.2 Dynamic inefficiency

The static setting offers only partial answers to the welfare implications of the leisure enhancing technological progress. A major distortion in the full dynamic model is that household's choices of leisure hours have a long-term impact on the on the growth rate of productivity. This effect is external to the household's decision, which results in a major inefficiency of the market solution.

To see this clearly, it is useful to set up the optimal control problem in the dynamic setting. For the purposes of this problem, it is useful to make a simplifying assumption that s is given. The problem of the planner is then:

assumed in the calibration makes the platform unprofitable unless the demand for advertising is boosted by the large number of consumer goods producers.

$$\max_{\{c,H_W\}_{t=0}^{\infty}} \int_0^{\infty} e^{-\rho t}\left(\log(c)+l\right)dt; \quad c \equiv \frac{C}{N},\ \frac{\dot{N}}{N}=n$$

subject to

$$C+X_B=Y-X=AL_Y\hat{\alpha}$$

$$l=H_L M^{\frac{1}{\nu-1}}$$

$$M=X_B$$

$$H_W+H_L=1$$

$$0\leqslant H_W\leqslant 1$$

$$\dot{A}=\zeta A^{\phi}L_A^{\lambda}$$

$$L_A=sH_WN,\ s\in[0,1]$$

where, in the first constraint, $\hat{\alpha}=\alpha^{\frac{\alpha}{1-\alpha}}-\alpha^{\frac{1}{1-\alpha}}$, C,X,X_B are all weakly positive and the second equality follows from substituting in the optimum quantity of each variety of intermediate goods from equation (3.17).

There are two control variables of this problem - consumption per capita c and hours worked H_W. The stock of knowledge A is the state variable. The problem's solution is not guaranteed to be interior: in particular, there are two constraints that may bind in the optimal allocation: $X_B\geqslant 0$ and $H_W\leqslant 1$. Let the Lagrange multipliers on these constraints be ψ_1 and ψ_2 and let μ denote the costate. We can then set up the following Hamiltonian (suppressing the time subscripts):

$$\mathcal{H}=\log(c)+(1-H_W)X_B^{\frac{1}{\nu-1}}+\mu\left(\zeta A^{\phi}(sH_WN)^{\lambda}\right)$$

and the Lagrangian:

$$\mathcal{L}=\mathcal{H}+\psi_1X_B+\psi_2(1-H_W).$$

The necessary conditions are as follows:

$$X_B:\quad -\frac{1}{cN}+\frac{1}{\nu-1}(1-H_W)X_B^{\frac{2-\nu}{\nu-1}}+\psi_1=0 \tag{1.37}$$

$$H_W:\quad \frac{1}{c}A(1-s)\hat{\alpha}-X_B^{\frac{1}{\nu-1}}+\mu\zeta\lambda A^{\phi}(sH_WN)^{\lambda-1}sN-\psi_2=0 \tag{1.38}$$

$$\psi_1X_B=0;\quad \psi_2(1-H_W)=0 \tag{1.39}$$

$$A:\quad \frac{1}{c}(1-s)H_W\hat{\alpha}+\mu\zeta\phi A^{\phi-1}(sH_WN)^{\lambda}=\rho\mu-\dot{\mu} \tag{1.40}$$

$$TVC:\quad \underset{t\to\infty}{lim}[e^{-\rho t}\mu A]=0 \tag{1.41}$$

Expressions in (1.39) represent the complementary slackness conditions associated with the two non-negativity constraints $X_B\geqslant 0$ and $H_L\geqslant 0$. It is easy to see that either both constraints bind or neither binds. Condition (1.41) is the transversality condition, requiring that the value of the state variable in utility terms is zero in the limit as $t \to \infty$.

To see how the planner corrects for the dynamic inefficiency present in the market equilibrium, consider equation (1.38). This condition equates the marginal benefit from decreasing hours worked - equal to the additional utility from leisure, the term $X_B^{\frac{1}{\nu-1}}$ - to the cost of doing so. The cost comprises of two components: first, as usual, lowering hours worked today immediately lowers the resources available for consumption (term $\frac{1}{c}A(1-s)\hat{\alpha}$); second, lowering hours worked today diminishes productivity growth going forward (term $\mu\zeta\lambda A^{\phi}(sH_W N)^{\lambda-1}sN$ in the equation). The planner accounts for this future cost of leisure, correspondingly lowering the present supply of leisure services.

1.5 Conclusion

This paper put forward the idea that leisure-enhancing technologies, ubiquitous today and present over the past century, are key to understand the growth process. Distinguishing between the traditional productivity-enhancing innovations and the leisure-enhancing technological progress can shed new light on some of the unanswered questions and puzzles in applied theoretical work and in policy. In particular, the theory can account for the trends in hours worked, interest rates and productivity. The framework also highlights the non-trivial welfare implications of the leisure sector. This preliminary effort opens up several avenues for further exploration.

It may be useful to recast the mechanisms highlighted here in alternative growth modeling frameworks, for instance in the spirit of quality innovations as in Grossman and Helpman (1991) and Aghion and Howitt (1992b) or in the learning and search framework of Lucas and Moll (2014). Furthermore, the model presented here was kept particularly parsimonious, especially when parametrizing production functions of the marketing platform, and future research may consider more detailed analysis of the functional forms and the key parameters. Also, the analysis rested on the assumption that leisure services are provided free of charge; future work could generalize the problem of the platform to explicitly model the provision of partly subsidized services, which could also help bring the model more directly to the data. Finally, the model could be employed in the analysis of the cyclical behavior of macroeconomic wedges, following on from Hall (2014).

One outstanding question is how the rise of the leisure sector interacts with heterogeneity, both at the household and at firm level. On the households side, the interesting question concerns time allocation decisions and their interplay with inequality, broadly defined. Disaggregated evidence on time allocation across the income distribution shows that poor individuals increased their leisure more than the rich (Boppart and Ngai, 2016). Allowing for household and income heterogeneity in the presence of leisure-enhancing progress could bring out new insights and aid the debate on leisure-inequality. Considering firm heterogeneity may be interesting too: the current setting is well suited to analyze equilibrium outcomes when heterogenous firms compete not only in prices but also through advertising. More productive firms may devote more resources to brand building, cementing their market share, with interesting implications for market power (De Loecker and Eeckhout, 2017).

A richer analysis of normative implications of leisure-enhancing technologies could be useful. The current paper is based on the revealed-preference argument. In reality, the issue is likely to be more complex. For example, recent research from psychology documents adverse effects that the recent technologies have on well-being (Verduyn et al., 2017). Over and above that, the current paper takes a neutral stance on the welfare impact of advertising itself, while in reality both positive and negative channels of influence arise. On the upside, advertising may provide consumers with valuable information or may have complementarities with consumption if consumers value consuming a good that is well advertised. On the downside, advertising supplied together with leisure services may diminish the utility from consuming these services. Given the prevalence of ad-supported business model, future efforts that focus on the balance of these forces would be useful.

Finally, the framework built here can help answer some of the policy questions, such as optimal taxation of the providers of free leisure services. I plan to pursue these ideas in future work.

Chapter 2

Will robots take our jobs? Automation and unemployment in general equilibrium

2.1 Introduction

Technology has been reshaping the economy in a profound way. Automation of routine work is widely credited as the key contributor to job polarization (Autor (2010), Acemoglu and Autor (2011)). The rise of industrial robots appears to have had a significant negative effect on employment and wages in local labor markets (Acemoglu and Restrepo, 2017) and may have contributed to the decline in the global labor share (Acemoglu and Restrepo, 2019). Technologies of the future – such as industrial robots or artificial intelligence – bring great opportunities, but also fears and uncertainties with regards to the future relationship between labor and capital, sources of demand and employment, and even risks of further political fragmentation.

Researchers who attempt to quantify the likelihood of a 'robot takeover' largely find that the risks are indeed significant. In one widely cited paper, Frey and Osborne (2013) assert that nearly half of all the existing jobs in the US are amenable to automation. McKinsey Global Institute (2017) reaches a similar conclusion looking at jobs across the globe. A more granular analysis at the task level suggests that no job will be left untouched by the latest technologies, although few will become fully automated (Brynjolfsson, Mitchell, and Rock (2018)).

The broader public seem to be aware of the risks. A Pew Research Center survey of over 4,000 U.S. adults conducted in 2017 has found that many Americans anticipate significant impacts from various automation technologies in the course of their lifetimes – from the widespread adoption of autonomous vehicles to the replacement of entire job categories with robot workers.[1] It reports that around 80% of Americans think that robots being able to do most jobs currently done by humans is a realistic prospect, and 75% are concerned or very concerned about this.

This evidence is largely at odds with the standard macroeconomic model of representative agent and labor-augmenting technologies. In such a framework, in the long-run capital always makes workers more productive, raising average wages (Caselli and Manning, 2018); moreover,

[1] Results are available at https://www.pewinternet.org/2017/10/04/automation-in-everyday-life/

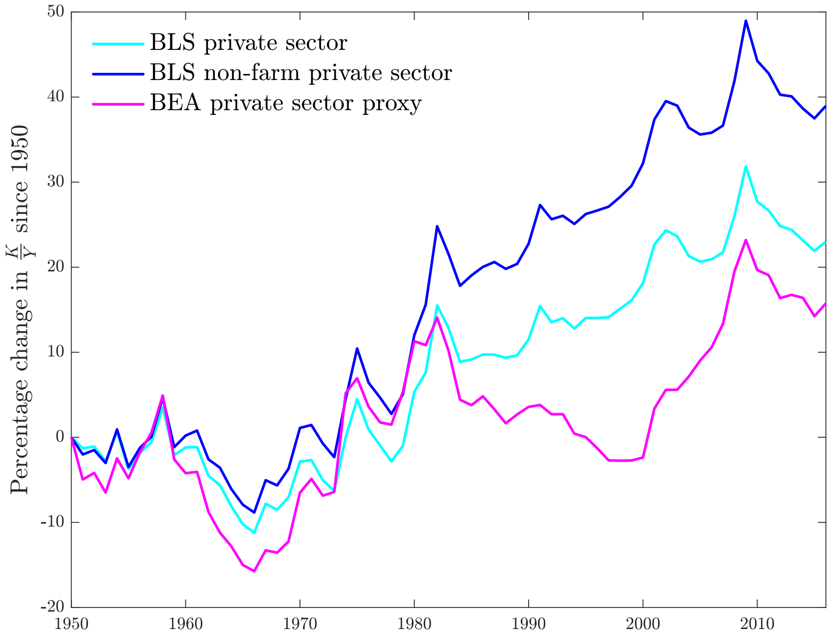

Figure 2.1: Capital-to-output ratio in the United States
Notes: The BLS lines are real capital services / real VA output of each sector, and the BEA line is current cost capital divided by current gross value added of the non-financial corporate business sector. The merits of these different measures depend on the context. The BLS measure is better for thinking about the importance of capital in production, while the BEA measure is a better measure of wealth.

perfect insurance at the individual level and no unemployment in equilibrium mean that the impact of technology is evenly spread across workers.

In this paper I develop a more realistic setting for analysis of technology, featuring equilibrium unemployment, incomplete insurance markets and endogenous technology choice. In this framework individual uncertainty and heterogeneity play an important role. I then attempt to answer two sets of questions.[2]

First, I want to understand the impact of automation technologies on the labor market. Under what circumstances could technological unemployment emerge? The process of automation is not new, and past experience tells us that the economy is capable of absorbing the resources that have been left idle by technological change, at least in the long run. For example, Figures 2.1 and 2.2 show that, in the United States, capital-to-output ratio has increased dramatically over the past 50 years, and yet unemployment rate has remained stable (Autor, 2015).[3] At the same time, average wages lagged behind measured productivity, driving a decline in the labor share. My aim is to use the model to shed new light on these salient facts.

The second set of issues that I study concerns the general equilibrium nature of the adjustment to technological changes. Technology brings about deep shifts in the economic environment. Examples include changes to individual job security, shifts in relative prices and in the distribution of wages. How do those effects play out in equilibrium? Are the general

[2]In a separate paper I consider the implications of automation for the distributions of income and wealth (Moll, Rachel, Restrepo (forthcoming)).

[3]The decline in labor force participation in the United States meant that employment-to-population ratio is currently below its early 2000s levels, but it is still higher than in the 1970s and 1980s. Nonetheless, participation margin – missing from the model in this paper – is potentially important. I plan to incorporate it in future work.

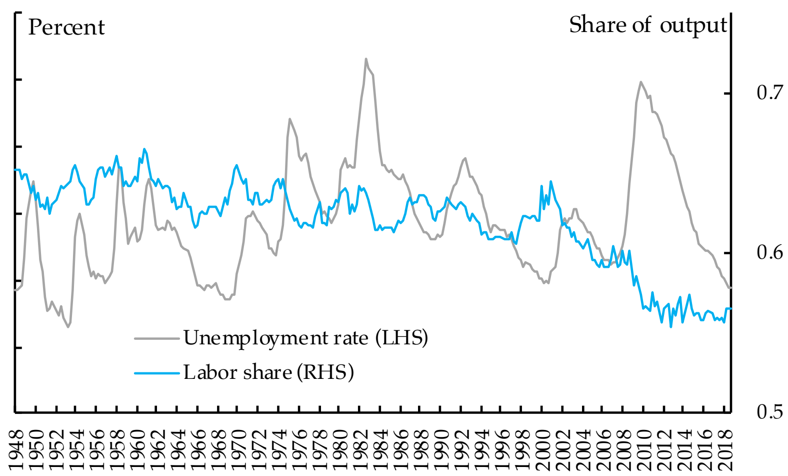

Source: FRED and the BLS. Note: The labor share is for the non-farm business sector.

Figure 2.2: The unemployment rate and the labor share in the United States

equilibrium effects amplifying or dampening the technological shocks? Is there a possibility of aggregate externalities, and, if so, what policies may be needed to deal with it?

To begin to tackle these questions, I analyze how two kinds of technological shifts – rising productivity of capital and invention of new automation technologies – affect the steady state and full dynamic adjustment path of the economy.

In parts of the economy which have been technologically automated, producers decide whether to use workers or machines. These decisions are driven by these technologies' relative costs. Firms that choose to rely on labor must post vacancies, search for workers, and bargain with them over wages. Firms that use machines rent them in a competitive capital rental market. Factor prices, determined in general equilibrium, can drive technology adoption.

Search frictions in the labor market mean that in equilibrium some workers find themselves unemployed. Workers lack state-contingent instruments to insure themselves against the risk of unemployment: they can save only in capital stock and in equities offering a certain rate of return. Lack of insurance means that individuals experience different income histories and thus accumulate different levels of wealth. The equilibrium features a full distribution of workers across the asset spectrum. In a stationary equilibrium this distribution is time invariant, even as at the micro-level agents continuously change their position in the distribution, accumulating assets through saving when employed and deploying their buffers when out of a job. Moreover, because individual workers bargain over the surplus from a match with their to-be-employers, the ex-post heterogeneity in assets translates into heterogeneity in wages: the equilibrium features a *wage schedule*. The intuition for this finding, originally due to Krusell, Mukoyama, and Sahin (2010), is that a surplus from a job depends on each worker's individual asset holdings. Richer workers are less desperate to find employment, which strengthens their effective bargaining position and results in higher wages. In this way, wealth inequality begets income inequality.

Given this setup and the questions I set to shed light on, I uncover several novel results.

First, I derive a relationship that links demand for automation to the interest rate. To do so, I solve firms' technology choice problem. This problem is non-trivial at the micro level, but is made simpler through the use of the ideal price condition that must hold in equilibrium. The solution features a relationship that determines the equilibrium adoption of machines only on the basis of the interest rate. I show that, under general conditions, this *automation demand schedule* is highly non-linear and convex, meaning that the sensitivity of automation to the cost of capital is increasing in the existing level automation. This is an important result that cautions about the inference on the impact of automation from historical episodes.

Second, I uncover two novel general equilibrium channels that operate when the economy is hit by automation waves. Both channels raise the desire to save, which tends to depress the interest rate and feeds back into further adoption of automation technologies. The first is the *risk-mitigation effect*: to the extent that automation leads to higher unemployment and more labor market uncertainty, it raises saving for precautionary reasons. The second is the *saving for higher wages* effect. I show that automation leads to a steepening of the wage schedule, as asset-rich workers are in a stronger bargaining position and thus fare better than their poor counterparts following a technology shock. This incentivizes workers to save in order to "climb up" the wage schedule. I offer some very preliminary illustrative evidence for some of these channels, and plan further work to bring them to the data. My favored interpretation for the moment is broader than the specific mechanisms embodied in my model: it is that at the level of an individual, savings may help insure against adverse future technology shifts; but at the macro level, a higher desire to save leads to further technology adoption and thus exacerbates the risk. Thus a broad lesson from this research so far is that the general equilibrium effects are interesting and testing the size of these effects must be a top priority for future work.

The third result concerns the path of unemployment following automation shocks. I show that, beyond the immediate shake-out, automation *reduces* unemployment in the short-to-medium-run. The reason behind this is related to the GE channels described above: with higher automation, workers' effective bargaining power deteriorates, leading to a decline in wages relative to productivity and increasing labor demand. In the long-run, asset accumulation restores workers' bargaining power, which may lead to higher unemployment.

Related literature. This paper builds on and contributes to several strands of the literature. It is most closely related to the literature on the economics of labor-replacing technologies. A classic reference in this literature is Zeira (1998), who develops a framework in which firms choose between labor-intensive and capital-intensive production technologies, and links these adoption choices to the long-run levels of development and the observed large differences in GDP per capita across countries. The model developed here builds on the core of that framework and extends it in several important respects. Most importantly, Zeira (1998) considers a small open economy case, whereby the interest rate is constant and exogenous; here instead the interest rate is determined in general equilibrium. This turns out to be central to the discovery of novel mechanisms at play. More recently, the important contributions to this literature are Hémous and Olsen (2018), Acemoglu and Restrepo (2018) and Aghion, Jones, and Jones (2017). Like the first of these papers, the present model analyzes on the role of factor prices, but unlike that paper, the focus is on the equilibrium interest rate (rather than high-skilled vs. low-skilled wages) and on technology adoption (rather than innovation). The latter two papers are fundamentally interested in the long-run growth properties of an economy where some tasks are automated. Acemoglu and Restrepo (2018) derive the balanced growth path of an economy in which existing tasks are automated and new tasks are created. The present paper complements theirs, making simplifying assumptions along some dimensions but extending the analysis along others. It abstracts from the process of the production of new tasks, and instead considers a detrended economy in which an automation shock is equivalent to "automation running ahead of the creation of new tasks" in the Acemoglu and Restrepo (2018) framework. In this sense this is the model of *relative – or net – automation*.[4] The model developed here provides a much more detailed treatment of the labor and insurance markets. Presence of equilibrium unemployment gives my model the ability to speak to the question of the link between automation and involuntary unemployment. Incomplete markets framework allows for the study of rich interactions between the economic environment and individual behavior in general equilibrium, and in particular on the links between the adoption of automation and labor market uncertainty and job security. Aghion, Jones, and Jones (2017) focus on the Baumol cost disease – namely the observation that growth is constrained by the activities with lowest productivity. An important role in this intuition is what happens to relative prices. My work here maintains this focus, even though it abstracts from the structural complementarity between tasks that is the key in driving their results.

Second, the paper contributes to the literature that studies economies with idiosyncratic risks and ex-post wealth heterogeneity in the Bewley-Hugget-Aiyagari tradition (see, for example, Oh and Reis (2012), Heathcote, Storesletten, and Violante (2009) and Quadrini and Ríos-Rull (2014)). In particular, this paper draws on the study by Krusell, Mukoyama, and Sahin (2010), who incorporate the labor market search frictions into an incomplete markets general equilibrium model. The key difference between their approach and the one developed here is that I take advantage of the developments of the continuous time techniques that allow for an analytically elegant and computationally efficient solution to these models,[5] and I further enrich the model by adding a task-based framework and endogenizing the choice of technology.

Finally, the third strand of the literature which this paper contributes to concerns the interaction of technology and the labor market. In particular, the paper combines the task-based framework of Autor, Levy, and Murnane (2003) and Acemoglu and Autor (2011) with the frictional labor market of Diamond-Mortensen-Pissarides tradition (Pissarides (1990)). The baseline version of the model assumes no skill heterogeneity, in contrast to the focus of the literature on skill-biased technical change (Katz and Murphy (1992), Acemoglu (1998), Caselli (1999), Krusell, Ohanian, Rios-Rull, and Violante (2000), Aghion, Howitt, and Violante (2002),

[4]My model is also simpler in that it assumes unitary elasticity of substitution between tasks in the production of the final good – although the bulk of the results would be unchanged with a more general production function. Indeed, Acemoglu and Restrepo (2018) show what assumptions are needed for the conclusions from the analysis of the Cobb-Douglas production function to extend to more general functional forms.

[5]The general tools and techniques are developed in Achdou et al. (2017), while the labor market equilibrium in continuous time is analyzed in Bardóczy (2017).

Acemoglu (2007)). Instead, worker heterogeneity stems only from idiosyncratic employment histories and the consequent saving behaviors. This is not to undermine the importance of skill heterogeneity in the context of adaptability to technology – clearly, skills play a critical role in this context. But abstracting from skill differences helps isolate one of the contributions of this paper, which is to highlight the importance of individual wealth in driving outcomes.

The rest of the paper is structured as follows. Section 2.2 outlines the model and defines the equilibrium. Section 2.3 considers how changes in automation affect labor demand. The next two Sections analyze the effects of two distinct technological shifts: Section 2.4 considers an increase in the efficacy of existing automation technologies, while Section 2.5 analyzes the effects of a shift in the technological automation threshold driven by the invention of new automation technologies. Finally, Section 2.6 concludes by summarizing the key findings and discussing avenues for future work.

2.2 The model

The model, written in continuous time, combines the Bewley-Hugget-Aiyagari framework of incomplete markets and precautionary saving and the frictional labor market model of Diamond-Mortensen-Pissarides with the task-based model of technology adoption which can be traced back to Zeira (1998). I now outline the details.

2.2.1 Technology

2.2.1.1 Final good

Perfectly competitive final sector firms use a Cobb-Douglas technology, combining a continuum of intermediate inputs (or tasks) indexed by $j \in [0, 1]$ in the production of the final good:

$$\ln(Y(t)) = \int_0^1 \ln(y(j,t))dj. \tag{2.1}$$

Normalizing the price of consumption good to 1 and denoting the price of task j by $p(j,t)$, the profit maximization of final sector firms gives the demand schedule for each task:

$$y(j,t) = \frac{Y(t)}{p(j,t)}. \tag{2.2}$$

For future reference, note that plugging (2.2) into (2.1) yields the ideal price condition:

$$\int_0^1 \ln(p(j,t))dj = 0. \tag{2.3}$$

2.2.1.2 Intermediate tasks

The technological assumptions for production of intermediate tasks j are based on the task-based framework of Zeira (1998), Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018). Tasks are produced by intermediate producers in a perfectly competitive product market with free entry. Let $\mu \in [0,1]$ denote the technologically feasible automation threshold. The technologically automated tasks, $j \in [0,\mu]$, can be produced using either the automated technology which only requires capital, or the manual technology which only requires labor. Other tasks can be produced only with the manual technology. In addition, tasks differ in how amenable they are to automation. Without loss of generality, assume that tasks are ordered according to how easy it is to automate them, so that task $j = 0$ is the easiest to automate. Formally, assume that unit capital requirement of task j is given by $\gamma(j)$, where function γ satisfies Assumption 1:

ASSUMPTION 1: *function* γ *is continuous and differentiable, and satisfies* $0 < \gamma < +\infty$ *and* $0 < \gamma' < +\infty$ *for all* $j \in [0,1]$.

Assumption 1 implies that the elasticity of the capital requirement to the movement along the task spectrum defined as

$$\epsilon_{\gamma(j)} \equiv \frac{\partial \gamma(j)}{\partial j}\frac{j}{\gamma(j)}$$

is positive and finite everywhere.

I assume that unit labor requirement is identical across tasks and equal to $\bar{l}$ (equivalently, there is constant marginal productivity of labor across tasks equal to z, with $z = \frac{1}{\bar{l}}$), and I normalize $\bar{l} = 1$. Summarizing, we have:

$$y(j,t) = \begin{cases} l(j,t) + \frac{k(j,t)}{\gamma(j)} & j \in [0,\mu] \\ l(j,t) & j \in (\mu,1] \end{cases}.$$

The assumption that the function γ is increasing means that labor has comparative advantage in production of tasks with a higher index. This guarantees that in equilibrium the tasks with a lower index will be automated, while those with the higher index will be produced with labor.

The task spectrum is depicted in Figure 2.3. Denote by $\tilde{m}(t)$ the unconstrained optimal level of automation – that is, the level of automation that would prevail given the current factor prices under a slack technological constraint ($\mu = 1$). Given the ordering assumption, *equilibrium automation* is the threshold task $m^*(t) \in [0,\mu]$ such that, in equilibrium, all producers in the interval $[0, m^*(t)]$ use the automated technology and all firms $j \in (m^*(t), 1]$ use the manual technology. There are two cases, depending on whether or not all existing automation technologies are adopted: technological constraint may be binding: $m^*(t) = \mu < \tilde{m}(t)$ or slack: $m^*(t) = \tilde{m}(t) < \mu$.

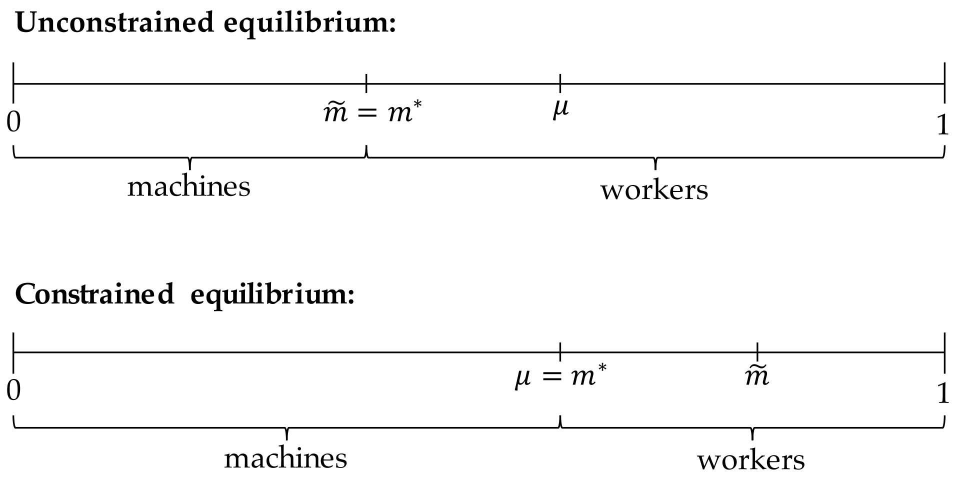

Figure 2.3: The task spectrum

2.2.2 Firms' pricing decisions

The market for physical capital is perfectly competitive and there is perfect competition in the product market within each task j. Hence profits are driven to zero and prices of automated tasks equal the unit cost of production:

$$p(j,t) = (r(t) + \delta)\gamma(j) \quad \forall j \in [0, m^*(t)]. \tag{2.4}$$

This equation states that price equals cost of production of task j, which is simply the cost of capital (the rental rate $r(t)$ plus the depreciation rate δ) times the unit capital requirement $\gamma(j)$.

All firms that employ workers have the same technology and they hire from the same pool of workers, thus they face the same expected cost of production. Free entry to vacancy creation (discussed below) and perfect competition within each j imply that they must charge the same price in equilibrium. Denoting this price by $\bar{p}$, we have:

$$p(j,t) = \bar{p}(t) \quad \forall j \in (m^*(t), 1]. \tag{2.5}$$

2.2.3 Technology choice

Perfect competition in product markets ensures that firms always use the cheapest technology available. The ordering of tasks – with labor having comparative advantage in tasks with a higher index – guarantees that there exists a unique threshold $\tilde{m}(t)$ defined by:

$$(r(t) + \delta)\gamma(\tilde{m}(t)) = \bar{p}(t). \tag{2.6}$$

Without any further constraints, tasks below $\tilde{m}(t)$ would be produced with machines, and tasks above it with labor. However, automation is feasible only for the tasks at or below μ and so:

$$m^*(t) = \min\{\tilde{m}(t), \mu\}. \tag{2.7}$$

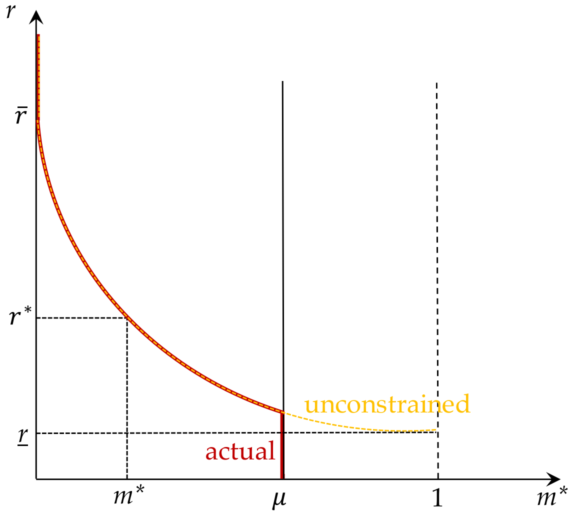

Figure 2.4: Demand for automation

How do firms decide which technology to use? Frictions in the labor and asset markets make this a difficult problem: each individual producer needs to make this decision by comparing the cost of capital $r+\delta$ with expected cost of opening and filling a vacancy, which depends on expected wages, expected incurred search costs, labor market tightness and so on.

It turns out that it is easier to solve this problem in general equilibrium rather than at a level of an individual producer. This can be achieved by noting that we only need to solve this problem in the unconstrained case (as otherwise we have $m^* = \mu$). Making use of the ideal price condition (2.3), we can eliminate all the labor market variables from the problem, leaving the interest rate as a sole determinant of equilibrium automation.

Combining equations (2.3), (2.6) and (2.7) gives:

$$0 = \int_0^{m^*(t)} \ln[(r(t)+\delta)\gamma(j)]dj + \int_{m^*(t)}^1 \ln[(r(t)+\delta)\gamma(m^*(t))]dj \qquad \forall m^*(t) \leqslant \mu.$$

Rearranging this expression yields the *unconstrained demand for automation*:

$$r(t) = \exp\left[(m^*(t)-1)\ln\gamma(m^*(t)) - \int_0^{m^*(t)} \ln\gamma(j)dj\right] - \delta \quad \forall m^*(t) \leqslant \mu. \tag{2.8}$$

This schedule traces out the marginal task produced with machines for every possible interest rate. In the $m^* - r$ space it takes a form of a downward sloping curve, illustrated diagrammatically in Figure (2.4).[6] The following Proposition describes its striking property:

Definition 11 *Assume that Assumption 1 holds, implying that ϵ_γ is positive and finite. Then the automation demand curve is highly non-linear. Specifically, the interest rate changes matter relatively little when interest rates are high and desired automation is low; but they matter a*

[6]The slope of the schedule is $-(r+\delta)\frac{1-m^*}{m^*}\epsilon_{\gamma(m^*)}$ where $\epsilon_{\gamma(m^*)}$ is the elasticity of γ (wrt j) evaluated at m^*. This is negative as long as the interest rate is larger than $-\delta$. As I verify below, this must be the case in equilibrium.

great deal when interest rates are low and automation adoption is already high. In fact, the demand for automation becomes inelastic to the changes in the interest rate when $m^* \to 0$, and becomes infinite as $m^* \to 1$.

Proof. For the purpose of the proof, I drop the time subscripts. The sensitivity of automation demand to changes in the cost of capital is characterized by the slope of the demand curve: $\frac{dm^*}{dr}$. By the chain rule we have

$$\frac{dm^*}{dr} = \left(\frac{dr}{dm^*}\right)^{-1}.$$

Differentiating (2.8) gives:

$$\frac{dr}{dm^*} = (r + \delta)\left[\epsilon_{\gamma(m^*)}\left(1 - \frac{1}{m^*}\right)\right]$$

The sensitivity of the demand schedule is thus

$$\frac{dm^*}{dr} = -\frac{1}{(r + \delta)\left[\epsilon_{\gamma(m^*)}\left(\frac{1}{m^*} - 1\right)\right]}.$$

Which gives $\lim_{m^* \to 0} \frac{dm^*}{dr} = 0$ and $\lim_{m^* \to 1} \frac{dm^*}{dr} = -\infty$. ■ Proposition 11 means that a given change in the cost of capital may have very different consequences for equilibrium automation, depending on how prevalent the use of machines is. To grasp the intuition behind this non-linearity first note that the demand curve is downward sloping because lower r translates into macro-level cost-saving from the cheaper production of the tasks already automated, which in turn allows for more expensive (higher $\gamma(j)$) technologies to be employed. These cost savings are relatively modest when automation is low, as only a small range of already-automated tasks benefits directly from this change. But when automation is high, a decline in the cost of capital constitutes a larger fall in the aggregate cost of production, allowing for higher automation adoption in equilibrium.

It is straightforward to pin down the values of thresholds $\bar{r}$ and $\underline{r}$ for which neither or all sectors in the economy want to produce using machines. First, none of the automation technologies are adopted if the interest rate is greater than the upper threshold $\bar{r}$:

$$\bar{r} = \frac{1}{\gamma(0)} - \delta.$$

If technological automation is complete, i.e. $\mu = 1$, then there will be full automation if the interest rate is less than or equal to:

$$\underline{r} = \exp\left[-\int_0^1 \ln\gamma(j)dj\right] - \delta.$$

The comparative statics on these thresholds are intuitive: the upper threshold $\bar{r}$ depends only on the productivity of the the most advanced automation technology. The higher is this productivity – the lower is the capital requirement $\gamma(0)$ – the higher will be the interest rate that would prevent any automation technologies to be adopted. The lower threshold $\underline{r}$ depends inversely on capital requirements of all technologies.

2.2.4 The labor market

There are search frictions in the labor market: to hire a worker, a firm must first post and advertise a vacancy, which has a flow cost of ξ for as long as the vacancy remains unfilled. The number of successful matches depends on the number of unemployed workers seeking jobs and on the number of vacancies. More specifically, the number of matches is governed by the matching function $M(u, v)$, assumed to be homogenous of degree one. Let θ denote the labor market tightness, q denote the vacancy filling rate, and f denote the job finding rate, so that $\theta = \frac{v}{u}$, $q \equiv \frac{M(u,v)}{v} = q(\theta)$ and $q' < 0$, $f \equiv \frac{M(u,v)}{u}$. Job matches can separate for two reasons. First, all jobs are hit by an exogenous job destruction shock at a rate σ. Second, a job match is terminated when a firm decides to switch from a manual towards an automated technology. I now derive an expression for separation rate in this economy.

Let $\bar{l}(t)$ denote mass of employees per task. The aggregate employment in this economy is:[7]

$$N(t) = (1 - m^*(t))\,\bar{l}(t). \tag{2.9}$$

Denote the (endogenous) job finding rate with $f(t)$. Differentiating (2.9) with respect to time, we have:

$$\dot{N}(t) = -\dot{m}^*(t)\bar{l}(t) + (1 - m^*(t))\left(-\sigma\bar{l}(t) + f(t)\frac{1 - N(t)}{1 - m^*(t)}\right)$$

Rearranging:

$$\dot{N}(t) = -\sigma N(t) + f(t)(1 - N(t)) - N(t)\frac{\dot{m}^*(t)}{1 - m^*(t)} \tag{2.10}$$

Equation (2.10) shows that the actual separation rate – denoted by ς – is the sum of the exogenous separation rate σ and an endogenous separation due to automation:

$$\varsigma = \sigma + \frac{\dot{m}^*(t)}{1 - m^*(t)} \tag{2.11}$$

In a stationary equilibrium, when technology adoption does not change, the separation is driven entirely by the exogenous separation rate. But when firms adopt automation technologies, separation increases reflecting the shake-out as workers previously employed in those firms lose their jobs. Expression (2.11) implies that the initial labor market shakeout due to technology is larger with higher starting level of m^*. This is intuitive: when workers are crowded in on a small spectrum of tasks, the mass of employees per task is high and the shake-out associated with a given change in adoption of technology is large.

Outside of the stationary equilibrium, the knowledge of future automation of a particular task j currently employing a worker could, in principle, impact on firm's and worker's behavior. To limit the complexity of the model, I assume the following:

[7] Normalizing the labor force to unity, the unemployment rate is given by $u(t) = 1 - N(t)$.

ASSUMPTION 2: *Firms and workers know only the macro job destruction rate and not the likelihood or timing of destruction of their particular task j.*

Assumption 2 makes a worker indifferent between all potential matches.[8]

2.2.4.1 Hamilton-Jacobi-Bellman equations

Once matched, a worker and a firm engage in bargaining to determine the wage. For any worker, the value she attaches to having a job depends on her current wealth. Thus personal wealth holdings will influence her bargaining position and her wage. This result, originally due to Krusell, Mukoyama, and Sahin (2010), means that a worker's wage is a function of her wealth: there is a wage schedule $w(a)$ in equilibrium. A consequence of this dependence is that the value of any filled job will depend on the asset holdings of the individual worker employed in that job.

The following HJB equations characterize the value of a filled job and value of a vacancy, respectively:

$$r(t)J(a,t) = \bar{p}(t) - w(a,t) + J_a(a,t)\dot{a}(a,s_e,t) + \varsigma\left(V(a,t) - J(a,t)\right) + \frac{\partial}{\partial t}J(a,t) \quad (2.12)$$

$$r(t)V(t) = -\xi + q(t)\int_{\underline{a}}^{\infty} J(a,t)\frac{g(a,s_u,t)}{u(t)}da + \frac{\partial}{\partial t}V(t). \quad (2.13)$$

The flow value of the filled job is the flow of profit that the job generates (revenue $\bar{p}$ minus the wage $w(a,t)$), plus "capital gain", which consists of three parts: the change in the value of the job due to changes in worker's asset holdings, the change in value due to possible separation of the firm and the worker, and the change in the value over time, due to some exogenous shocks.

Equation (2.13) has a similar structure and intuition. It says that the flow return to a vacancy is a negative of the vacancy posting cost plus the expected change in value when the vacancy is filled. This expected change in value when a firm is matched with a worker is equal to the vacancy-filling probability $q(t)$ times the expected value of the job, where the expectation is taken across all unemployed searching for a job ($\underline{a}$ denotes the minimum level of assets that workers can hold, and $g(a,s_u,t)$ is the distribution of unemployed workers across the asset space).[9]

Relative to a textbook DMP model, the novel elements here are (i) the dependence of the value of the filled position on assets held by the worker occupying the job, through the dependence of the wage on worker's asset position; (ii) the fact that the value of the job changes as

[8] This is a simplifying assumption but it is not consequential for the economics I discuss in this paper. In future work I plan to consider cases where knowledge of future automation of a given task affects value functions and impacts on the bargaining in that task. This may bring out additional channels through which automation affects the labor market.

[9] Division by u ensures $\frac{g(\cdot)}{u}$ is a density function, i.e. that it integrates to one.

workers' assets change (the term $\frac{\partial}{\partial a}J(a,t)\dot{a}(a,e,t)$); and (iii) the fact that firms base their job creation decisions on the wage they expect to pay (rather than a single scalar wage).

Free entry implies:

$$(r(t)+\varsigma)J(a,t)=\bar{p}-w(a,t)+J_a(a,t)\dot{a}(a,s_e,t)+\frac{\partial}{\partial t}J(a,t) \tag{2.14}$$

and

$$\xi=q(\theta)\int\limits_{\underline{a}}^{\infty}J(a,t)\frac{g(a,s_u,t)}{u(t)}da+\frac{\partial}{\partial t}V(t). \tag{2.15}$$

2.2.5 Households

All workers are employed by the manual sector of the economy. They are infinitely lived, supply labor inelastically and choose consumption and savings to maximize the stream of consumption utility. Crucially, they face uninsurable unemployment risk – markets are incomplete and workers can only self-insure against unemployment through their saving decisions. When unemployed, a worker generates home production output of value h.

Let $W(a_0,s_0)$ denote the value function of a worker whose starting level of assets is a_0 and initial employment status is s_0. A worker whose discount rate is ρ faces the following problem:

$$\begin{aligned}
W(a_0,s_0) &= \max_{\{c(t)\}_{t\geqslant 0}} \mathbb{E}_0\int\limits_0^{\infty}e^{-\rho t}u(c(t))dt\\
\dot{a}(t)+c(t) &= y(a,s,t)+r(t)a(t)-T\\
s &\in \{s_e,s_u\}\\
y(a,s,t) &= \begin{cases} w(a,t) & if\ s=s_e\\ h & if\ s=s_u\end{cases}\\
a(t) &\geqslant \underline{a}.
\end{aligned}$$

The first equation simply says that the value function is a discounted infinite sum of future utility from consumption. Utility function is continuous and twice differentiable, and satisfies $u'>0$ and $u''<0$. The second equation is the flow budget constraint, saying that saving plus consumption equals total income net of lump-sum taxes which the government uses to finance the unemployment benefits. Individual's income is composed of asset income, equal to the net return $r(t)$ times the asset holdings $a(t)$, and labor income $y(a,s,t)$. The notation explicitly highlights labor income's dependence on individual asset holdings a: this dependence is a result of the interaction between incomplete insurance markets and labor market frictions. The third expression specifies that there are two employment states: individuals are either employed or unemployed. From the perspective of an individual worker, changes in employment states are Poisson processes and occur with intensities that depend on the conditions in the labor market. In particular, unemployed workers find jobs with intensity $f(t)$ (the job finding rate), and employed workers lose their job with intensity ς. Finally, the fifth inequality is a borrowing constraint. It says that individual assets cannot be lower than a lower bound $\underline{a}$. This is where market incompleteness shows up: at times, some consumers may want to run down their assets below $\underline{a}$ to cushion their consumption, but they are prohibited from doing so.

It is convenient to write this problem in a recursive form using HJB equations. For an employed worker, the HJB equation is:

$$\rho W(a, s_e, t) = \max_c \{u(c) + W_a(a, s_e, t)\,[w(a,t) + (r(t)a - c)] + \\ + \sigma\,[W(a, s_u, t) - W(a, s_e, t)] + \frac{\partial}{\partial t} W(a, s_e, t)\} \quad (2.16)$$

This equation says that the required return to an employed worker equals the maximized value of the flow return (i.e. consumption utility) plus changes in the value function due to saving, potential loss of employment, and any other influences that change the value function over time.

The first order condition of this problem

$$u'(c) = W_a(a, s_e, t) \quad (2.17)$$

continues to hold at the borrowing constraint, which is the consequence (and the key advantage) of the continuous time formulation. At the borrowing constraint saving cannot be negative: $\dot{a}(a, s, t) \geqslant 0$, and so consumption must not exceed income:

$$c \leqslant w(\underline{a}, t) + r(t)\underline{a}. \quad (2.18)$$

Together with strict concavity of the utility function, inequality (2.18) implies that:

$$W_a(\underline{a}, s_e, t) \geqslant u'(w(\underline{a}, t) + r(t)\underline{a}). \quad (2.19)$$

Equation (2.19) is the state boundary condition of this problem. The first order condition (2.17) readily gives consumption and savings policy functions:

$$c(a, s_e, t) = u'^{-1}(W_a(a, s_e, t)) \quad (2.20)$$

$$\dot{a}(a, s_e, t) = y(a, s_e, t) + r(t)a - c(a, s_e, t) \quad (2.21)$$

The analogous conditions for the unemployed are

$$\rho W(a, s_u, t) = \max_c \{u(c) + W_a(a, s_u, t)\,[h + r(t)a - c] + \\ + f(t)\,[W(a, s_e, t) - W(a, s_u, t)] + \frac{\partial}{\partial t} W(a, s_u, t)\} \quad (2.22)$$

$$W_a(\underline{a}, s_u, t) \geqslant u'(h + r(t)\underline{a}). \quad (2.23)$$

$$c(a, s_u, t) = u'^{-1}(W_a(a, s_u, t)) \tag{2.24}$$

$$\dot{a}(a, s_u, t) = y(a, s_u, t) + r(t)a - c(a, s_u, t) \tag{2.25}$$

2.2.5.1 The labor market and saving behavior

The labor market structure affects household's intertemporal saving behavior in two ways: through the uncertainty with regards to the future employment status, and through the dependence of wages on individual level of assets. The modified Euler Equations of the employed and unemployed, respectively, are given by:

$$\frac{\dot{c}_e}{c_e} = \frac{1}{\psi}\left[r - \rho + w'(a) + \sigma\left(\frac{u'(c_u)}{u'(c_e)} - 1\right)\right] \tag{2.26}$$

and

$$\frac{\dot{c}_u}{c_u} = \frac{1}{\psi}\left[r - \rho - f\left(1 - \frac{u'(c_e)}{u'(c_u)}\right)\right] \tag{2.27}$$

where $\psi = -\frac{u''(c)c}{u'(c)}$ is the intertemporal elasticity of substitution. Equations (2.26) and (2.27) make precise the influences on household saving decisions. As in an economy with certainty, $r-\rho$ measures the importance of the return vis-a-vis the rate of impatience in determining saving. The two new influences are the wage schedule and the labor market frictions. An upward sloping wage schedule and a possibility of losing their job drive up saving of the employed workers (reducing current vs. future consumption, and pushing consumption growth up). The possibility of finding employment reduces the saving of the unemployed, so that a weaker labor market delivers an increase in their saving.

2.2.5.2 Law of motion for the distribution of workers

The evolution of the distribution of assets and employment states is characterized by the standard Kolmogorov Forward equations (KFEs) for the employed and the unemployed. Let $g(a, s, t)$ denote the distribution of assets over time for all workers in state $s \in \{s_e, s_u\}$ and g_t denote its time derivative. Then the KFEs are

$$g_t(a, s_e, t) = -\frac{\partial}{\partial a}[\dot{a}(a, s_e, t)g(a, s_e, t)] + f(t)g(a, s_u, t) - \sigma g(a, s_e, t) \tag{2.28}$$

for the employed and

$$g_t(a, s_u, t) = -\frac{\partial}{\partial a}[\dot{a}(a, s_u, t)g(a, s_u, t)] + \sigma g(a, s_e, t) - f(t)g(a, s_u, t) \tag{2.29}$$

for the unemployed.

2.2.5.3 Wage setting

Every worker-firm match generates a surplus that needs to be shared between the two parties. As already highlighted by the discussion above, individual asset holdings a determine the effective bargaining power of a worker, meaning that a solution to the bargaining problem under perfect information is a wage schedule $w(a)$. In what follows, I assume that the bargaining game is one of complete information: in particular, the firm knows the circumstances of a worker and in particular has full knowledge of her asset holdings. I also consider a regime in which a wage identical for all workers is set by a government or a union.

The literature has considered several wage setting protocols, which give equivalent results under risk neutrality, but not under the presence of risk-aversion (L'Haridon, Malherbet, and Pérez-Duarte (2013)). By far the most popular is Nash bargaining, which follows from the maximization of the expected present value of the job to the worker and to the employer, net of the value of searching for an alternative partner. Mathematically, wages are chosen so as to maximize the product of each partner's net return from agreement and cooperation – that is, to maximize the Nash product:

$$w(a) = \arg\max_{w} [W(w, a, s_e) - W(a, s_u)]^{\beta} [J(w, a)]^{1-\beta}, \tag{2.30}$$

where I explicitly denote the dependence of the value functions on the wage. The first order condition of this problem is:[10]

$$\frac{\beta \partial_w W(w, a, s_e)}{W(w, a, s_e) - W(a, s_u)} + \frac{(1-\beta)\partial_w J(w, a)}{J(a)} = 0. \tag{2.31}$$

The partial derivatives are $\partial_w W = \frac{\partial_a W}{\rho+\sigma}$ and $\partial_w J = \frac{-1+\partial_a J(a)\cdot\partial_w \dot{a}}{r+\sigma}$ where $\partial_w \dot{a}(a, s_e) = 1-\partial_w c_e = 1 - \partial_a c_e$ is the marginal increase in saving of the worker following a rise in the wage, and the second equality follows because a one-off \$1 increase in the wage is equivalent to an increase in personal wealth by \$1. Substituting these in (2.31) yields

$$\beta \frac{(\sigma + r) J(a) \partial_a W}{1 - \partial_a J(a)(1 - \partial_w c_e)} = (1-\beta)(\rho + \sigma)[W(a, s_e) - W(a, s_u)]. \tag{2.32}$$

Finally, using (2.12) and (2.16), we obtain a closed-form expression for the wage schedule under Nash bargaining:

$$w^{Nash}(a) = \beta \frac{\bar{p} + \partial_a J(a)[ra - c_e]}{1 - \partial_a J(a)(1 - \partial_a c_e)} - (1-\beta)\frac{u(c_e) + W_a[ra - c_e] - \rho W(a, s_u)}{W_a(a, s_u)}. \tag{2.33}$$

An alternative to Nash bargaining is the Egalitarian bargaining protocol (Kalai, 1977). In that setting, the agreed wage achieves a split of the surplus across the two parties that is proportional to their relative bargaining power:

$$\frac{W(a, s_e) - W(a, s_u)}{J(a)} = \frac{\beta}{1-\beta}. \tag{2.34}$$

By using (2.12) and (2.16) in (2.34) we derive the Egalitarian bargaining wage schedule $w(a)$ in closed form:

[10] To see this immediately, take logs of the right hand side of (2.30).

$$w^{Egalitarian}(a)\left[(1-\beta)\frac{W_a(a,s_e)}{\rho+\sigma}+\beta\frac{1-J_a(a)}{r+\sigma}\right]=\beta\frac{\bar{p}+J_a(a)[ra-c_e]}{r+\sigma}$$
$$-(1-\beta)\frac{u(c_e)+W_a(a,s_e)[ra-c_e]-\rho W(a,s_u)}{\rho+\sigma}. \quad (2.35)$$

I solve the model using both of these bargaining protocols plus the regime where wage is set at a union level. I discuss the results in more detail in subsection 2.2.8.2.

2.2.6 Asset market

Assets in positive net supply take the form of perfectly substitutable capital and equity (i.e. the rights to the stream of corporate profits[11]). The no-arbitrage condition between the two assets dictates that both have the same rate of return:

$$r(t)=\frac{d(t)+\dot{p^e}(t)}{p^e(t)}, \quad (2.36)$$

where $d(t)$ is the aggregate dividend and $p^e(t)$ is the price of equity. This equation says that the rate of return to capital, $r(t)$, is equal to the rate of return on equity. In aggregate, the dividend flow is equal to the total flow of profit across all filled jobs less the cost of recruiting:

$$d(t)=\int_{\underline{a}}^{\infty}\pi(a,t)g(a,s^e)da-\xi v(t). \quad (2.37)$$

2.2.7 Market clearing and equilibrium

In the labor market, free entry to vacancy creation translates into an equilibrium condition $V=0$, and so:

$$\frac{\xi}{q(\theta)}=\int_{\underline{a}}^{\infty}J(a)\frac{g(a,s_u)}{u}da. \quad (2.38)$$

The equilibrium in the asset market requires that asset supply (aggregate household saving) equals asset demand (firms' capital stock and value of equity):

$$\int_{\underline{a}}^{\infty}a\left[g(a,s_e,t)+g(a,s_u,t)\right]da=K(t)+p^e(t). \quad (2.39)$$

The aggregate capital stock can be calculated in the following way. Using equation (2.2) we can substitute out $y(j,t)$ in equation (2.38) to get an expression for aggregate output:

$$Y(t)=\frac{1-u(t)}{1-m^*(t)}z\bar{p}(t). \quad (2.40)$$

The capital employed by producer j (where $j\leqslant m^*(t)$) is $k(j,t)=y(j,t)\gamma(j)$. Plugging this

[11] In equilibrium manual technology producers turn a positive profit just sufficient to cover the recruiting costs.

back into the (2.2) and using (2.4) yields an expression for capital at firm j that depends only on aggregate variables:

$$k(j,t) = \frac{Y(t)}{r(t)+\delta}.$$

Aggregate capital demand is the sum of individual capital demands across all the producers using the automated technology. The capital market clearing then implies that:

$$K(t) = m^*(t)\frac{Y(t)}{r(t)+\delta}. \tag{2.41}$$

Equations (2.40) and (2.41) together give the expression for the equilibrium aggregate capital stock:

$$K(t) = \frac{m^*(t)}{1-m^*(t)}\frac{(1-u(t))\bar{p}(t)}{r(t)+\delta}. \tag{2.42}$$

The formal definition of the equilibrium is as follows:

Definition 12 *The* competitive equilibrium *is a set of value functions* $\{W(a,s,t),\ J(a,t),\ V(t),\ Q(j,t)\}$*; a set of policy functions* $\{c(a,s,t);\ \dot{a}(a,s,t)\}$*; a distribution over assets and employment* $g(a,s,t)$*; a set of prices* $\{r(t),\ p(j,t),\ w(a,t),\ p^e(t)\}$*; automation threshold* $m^*(t)$*; aggregate capital* $K(t)$*, aggregate output* $Y(t)$*, labor market tightness* $\theta(t)$ *and the unemployment rate* $u(t)$ *consistent with optimal behavior and market clearing as defined by equations (2.16)-(2.42).*

2.2.8 Stationary equilibrium

A stationary equilibrium is a competitive equilibrium in which the distribution of workers over assets and employment states is constant. Mathematically, the conditions for stationary equilibrium are obtained by setting time derivatives to zero and dropping time dependance in equations (2.16)-(2.42).

It is well-known that a model featuring uninsured income risk does not permit an analytical solution. To illustrate the workings of the model in equilibrium, I solve the model numerically. To do so, I parametrize the model and I develop a numerical solution algorithm, building on the continuous time computational methods of Achdou et al. (2017) and the extension with search and matching of Bardóczy (2017). The structure of the solution is much more complex than a standard incomplete markets model, which only searches for a scalar interest rate that equilibrates the asset market. Instead, my contribution is a fast method that solves for a general equilibrium solution over the set of admissible wage functions $w(a)$, the real interest rate r, the equilibrium technology choice m^* and equilibrium labor market tightness θ. I provide the details of the computational algorithm in Appendix 2.A.

2.2.8.1 Parametrization

I choose the following functional forms. The unit capital requirement for task j is exponential: $\gamma(j) = \kappa \cdot exp(j)$ where κ is a scalar (which can be shocked to deliver a change in capital productivity). The matching function is the standard constant-returns Cobb-Douglas:

$M(u,v) = \chi u^{\eta} v^{1-\eta}$, with χ a positive parameter, which implies that the vacancy-filling rate is $q(\theta) = \chi\theta^{-\eta}$ and the job finding rate is $f = \chi\theta^{1-\eta}$. I assume that instantaneous utility is CRRA, $u(c) = \frac{c^{1-\psi}}{1-\psi}$ with the intertemporal elasticity of substitution given by $\frac{1}{\psi}$.

Calibration of the parameters is fairly standard and corresponds to a quarterly frequency. I set the discount rate $\rho = 0.01$; the matching efficiency $\chi = 1$; separation rate $\sigma = 0.1$ (Shimer (2005)) and the risk aversion coefficient $\psi = 3$, implying intertemporal elasticity of substitution of $\frac{1}{3}$. I then set the Cobb-Douglas share of unemployment in the matching function equal to workers' bargaining power: $\eta = \beta = \frac{1}{2}$. Unemployment benefit is $h = 0.4$ targeting the replacement rate of around 40%. Capital depreciates at rate $\delta = 0.021$ or around 8% per annum. The technology parameter κ are set such that m^* is in the interior and specified below. I set the borrowing constraint at $\underline{a} = -2$ which is tighter than the natural borrowing limit.

2.2.8.2 The key features of the stationary equilibrium

Figure 2.5 illustrates the stationary equilibrium using a familiar diagram of Aiyagari (1994). On the vertical axis is the interest rate; on the horizontal axis are the aggregate assets in the economy. The dashed vertical line shows the borrowing limit $\underline{a}$. The equilibrium, labelled $EQ1$, is the point where the asset demand and the asset supply schedules cross.

Aggregate asset demand consists of two parts: automated firms' physical capital and the value of firms' equity.[12] This dual nature of asset demand means that the demand schedule has a kink at $\bar{r}$.[13]

The upward sloping line is the supply of assets. As is standard in the Bewley-Hugget-Aiyagari economies with uninsured income risk, higher interest rate makes saving cheaper and encourages asset accumulation, making the supply schedule upward sloping. It diverges to infinity at the discount rate ρ: in the limit, the abundance of assets generates plenty of self-insurance opportunities and brings the equilibrium interest rate close to the complete markets benchmark.

2.2.9 Relation to the literature

Having outlined the model, I now take a step back and provide a map between this model and some well-known models in the literature:

Definition 13 ***Benchmark A: Perfect insurance.*** *With perfect unemployment insurance the interest rate is pinned down by the discount rate: $r = \rho$. Equilibrium automation is $m^* =$*

[12]By the no-arbitrage condition (2.36), in stationary equilibrium the value of firms' equity is the annuity value of their dividend flow: $p^e = \frac{d}{r}$. This is the thin-lined hyperbola in the Figure, which asymptotes to infinity as the interest rate goes to zero. This places a natural lower bound on the interest rate in equilibrium.

[13]When interest rates are this high, no producer wishes to use the automated technology, and total asset demand is simply the value of firms' equity. The Figure also depicts $\underline{r}$, the threshold at which each producer would prefer to automate their production process. As long as this threshold is negative, however, full automation cannot be an equilibrium, because of the divergent valuation of firms' equity at $r = 0$.

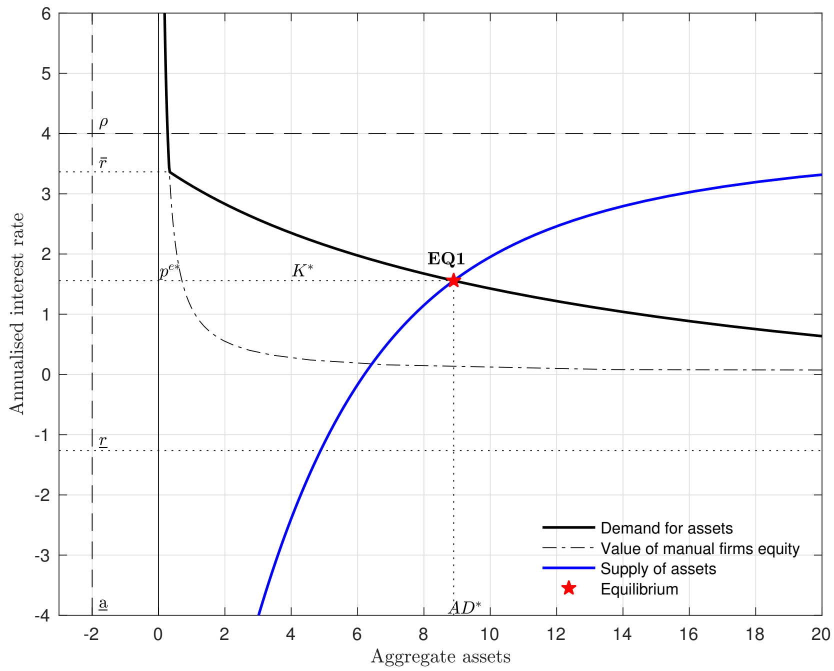

Figure 2.5: Stationary Equilibrium

$\max\{0, \min\{\mu, \hat{m}\}\}$ *where* $\hat{m}$ *solves*

$$\rho + \delta = \exp\left[(x-1)\ln\gamma(x) - \int_0^x \ln\gamma(j)dj \right].$$

Aggregate output is

$$Y = B\left(\frac{K}{m*}\right)^{m*}\left(\frac{L}{1-m*}\right)^{1-m*} \tag{2.43}$$

where $B \equiv exp\left(-\int_0^{m*} \ln\gamma(j)dj\right)$. *The labor market is described by the basic textbook model of search and matching, such as the one in Pissarides (1990).*

Definition 14 ***Benchmark B: Frictionless labor market.*** *With no frictions in the labor market – when vacancy posting is costless and when matching is infinitely efficient – the job finding rate is infinite and unemployment rate is zero, implying that there is no idiosyncractic risk. The economy behaves as-if the markets were complete. There is a representative agent who solves*

$$\rho W(K) = \max_C u(C) + W'(K)\dot{K}$$

subject to

$$\dot{K} = Y - C - \delta K.$$

The equilibrium interest rate is $r = \rho$. *Automation is determined as in the previous Proposition.*

Definition 15 ***Benchmark C: No automation possible.*** *If* $\gamma(0) \geqslant \frac{1}{\rho+\delta}$ *or if* $\mu = 0$, *no firm operates the automated technology in equilibrium. With a borrowing limit that is sufficiently*

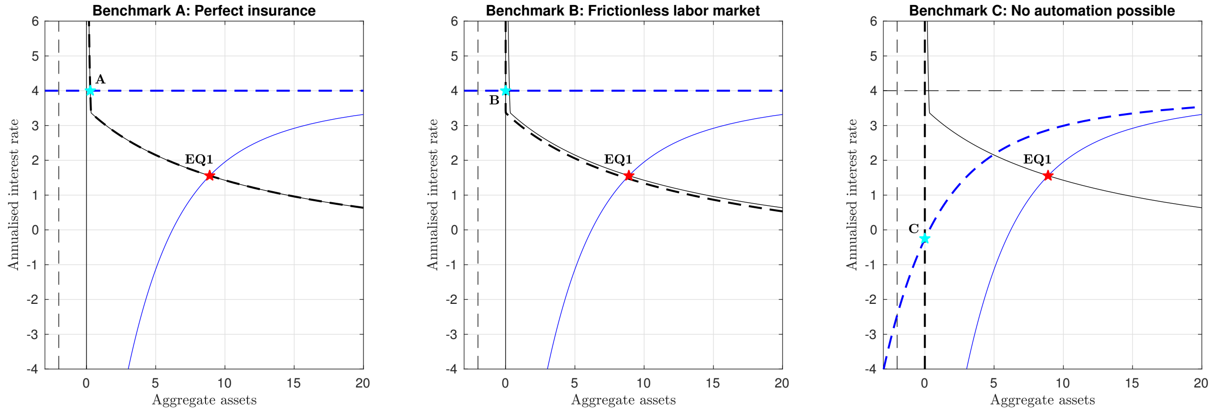

Figure 2.6: The stationary equilibrium vs. three benchmark economies

loose, the model economy resembles a Huggett / Bewley economy. If, in addition, conditions in Proposition 14 hold, then the model is identical to an endowment economy with no uncertainty.

These Propositions place the present model within the context of some of the workhorse frameworks of modern macroeconomics. Proposition 13 shows that without market incompleteness, the model collapses to that of Zeira (1998) with a canonical DMP structure in the labor market.

Proposition 14 illustrates that two of the model ingredients – incomplete markets and search frictions – are intricately linked. Without frictions, the labor market clears, unemployment rate is zero and individual workers face no uncertainty. The model collapses to the simple economy similar to that analyzed by Zeira (1998) and Acemoglu and Restrepo (2018).[14]

Proposition 15 shows that when automation is inefficient and/or prohibitively expensive there will be no automation and hence no capital in equilibrium and the model is simply a search and matching economy with incomplete markets. When, in addition, there are no frictions in the labor market, the economy is a simple endowment economy with no uncertainty.

Figure 2.6 contains stylized comparisons of the equilibrium in the model of this paper with the three benchmark economies outlined above. In all of the three panels, the thin lines and the point marked $EQ1$ represent the model in this paper, as in Figure 2.5. The left panel shows the determination of the equilibrium if perfect insurance was available. In that case, asset supply is perfectly elastic at $r = \rho$ (dashed horizontal line). Asset demand is the same as in the model economy[15] and the equilibrium is labelled A.

The middle panel shows the equilibrium under the frictionless labor market. In this case, dividends and value of manual technology firms are both zero, and asset demand coincides with demand for physical machines. Given zero unemployment, households face no uncertainty.

The right panel shows the equilibrium when no automation technologies are feasible (either because none have been invented, so that $m = 0$, or because they are very expensive, so that $\gamma(0)$ is large). With no capital in the model, the only potential asset in positive net supply is

[14] Similar result holds in a stationary equilibrium when the separation rate σ is zero.

[15] This similarity is approximate. Asset demand will depend on other endogenous variables, such as wages, that will be different across the two equilibria. In all illustrations here, both demand and supply schedules of the main model and benchmark economies are drawn for a given value of wages (equal to their value in equilibrium $EQ1$).

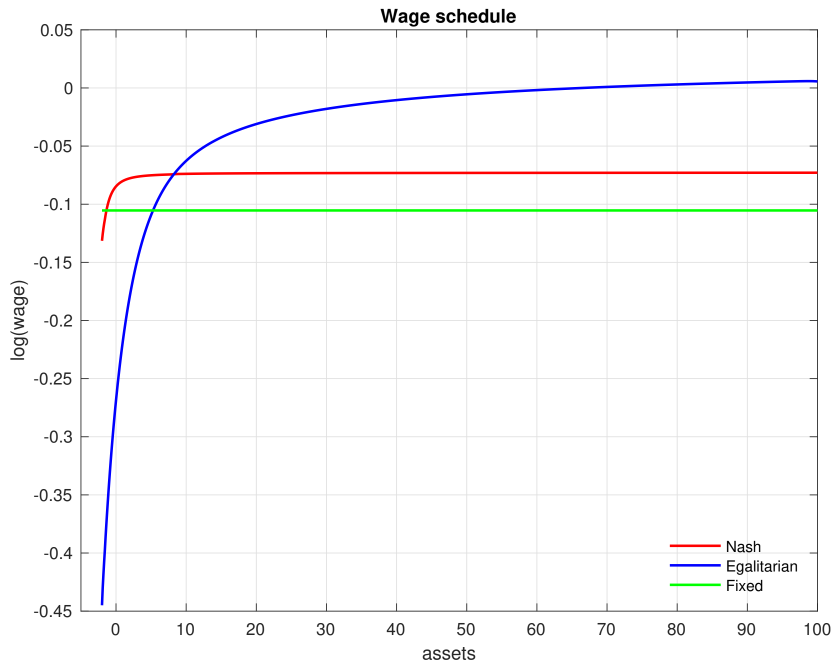

Figure 2.7: Wage schedules under alternative bargaining protocols

firms' equity. . The economy then corresponds to a Huggett (1993) economy, with vertical asset demand.

2.2.10 The wage schedule under alternative bargaining protocols

As flagged above, there are several ways to close the wage setting side of the model. One contribution of this paper which I now turn to is to provide some insight into the structure of the equilibrium under the different wage bargaining protocols.

Figure 2.7 compares the wage schedules in the stationary equilibrium of this economy computed under three bargaining regimes: the Nash and Egalitarian regimes, and the fixed wage regime in which an economy-wide wage is set by a government or a union.

As discussed in Krusell, Mukoyama, and Sahin (2010), under Nash bargaining the wage schedule is upward sloping only in the region close to the borrowing constraint, with the wage fast-approaching the wage that would have been obtained under risk-neutrality. This is no longer the case under Egalitarian regime, in which the wage schedule is significantly upward sloping throughout the asset spectrum. The underlying reason for this difference is that the two regimes differ in their scale invariance property (L'Haridon, Malherbet, and Pérez-Duarte (2013)). Intuitively, Nash bargaining scales worker's surplus from a match by her marginal utility of consumption. Higher assets raise reservation utility, requiring higher wages to restore the bargaining equilibrium. But they also raise consumption and thus lower marginal utility, increasing the scaled surplus and making the required wage adjustment smaller. As a result, the wage schedule is practically flat on much of the asset spectrum. Under Egalitarian bargaining, worker's surplus is not scaled, and to restore the equilibrium in the bargaining game following an increase in worker's assets, wages have to increase, even for high asset holdings.

The best way to choose between the alternative parametrizations of the bargaining protocol is to match the equilibrium wage schedule to that in the data. While detailed analysis of this question is beyond the scope of this paper at this stage, Appendix 2.B provide some preliminary indicative evidence of the existence of wage schedule in the data using the United States Survey of Consumer Finances. Of course, such analysis is plagued by endogeneity of assets to wages: most obviously, high skill workers will tend to earn more and accumulate higher assets. This is an obvious channel confounding the mechanism in my theory, which is that the level of individual savings drive variation in wages of otherwise identical workers.

There are other strands of literature that suggest that the link is there, however. A series of studies have found a significant effects of individual wealth on unemployment duration and reservation wages (Stancanelli (1999), Lentz (2009), Bloemen and Stancanelli (2002), Alexopoulos and Gladden (2006), Lammers (2014)). In an recent study, Krueger and Mueller (2016) use a survey of workers in New Jersey and show that reservation wages increase in individual savings. There is also some nascent evidence on the causal impact of wealth on wages using natural experiment-type variation. Giupponi (2018) considers a natural experiment provided by the Italian pension reform, and reports that a decline in wealth of EUR1,000 has led to a decline in wages by around EUR2.2, or 3%, relative to the control group.

While far from definitive, this evidence suggests that the link between assets and wages identified by the theoretical model is a plausible one. Nonetheless, pinning down the exact mechanisms and calibrating the magnitudes requires further empirical work. In future work I plan to explore this issue further through both empirical investigation and more detailed theoretical modeling of the relationship between assets and earnings. In the remainder of the paper I take the link between wages and assets as given and explore some of the implications.

2.3 Technology and labor demand

The model outlined above is a rich laboratory for the study of the nexus of automation, unemployment and inequality. The key benefit of this richness is that it allows for an exploration of some novel channels and interesting macro feedback loops through which technology affects the macroeconomy at large.

This richness comes at the cost, however: interpretation of the multilayered forces and interactions in general equilibrium necessitates some careful thought. To develop the intuition, this Section discusses the determinants of labor demand in the model.

2.3.1 Four forces that drive job creation

In stationary equilibrium, job creation is pinned down by the free entry condition to vacancy creation, reproduced here from equation (2.15):

$$\frac{\xi}{q(\theta)} = \int_{\underline{a}}^{\infty} J(a)\frac{g(a, s_u)}{u} da. \tag{2.44}$$

A firm that considers posting a vacancy does not know in advance how wealthy is the worker it is going to match with; instead it makes its vacancy posting decision on the basis of the expectation of the value of the job given the distribution of assets across the unemployed workers in the economy. This expectation is the probability weighted sum of possible job values $J(a)$, where the relevant probabilities are given by the density function of the distribution of potential hires, $\frac{g(a,s_u)}{u}$. So there are two factors that determine the labor market response: the value function $J(a)$ and the distribution of the unemployed $g(a, s_u)$.

To unpack the determinants of $J(a)$, note that the stationary version of equation (2.14) is

$$(r + \sigma)J(a) = \bar{p} - w(a) + J_a(a)\dot{a}(a, s_e; r), \tag{2.45}$$

where $\dot{a}(a, s_e; r)$ is the optimal saving policy function of the employed agents. Equation (2.45) is a PDE, which can be approximated by discretizing the asset spectrum as follows:[16]

$$(r + \sigma)J(a_n) = \bar{p} - w(a) + \frac{J(a_{n+1}) - J(a_n)}{da}\dot{a}(a, s_e; r), \tag{2.46}$$

where a_n is an n-th element on a discrete asset grid. Writing this in matrix notation we get:

$$(r + \sigma)\mathbf{J}(\mathbf{a}) = \bar{p} - \mathbf{w}(\mathbf{a}) + \mathbf{A}(s_e; r, w(a))\mathbf{J}(\mathbf{a}) \tag{2.47}$$

where $\mathbf{A}(s_e; r, w(a))$ is the matrix of saving intensities of the employed. Thus the expression that characterizes the (vector of) value of filled jobs is:

$$\mathbf{J}(\mathbf{a}) = [(r + \sigma)\mathrm{I} - \mathbf{A}(s_e; r, w(a))]^{-1}(\bar{p} - \mathbf{w}(\mathbf{a})). \tag{2.48}$$

Furthermore, note that the Kolmogorov Forward Equation equation in the steady state is:

$$\mathbf{0} = \mathbf{A}(s_u; r, w(a))^{\mathrm{T}}\mathbf{g}_u(a). \tag{2.49}$$

That is, in the steady state, the distribution of the unemployed is the zero eigenvector of the transpose of the transition matrix of the unemployed workers $\mathbf{A}(s_u, r, w(a))$. Economically and intuitively, this highlights that the distribution and the optimal saving decisions are intricately linked: naturally, the equilibrium distribution is a result of the optimal saving behavior.

Together, equations (2.48) and (2.49) bring out the four forces that matter for job creation:

1. price of manual tasks $\bar{p}$;
2. wages;
3. the interest rate;
4. saving behavior.

[16]This is a "finite difference method", whereby the continuum of assets is discretized on a grid and the derivative is approximated as the slope of the $J(\cdot)$ function between the grid points. The forward difference is used here because employed workers always save a positive amount in equilibrium.

2.3.2 Automation and $\bar{p}$

The key variable one wants to look at following an automation shock is the price of goods and services produced by human workers, $\bar{p}$. It informs us about what happens to the demand for those goods. As I show now, it is possible to characterize this response analytically, and the response can be rather different depending on whether the economy is in a constrained or an unconstrained equilibrium.

2.3.2.1 Unconstrained equilibrium and a higher adoption of existing technologies

When technological constraint is slack $(m^* < \mu)$, any reduction in the cost of using machines will translate into higher adoption on the margin. The following Proposition shows that, irrespectively of the source of this change, changes in equilibrium adoption and equilibrium price of manual tasks have the same sign:

Definition 16 ***Price of manual tasks when technological constraint is slack:*** *Suppose that in equilibrium not all existing automation technologies are adopted: $m^* < \mu$. Then the equilibrium price of manually produced goods $\bar{p}$ is increasing in the (endogenous) automation threshold m^*. Any change in the environment that causes a marginal increase in m^* raises $\ln \bar{p}$ by the value of the elasticity of $\gamma(j)$ evaluated at m^*.*

Proof. Dropping time subscripts, Equations (2.8) implies that

$$\ln \bar{p} = m^* \ln \gamma(m^*) - \int_0^{m^*} \ln \gamma(j) dj.$$

Thus the derivative of the log price is:

$$\frac{\partial \ln \bar{p}}{\partial m^*} = \frac{m^*}{\gamma(m^*)} \gamma'(m^*) = \epsilon_{\gamma(j)} \mid_{j=m^*},$$

where $\varepsilon_{\gamma(j)} \mid_{j=m^*}$ is the notation for the elasticity of $\gamma(j)$ evaluated at the equilibrium automation threshold. Assumption 1 ensures that $\gamma(j)$ is non-negative and increasing, implying that $\varepsilon_{\gamma(j)} > 0$. ■ In an unconstrained equilibrium, higher adoption always raises the price of tasks produced by labor and, all else equal, increase labor demand. The intuition is that further automation technologies are adopted only if they become relatively cheaper, so that manual tasks must become (relatively) more expensive. Figure 2.8 illustrates this diagrammatically.[17] The upward sloping solid curve is the cost of producing a given task j with the automated technology. The initial equilibrium is $(m^*, \bar{p})$, with $(r + \delta)\gamma(m^*) = \bar{p}$ (Proposition 11). A change in the economic environment such that machines become cheaper – e.g. a fall in the depreciation rate δ or an exogenous shift in the discount rate ρ – shifts the cost of automation schedule down, and the equilibrium automation rises to $m^{*'}$. The ideal price condition means that the price $\bar{p}$ adjusts to make area A and area B equal in size.

[17] The Figure is drawn for a simple case where the technological constraint is $\mu = 1$.

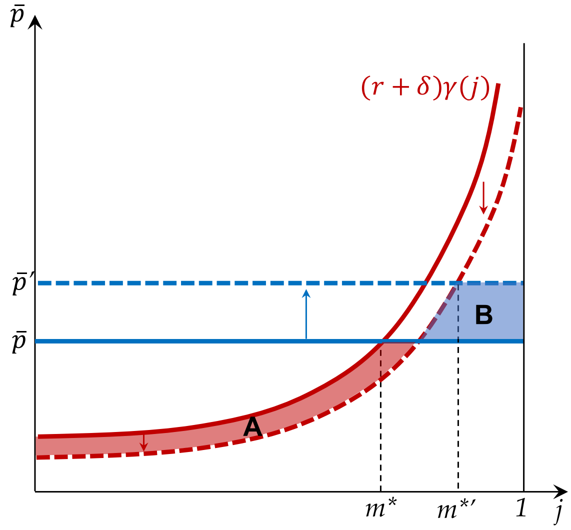

Figure 2.8: Price of manual goods $\bar{p}$ in an unconstrained equilibrium.

2.3.2.2 Constrained equilibrium: the invention of new automation technologies

I now turn to the case where all existing automation technologies are adopted: $m^* = \mu$. The price schedule is no longer continuous; instead, $p(\mu) < \bar{p}$. The ideal price condition (2.3) becomes:

$$\ln \bar{p} = \frac{1}{\mu - 1}\left(\int_0^\mu \ln[(r(\mu) + \delta)\gamma(j)]dj\right). \tag{2.50}$$

In this equilibrium the only relevant "automation shock" is the change in the automation technology frontier, μ. What happens to the price of manually produced goods as technology moves forward? The following Proposition characterizes the behavior of $\bar{p}$ in the short-run:

Definition 17 ***Price of manual tasks in the short run when technological constraint is binding.*** *Suppose that the technological constraint is binding in equilibrium, so that $m^* = \mu$. Then the short-run effect of an increase in μ (holding K and L fixed) is ambiguous. It consists of two opposing effects: the productivity effect and the displacement effect:*

$$\frac{d\ln\bar{p}}{d\mu}\,|_{K,L} = \underbrace{\ln\left(\frac{\bar{p}}{(r+\delta)\gamma(\mu)}\right)}_{productivity\ effect>0} + \underbrace{\frac{1}{\mu-1}}_{displacement\ effect<0}$$

Proof. The labor market clearing implies that, in the technologically constrained equilibrium,

$$\ln\bar{p} = \ln(1-\mu) + \ln Y - \ln L.$$

For fixed level of K and L, we have:

$$\frac{d\ln\bar{p}}{d\mu}\,|_{K,L} = \frac{d\ln Y}{d\mu}\,|_{K,L} + \frac{1}{\mu-1}. \tag{2.51}$$

The first term is the productivity effect, and it remains to be proven that this is always positive. Note that individual-firm capital and labor demands imply the aggregate capital and labor demands:

$$k(j) = \frac{Y}{p(j)}\gamma(j) = \frac{Y}{r+\delta} \Longrightarrow K = \int_0^\mu k(j)dj = \mu \frac{Y}{r+\delta}. \tag{2.52}$$

$$l(j') = \frac{Y}{\bar{p}} \Longrightarrow L = \int_\mu^1 l(j')dj' = (1-\mu)\frac{Y}{\bar{p}} \tag{2.53}$$

Using these equations to substitute out factor prices in (2.50), we have:

$$Y = B\left(\frac{K}{\mu}\right)^\mu \left(\frac{L}{1-\mu}\right)^{1-\mu} \tag{2.54}$$

where $B \equiv exp\left(-\left[\int_0^\mu \ln\gamma(j)dj\right]\right)$. Consider the log-transformation of the aggregate production function:

$$\ln Y = -\left[\int_0^\mu \ln(\gamma(j))dj\right] + \mu\ln K - \mu\ln\mu + (1-\mu)\ln L - (1-\mu)\ln(1-\mu).$$

Differentiating and rearranging gives:

$$\frac{d\ln Y}{d\mu} \mid_{K,L} = \ln\left(\frac{1-\mu}{\mu}\frac{K}{L}\frac{\bar{l}}{\gamma(\mu)}\right) = \ln\left(\frac{\bar{p}}{(r+\delta)\gamma(\mu)}\right) \tag{2.55}$$

where the second equality follows from $\frac{K}{L} = \frac{\mu}{1-\mu}\frac{\bar{p}}{r+\delta}$. Because in the technologically constrained equilibrium $\bar{p} \geqslant (r+\delta)\gamma(\mu)$, $\frac{d\ln Y}{d\mu}$ is always positive. ∎ The productivity effect has been discussed in the recent literature in the context of technological change (Acemoglu and Restrepo, 2018) and offshoring (Grossman and Rossi-Hansberg, 2008). The intuition is that the invention of automated technologies presents an opportunity to enhance the efficiency of the production process, which translates into higher overall demand, including for the manually produced goods.

Counteracting the positive productivity effect is the displacement effect: workers scramble for jobs that are less plentiful. This effect is always negative, and becomes stronger as automation progresses: in the limit, as $\mu \to 1$ and there are very few jobs left, the price of manually produced tasks falls unboundedly.

The long-run effect of a higher μ is ambiguous. It depends on whether technological automation brings the overall cost of production of tasks down or not. If the cost of producing automated tasks in the new equilibrium is lower, then the price of manual tasks increases and vice-versa. The overall change in cost comprises of two effects First, the newly automated tasks are produced using cheaper technologies, and so prices of tasks between μ and μ' unambiguously decline. On the other hand, capital accumulation and higher capital demand is associated with a *higher* interest rate in equilibrium in an economy with incomplete markets. If the cost saving effect dominates, this raises the relative price of the tasks produced by labor. Otherwise, technological automation on the whole raises the cost of producing with capital, and the relative price of labor tasks falls.

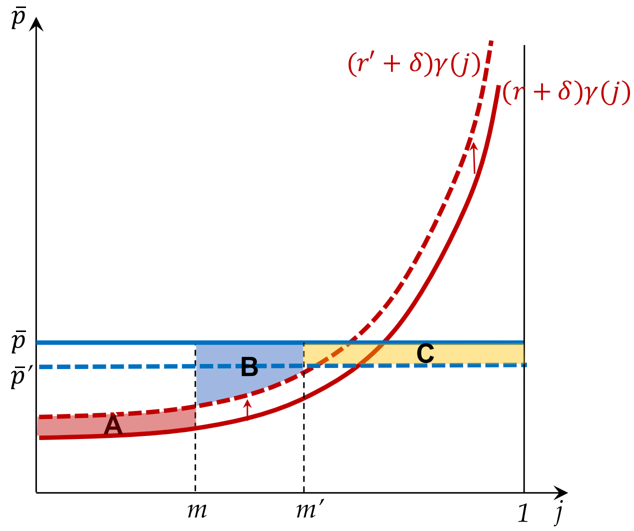

Figure 2.9: Price of manual goods $\bar{p}$ in a constrained equilibrium.

Figure 2.9 illustrates this graphically. The upward sloping line traces out the cost of producing with automation technologies. This curve shifts upwards in the long run following the technological innovation, due to the increase in the interest rate. Tasks in the range $[0, \mu]$ are now more expensive. The newly automated tasks $(\mu, \mu']$ are cheaper. The net effect is the difference of the two areas labelled A and B. The ideal price condition requires that in the long-run $\bar{p}$ adjusts so that the overall price index is unchanged, that is, area C compensates for any difference in areas A and B.

Which effect dominates depends on the magnitudes of the increase in the interest rate, the change in technological constraint μ, and the initial distance between the price of labor tasks and automation costs.[18]

2.3.3 Automation and wages

The previous discussion showed that $\bar{p}$ may go up or down with automation, depending on the source of the change that brings it about. I now discuss what happens to wages after the change in this price, focusing on a single firm and a single worker. This partial equilibrium analysis is useful to build intuition of the channels at play.

For concreteness, suppose that $\bar{p}$ increases following the shock. The direct effect of this change is to raise the value of manual jobs $J(a)$ as the profit of firms operating the manual technology $\pi(a) = \bar{p} - w(a)$ goes up. This first-round effect is the same for all firms irrespective of the assets held by their workers. This additional surplus is shared through bargaining. The following Proposition characterizes the outcome under Egalitarian bargaining:

Definition 18 *Wages of asset rich respond by more than wages of asset poor to the changes in the price of manual goods $\bar{p}$.*

[18] A small increases in μ are likely to raise the overall costs of production due to the concavity of the asset supply schedule: the interest rate increases early on in the automation process are likely to be the largest, while the cost savings are likely to be more modest. This intuition is confirmed in numerical simulations that follow.

Proof. From 2.35, the derivative of the wage with respect to $\bar{p}$ is:

$$\frac{\partial w(a)}{\partial \bar{p}} = \frac{\beta}{r+\sigma}\left[(1-\beta)\frac{W_a(a,s_e)}{\rho+\sigma} + \beta\frac{1-J_a(a)}{r+\sigma}\right]^{-1}. \tag{2.56}$$

Because W is concave, W_a is decreasing in a. Concavity of W also implies that J is decreasing and convex,[19] so that $1-J_a$ is declining towards 1 as $a \to \infty$. Thus the right-hand side of (2.56) is increasing in a: the wealthy workers get a bigger pay increase. ∎ Therefore the uniform "first-round" increase in $J(a)$ translates into a non-uniform increase in wages. The economic intuition for this result is that of a variable sensitivity of firm's and worker's surpluses to the wage: both surpluses are more sensitive to wages at low levels of assets. Consequently, a smaller increase in wages is required to restore equilibrium in the bargaining game for a worker with less assets. I explore these effects quantitatively in more detail below.

2.3.4 Automation and the interest rate

In an economy with incomplete markets, capital accumulation associated with higher automation drives up the interest rate both in the short-run and in the long-run. The short-run response is standard: supply of capital is inelastic, and higher demand shows up as an increase in the price of capital. In the long-run, capital is accumulated, but the interest rate nonetheless settles at a higher level to ensure that capital demand equals supply: with uninsurable income risk, the long-run capital supply is upward sloping.

In Moll, Rachel, and Restrepo (forthcoming) we show that for any constant returns to scale production function F, it is possible to decompose the gains from technology in the following way:

$$\underbrace{\frac{F_\theta d\theta}{F}}_{\text{TFP gains} > 0} = s_K d\ln R + \underbrace{\int_a s_a d\ln w_a}_{\text{change in average wage} \lesseqgtr 0}.$$

That is, technological gains accrue to capital and labor. In the standard complete markets model, the return to capital is pinned down by the representative household's rate of time preference, ρ. In other words, the long-run capital supply is perfectly elastic. Following a shock, capital is accumulated to the point where the return to capital is unchanged. Such capital deepening raises wages, and the above decomposition shows that in such a framework all of the TFP gains will accrue to the workers (the inelastic factor).

In the model with incomplete markets, capital supply is upward sloping, and technology that increases demand for capital will raise the interest rate not only in the short-run but also in the long-run. Higher return to capital and less capital deepening, relative to the standard model, will mean that wages increase by less. This combination of factor prices will tend to drive strong job creation in equilibrium.

[19] This is because the wage schedule is increasing and concave, and the job value is a mirror-image of the wage schedule.

2.3.5 Automation and saving behavior

The discussion so far points to three interrelated reasons for why automation affects households desired saving. First, a higher interest rate makes saving more attractive. Second, a steeper wage schedule encourages saving, shifting the asset supply curve to the right. Third, changes in labor demand that emerge in equilibrium affect the riskiness of the household income process.

2.4 The effects of increased adoption of existing technologies

Armed with this intuition developed in the previous section, in this and in the next section I explore macroeconomic adjustment to technology shocks numerically. I start with the unconstrained equilibrium and consider a one-off unexpected shock to productivity of all automation technologies (the next section studies the shift in μ in a constrained equilibrium case). In what follows, I assume Egalitarian bargaining protocol is in place.

2.4.1 Comparison across the two stationary equilibria

2.4.1.1 Equilibrium prices of tasks

Specifically, the technological shift that I consider is the decline in κ (recall that $\gamma(j) = \kappa \exp(j)$). The fall in κ means that producing using the automated technology becomes cheaper across all tasks.

The left panel of Figure 2.10 illustrates the experiment by comparing the two γ functions, and the right hand panel illustrates the equilibrium prices, marked with asterisks. The equilibrium automation threshold is the task where the pricing schedule has a kink. In the initial equilibrium it is equal to $m_0^* = 18\%$.

The technology shock shifts the upward sloping schedule down. In line with Proposition 16, the price of manual tasks goes up, visible in the shift up of the dotted horizontal line until the ideal price condition is restored. The new equilibrium features a higher automation adoption: at $m_1^* = 31\%$, nearly a third of all tasks are automated.

2.4.1.2 Equilibrium in the asset market

Figure 2.11 illustrates the long-run equilibrium in the asset market. With incomplete markets, the 'first round' effect of the shock is to shift the equilibrium from $EQ1$ to point D:[20] aggregate assets rise along the household saving supply curve. But this is not the end of the story: in this model, the asset supply curve shifts to the right due to the "saving for higher wages" and "risk-mitigation" effects. The new equilibrium, labelled $EQ2$, can be found at the intersection of the two dashed curves.

For the remainder of this Section, I look under the bonnet of these shifts, starting with the labor market.

[20] Points A and C in the Figure mark the hypothetical initial and final equilibria under complete markets.

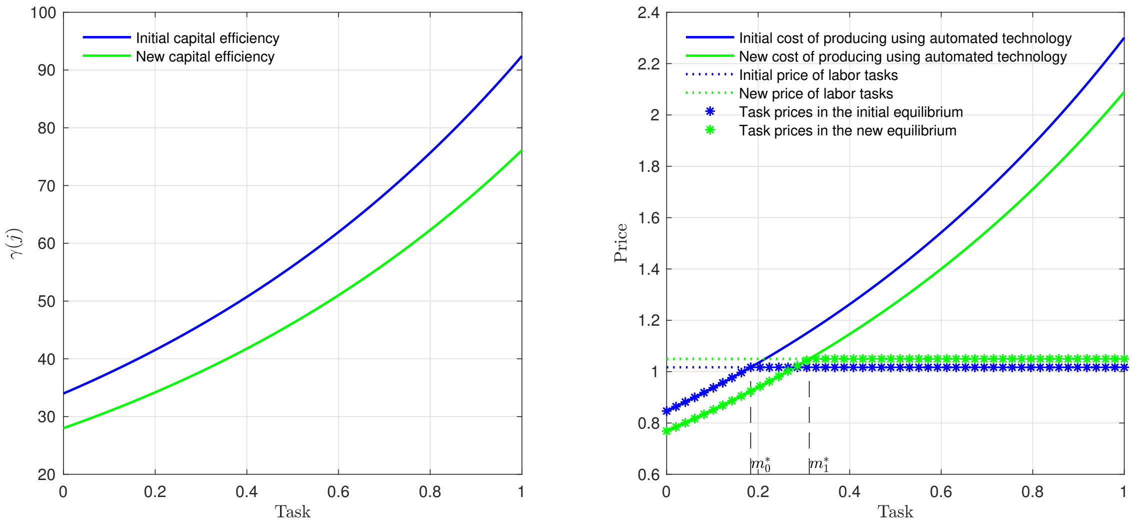

Figure 2.10: Prices of tasks in the new equilibria

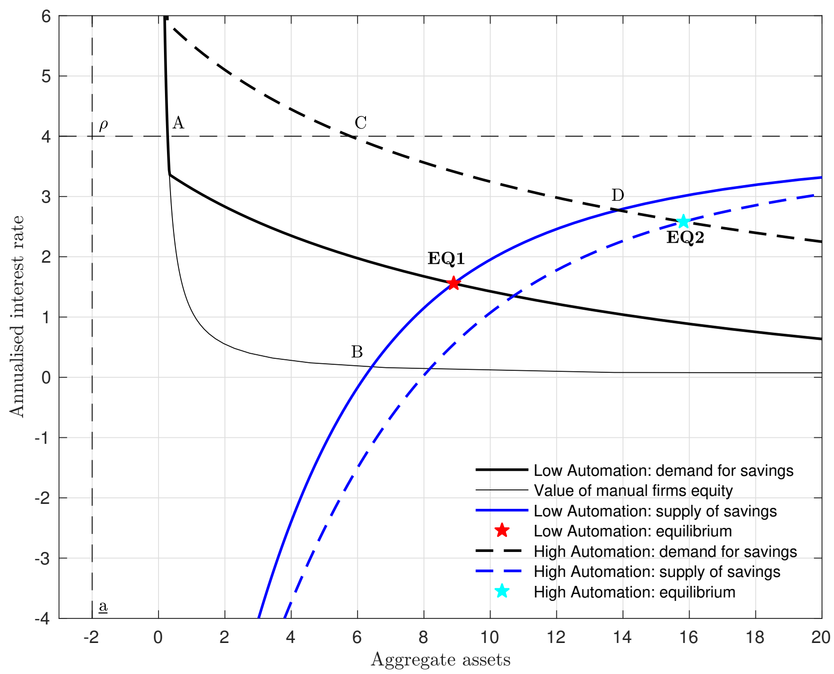

Figure 2.11: Determination of the new stationary equilibrium

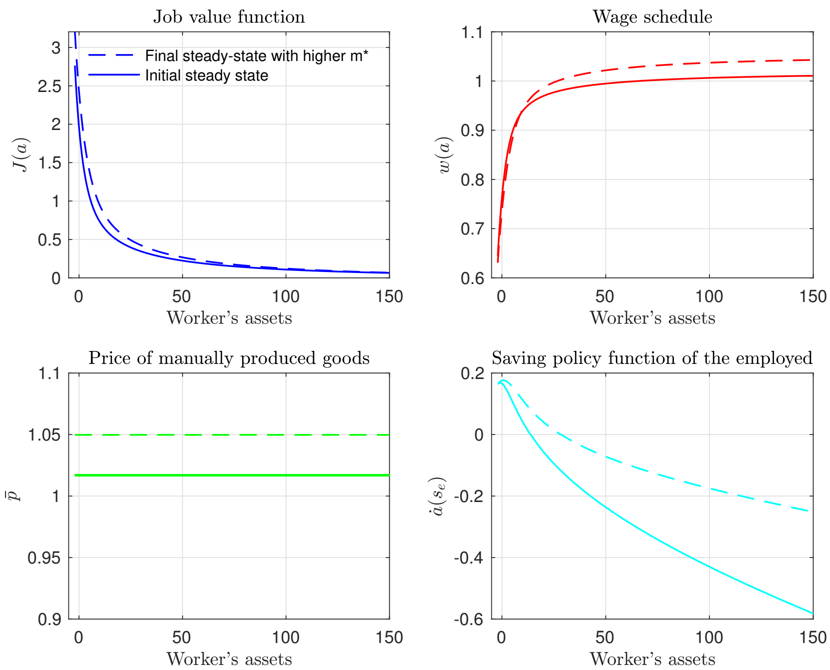

Figure 2.12: Value of the job and underlying components across the two equilibria

2.4.1.3 Labor demand

Under the present parametrization of the matching function, labor market tightness can be calculated explicitly from equation (2.44) as:

$$\theta = \left(\frac{\xi}{\chi} \int_{\underline{a}}^{\infty} J(a) \frac{g(a, s_u)}{u} da \right)^{\frac{1}{\eta}}.$$

As discussed in detail in the Section 2.3, tightness is a positive function of the weighted average of the job value $J(a)$, where the weights are given by the density of the unemployed $\frac{g(a,s_u)}{u}$.

From an employer's perspective, value of a match is higher in the new equilibrium, irrespectively of worker's individual asset holdings (top left panel of Figure 2.12). This is the net result of higher $\bar{p}$ (lower left) and changes in wages (top right). Wage schedule is visibly steeper. Finally, the saving policy function has shifted driven by higher interest rates, steeper wage schedule and higher unemployment risk (lower right).

Figure 2.13 uses equation (2.48) to decompose the increase in the job value (the difference between the dashed and solid lines in the first panel of Figure 2.12) into the various components. This decomposition indicates that there are two forces that are first order: (i) the rising prices of manually produced tasks; and (ii) wages, which fall for the poor and rise for the middle- and rich workers. Changes in the interest rate and the saving policy function of the employed are less important. Overall, $J(a)$ is higher everywhere, boosting vacancy posting and job creation. Had this been the end of the story, unemployment would decline.

But the story does not end here. Figure 2.14 compares the density $\frac{g(a,s_u)}{u}$ across the two stationary equilibria. The stationary distribution of the unemployed is shifted far to the right following the shock. In the aggregate, there is more capital in the new equilibrium, and thus

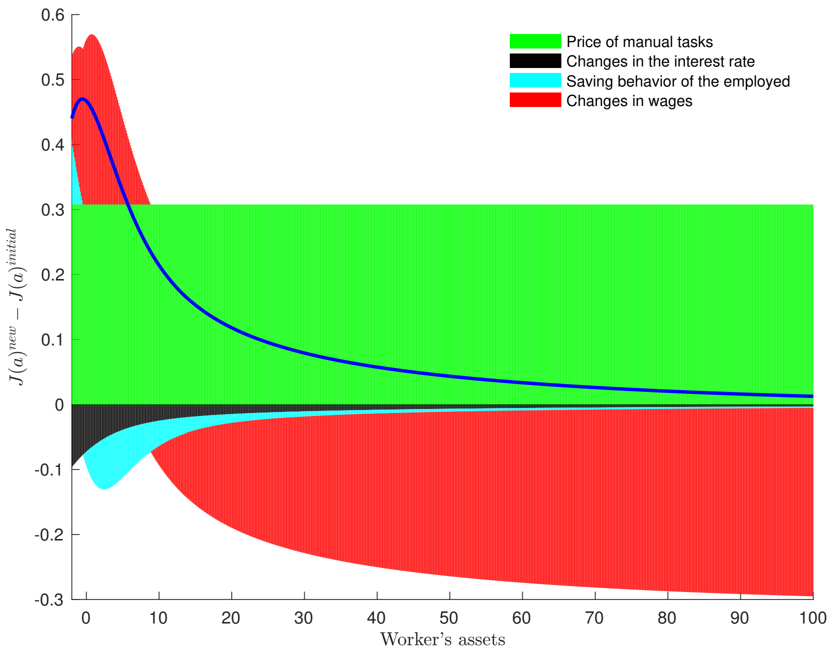

Figure 2.13: Decomposition of the change in the job value function between the two stationary equilibria

unemployed workers hold more of it. Because $J(a)$ is downward sloping in individual assets (firms are better off hiring an asset poor worker), the rightward shift of the distribution decreases the expected value of the job $J(a)\frac{g_u}{u}$. Thus there are two forces which work in opposing directions: for each level of worker's assets, value of the job is higher, but the shift of the entire distribution to the right means that when forming expectations about this value, firms attach higher weight to the region where the job values are lower. These results are summarized in the following Proposition:

Definition 19 *In general equilibrium, the rise in capital productivity has two offsetting effects on the steady state labor demand. First, for each level of individual asset holdings, higher price of manually produced tasks is only partially offset by rising wages, raising labor demand. Second, capital accumulation implies a shift of the density of the job-seekers and workers to the right, resulting in higher average expected wage, lowering labor demand.*

Table 1 compares the pre- and post-shock stationary equilibria in terms of the key variables. It reveals that, in this calibration, the latter effect dominates. The expected value of a match declines by around 12%, which is associated with a decrease in labor market tightness and an increase in the unemployment rate.

2.4.1.4 Dispersion of wages

As was already apparent from Figure 2.11, the rise in equilibrium automation leads to a *steepening of the wage schedule*, whereby asset poor workers receive a pay cut (or a smaller pay increase), while the asset-rich workers enjoy a pay increase (or a smaller pay decrease). To see the reasons behind this asymmetric response, recall that the equilibrium wage brings into

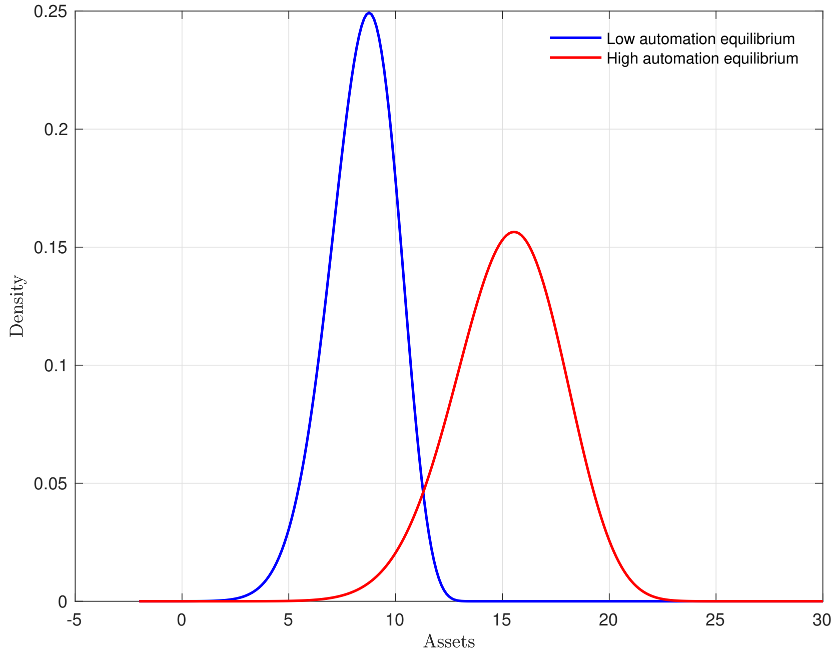

Figure 2.14: Distributions of the unemployed across the asset spectrum in the two equilibria

	Initial Eq.	Final Eq.
m^*	0.18	0.31
K	8.19	15.21
r	1.56	2.58
u	0.11	0.12
θ	0.68	0.52
f	0.82	0.72

Table 2.1: Stationary equilibria pre- and post-shock

balance worker's and firm's surplus (equation (2.35)). In stationary equilibrium, these surpluses are given by:

$$W(a, s_e) - W(a, s_u) = \frac{u(c_e) - u(c_u) + (u'(c_e) - u'(c_u))\, ra + u'(c_e)\,[w(a) - c_e] - u'(c_u)\,[h - c_e]}{\rho + f + \sigma}. \tag{2.57}$$

$$J(a) = \frac{\bar{p} - w(a) + \partial_a J(a)[w(a) + ra - c^e]}{r + \sigma}. \tag{2.58}$$

These expressions readily give the sensitivities of the surpluses to the wage. A \$1 increase in the wage raises worker's surplus by

$$\frac{u'(c_e)}{\rho + f + \sigma} \tag{2.59}$$

and lowers firm's surplus by

$$\frac{\partial_a J(a) - 1}{r + \sigma}. \tag{2.60}$$

Because worker's consumption policy function is non-decreasing in assets, and the utility function is concave, the first of these 'multipliers' (2.59) is decreasing in worker's asset holdings. This means that a given change in the wage changes worker's surplus by a larger amount when the worker is asset poor. Furthermore, because the wage schedule is upward sloping, the absolute value of the second multiplier (2.60) is decreasing in a as well. Both of these facts mean that wage is a more powerful equilibrating force at low level of assets.

As discussed in the previous subsection, an increase in automation leads to two changes in the stationary equilibrium: the labor market is weaker ($f \downarrow$), and the interest rate is higher ($r \uparrow$). Expression (2.57) clarifies how worker's surplus depends on f and r both directly, and also indirectly, through changes in c^e and c^u, and the right panel of Figure 2.15 illustrates this quantitatively. The direct effect of a fall in f is to raise the surplus: having a job becomes more valuable when the labor market is weaker. This effect is stronger for the asset-poor workers who rely on labor income to a greater extent.[21] The effect of a higher interest rate goes the other way: it lowers the value of having a job, as long as a worker has positive assets. Intuitively, a higher r means higher asset returns, making workers less dependent on wage income.[22] Quite naturally, this effect disappears for those with zero assets, and reverses for individuals with negative net worth. Finally, changes in the economic environment – a lower f and a higher r – mean that workers adjust their behavior. This adjustment is largest for asset-poor, unemployed workers.[23] The effect on the job value (the right panel) is dominated by the changes in $\bar{p}$.

These changes in surpluses and the non-linear sensitivity of the surpluses to wages translate into a steepening in the wage schedule. The three panels in Figure 2.16 show how this works. The left-most panel collates the effects from Figure 2.15, inverting the forces acting on worker's

[21] Mathematically, the numerator of the right-hand side of equation (2.57) is a positive and decreasing in a.
[22] Mathematically, $u'(c^e) - u'(c^u)$ is negative everywhere.
[23] Mathematically, $u(c^u)$ falls short of $u(c^e)$.

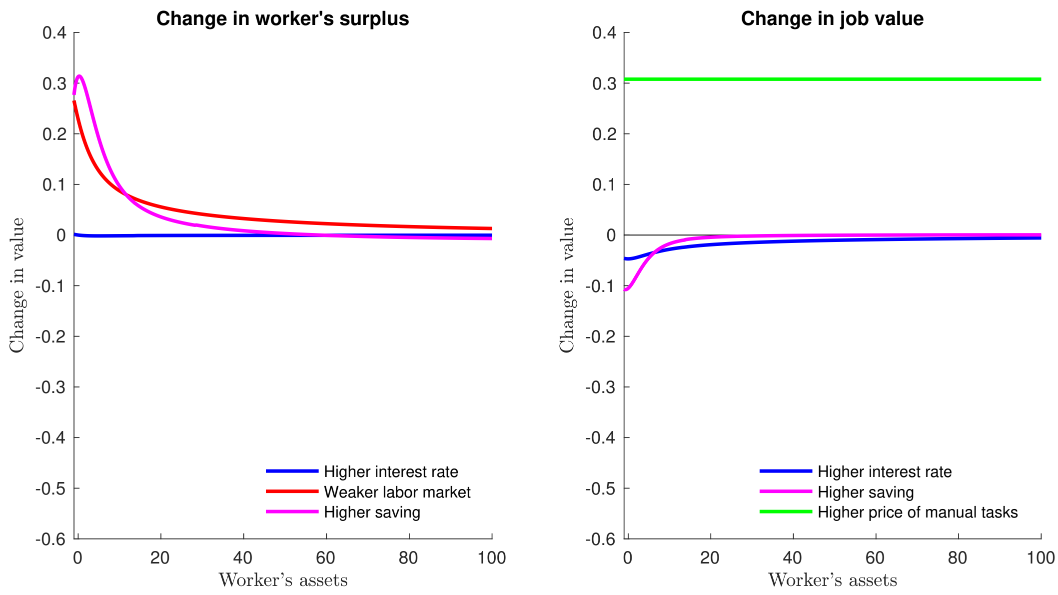

Figure 2.15: Forces driving changes in the worker's and firm's surpluses from a match

surplus so that a positive impact tends to push up on the wage and a negative one tends to depress it. The middle panel plots the multipliers (2.59) and (2.60). The dotted line traces out the quantity $\frac{\partial[W(a,s_e)-W(a,s_u)]}{\partial w(a)} - \frac{\partial J(a)}{\partial w(a)}$, which is the relevant scaling factor that one needs to apply to the forces in the left panel to arrive at their impact on the wage curve. Finally, the right-panel shows the decomposition of changes in the wage schedule across the two stationary equilibria. The solid line shows changes in wages, and colored bars show different forces at play. There are two key reasons why poorer workers lose out: first, firms need to increase their wages by less to compensate them for increased revenue product following the rise in $\bar{p}$, because even a small pay rise will yield a substantial utility benefit for these workers. Second, changes in the macro environment – higher interest rates and weaker labor market – mean that poorer workers are more desperate to avoid unemployment, reducing their effective bargaining power.

2.4.2 Transition dynamics

Many interesting questions about the impact of automation are about the short- and medium-run dynamics. To investigate those, I develop an algorithm to compute the dynamic equilibrium of the model (Figure 2.17).

On impact, the unemployment rate spikes up as the separation rate increases endogenously and displaced workers join the unemployment pool. The rise in interest rates makes saving more attractive, and the effective bargaining power of the workers deteriorates. This, combined with the increase in the price of manual goods, makes hiring particularly attractive, so that vacancies increase. Indeed, the rise in vacancies is larger than the rise in unemployment on impact, so that tightness goes up. As a result of strong hiring activity, unemployment falls back rapidly and remains low for around a decade.

Capital accumulation drives the long-term adjustment. Workers have higher holdings of capital and consequently demand higher wages. In line with the discussion above, the shift of

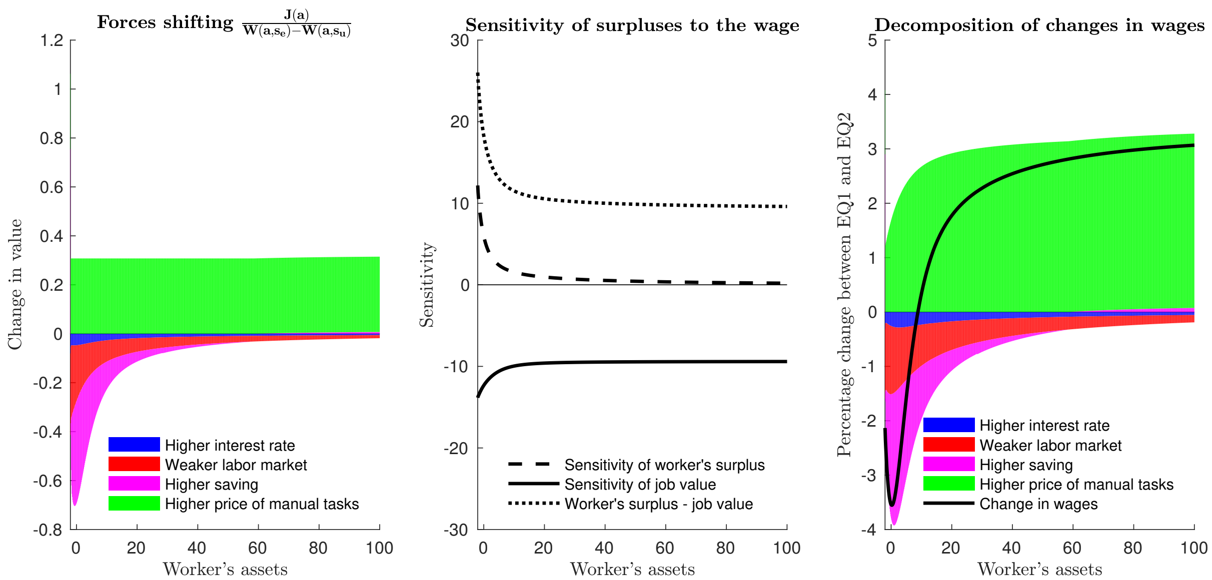

Figure 2.16: Arriving at a decomposition of the change in the wage across the asset distribution

the entire distribution of workers to the right dominates the increase in the value of a match for every level of worker's assets, resulting in weaker labor market in the long-run.

Figure 2.18 shows the response of aggregate variables. Capital stock, a state variable, is unchanged on impact. But demand for capital has increased, and to clear the asset market, interest rate jumps. Automation increases on impact, but the fall in wages and the rise in the interest rate provide some offset, so that the immediate increase is not very large. Output increases with the positive technology shock but so does saving, with consumption falling initially.

Top panel in Figure 2.19 illustrates the dynamic transition path of the distribution of workers, and the bottom panel tracks the changes in wages. Initially, wages across the entire distribution decline with the shock – this is purely a result of the shift in bargaining power away from workers and towards firms. Over time workers accumulate savings thus regain the bargaining power.

2.5 Macroeconomic effects of new automation technologies

2.5.1 Asset demand when technological constraint is binding

When the economy is constrained by the availability of the automation technologies, the demand for physical capital has a kink at the interest rate at which $\tilde{m} = \mu$ (Figure 2.20).[24] The demand for assets is a horizontal sum of the value of manual firms' equity and demand for machines.

[24]For interest rates above that level, changes in the interest rate affect demand for capital both through the intensive and the extensive margin: a lower r leads to firms already using automation technologies to expand, and to new firms switching from the manual to the automated technology. When the constraint is binding, capital demand becomes steeper as changes in the interest rate only work through the intensive margin.

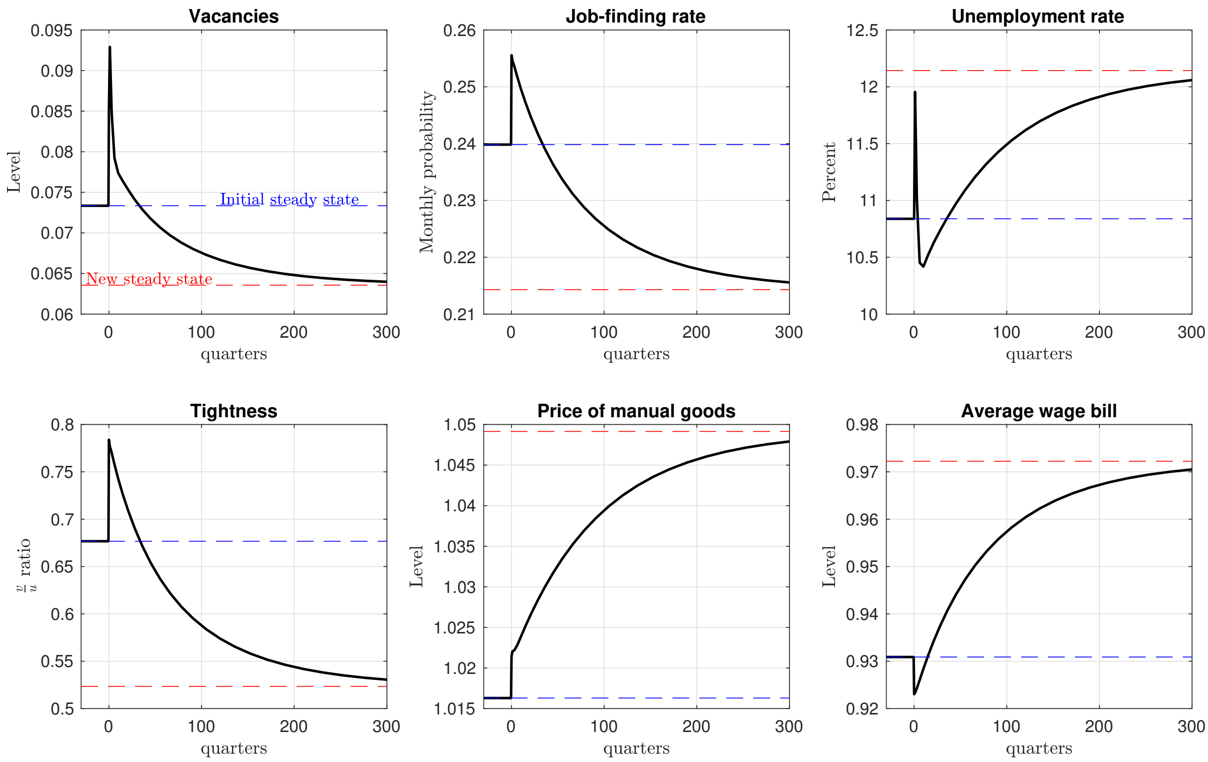

Figure 2.17: Dynamic path of the labor market following the rise in automation efficiency

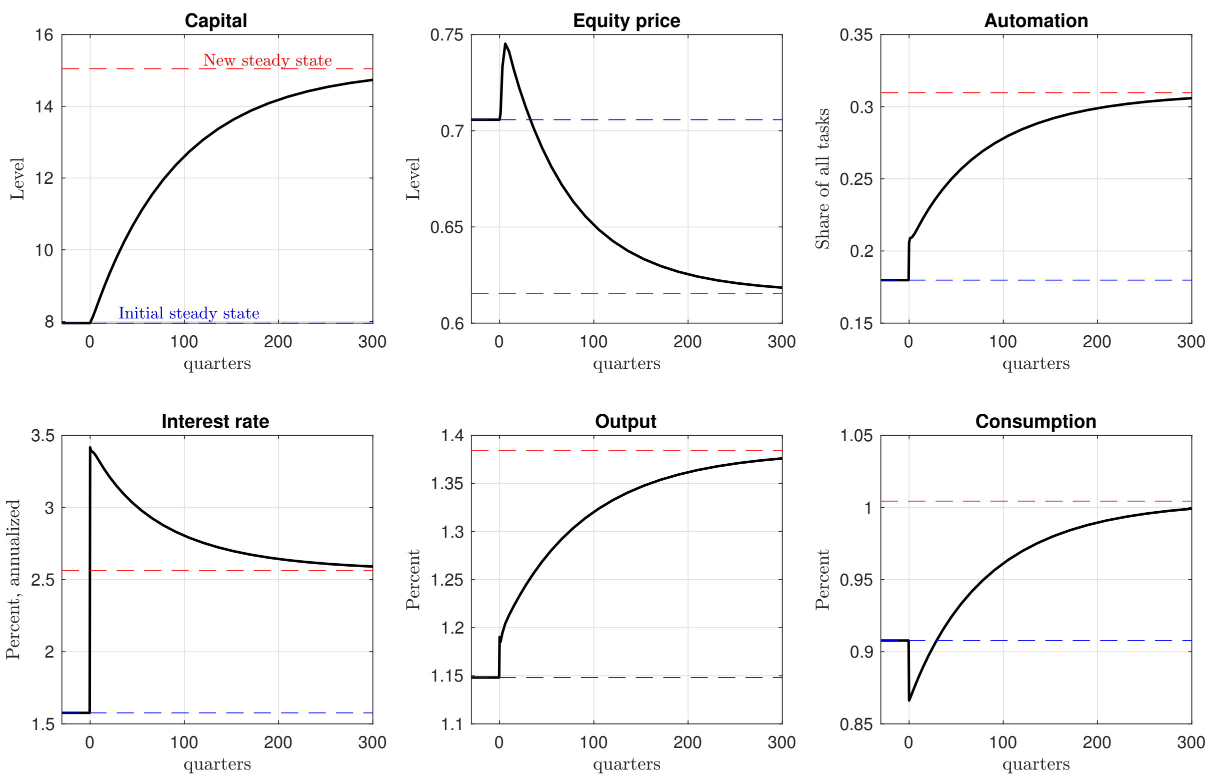

Figure 2.18: Transition path of the aggregate variables

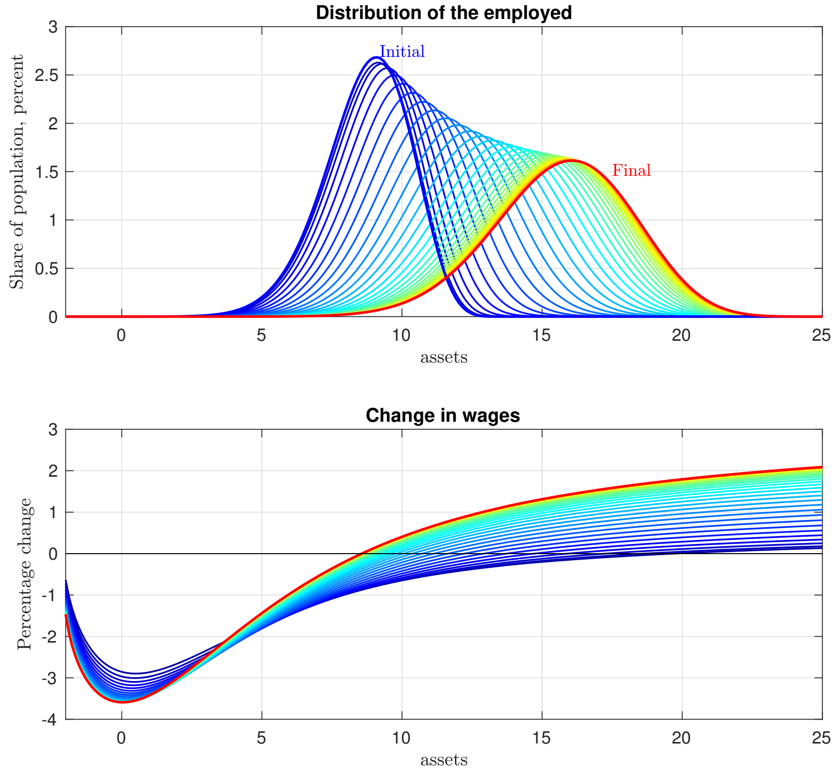

Figure 2.19: Transition of the distribution and the wage schedule
Note: the thin lines show the distribution during the transition. The thick lines are the initial and the final steady states.

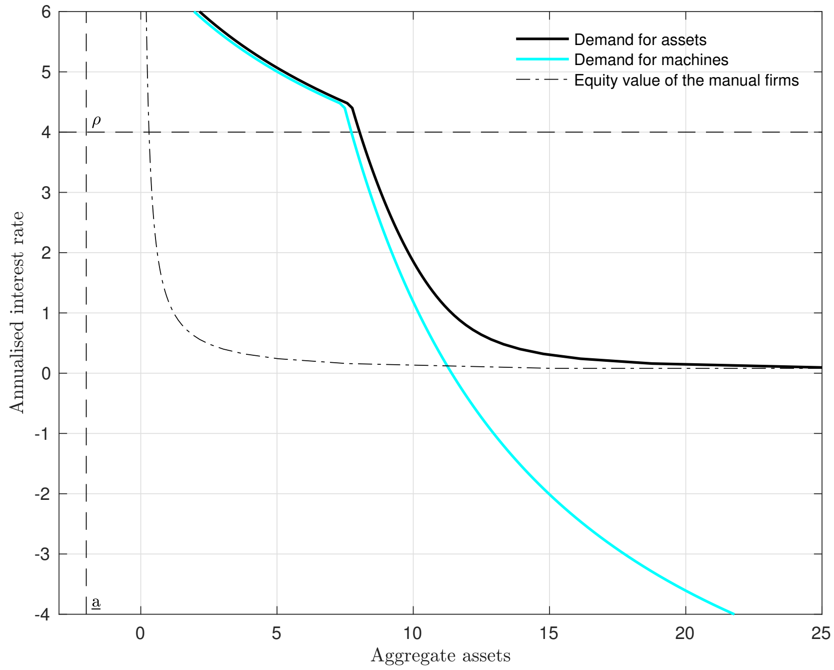

Figure 2.20: Demand for assets with binding technological constraint

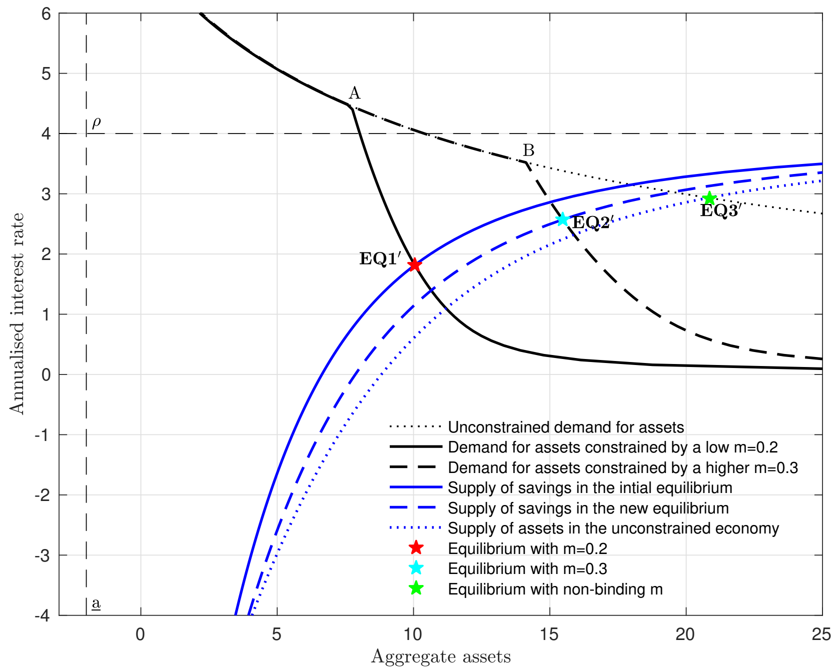

Figure 2.21: Stationary equilibria that are technologically constrained

2.5.2 Stationary equilibrium

I now use the model to trace out adjustment following a shift in the technological threshold μ. Figure 2.21 illustrates how the stationary equilibria are determined when the technological constraint is binding. The initial constrained equilibrium, labelled $EQ1'$, occurs where the two solid lines cross. Automation is constrained by the technological threshold of $\mu_0 = 0.2$. A wave of new automated technologies shifts the constraint from μ_0 to $\mu_1 = 0.3$ which moves the kink along the asset demand schedule and to the right, from point A to point B. As I show momentarily, the labor market conditions deteriorate in the new stationary equilibrium which shifts the asset supply curve to the right, based on the familiar logic from the previous section. Depending on the precise size of the shock and the response of factor prices, the final equilibrium can be constrained by the new technology threshold, as the $EQ2'$ in the Figure. Alternatively, the new constraint can be slack in the final equilibrium, as in $EQ3'$. The strength of the GE forces shifting the asset supply schedule it the right is different across the two cases.

2.5.3 Dynamics

I now briefly highlight the key features of the dynamic adjustment path, focusing on the case where the new steady state is unconstrained by technology ($EQ3'$).[25]

The key difference compared to the previous Section is that the price of manual tasks $\bar{p}$ declines, as the displacement effect dominates the productivity effect (Proposition 17). As a

[25]The transition dynamics that follow the shock which brings the economy from one constrained equilibrium to the other (e.g. a shock that shifts the economy from $EQ1'$ to $EQ2'$) are somewhat intricate, and computing the dynamic transition path is computationally demanding. For these reasons I present these details in the Appendix 2.C. Here instead I focus on a shock that raises the technological constraint μ sufficiently so that the technological constraint ceases to be binding in the new steady state.

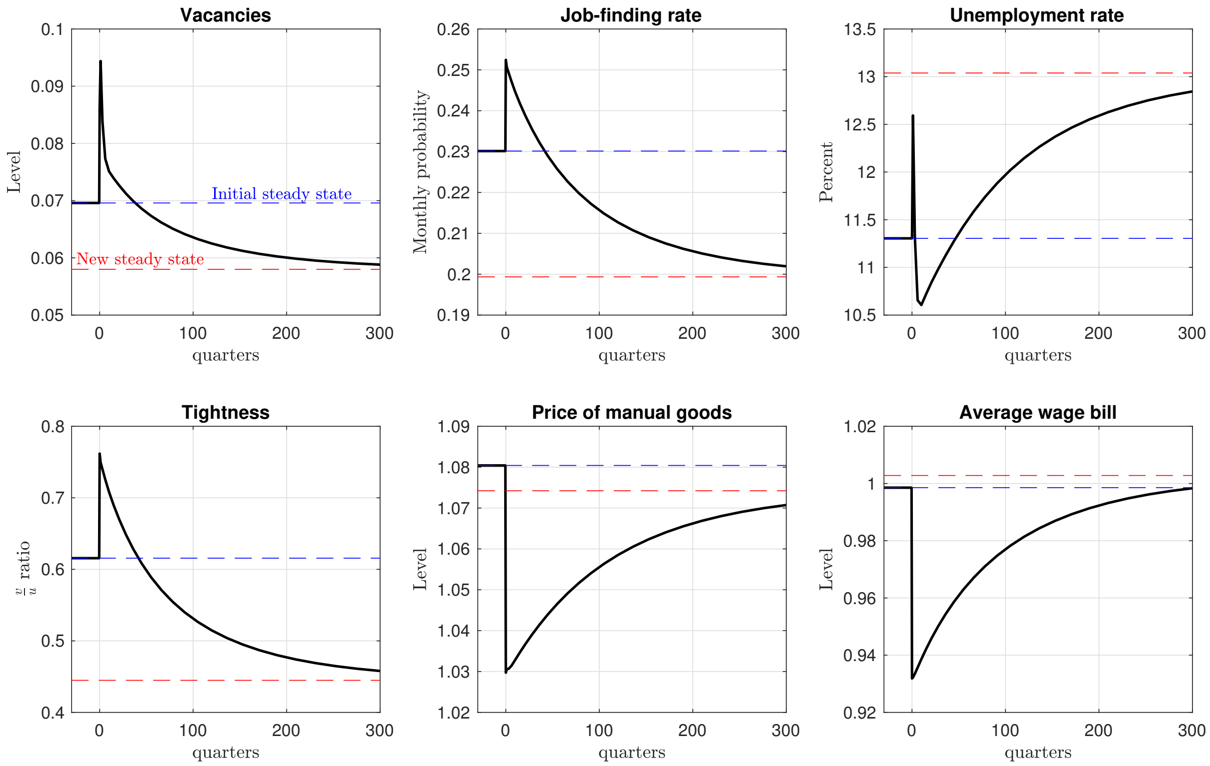

Figure 2.22: Dynamic path of the labor market following the invention of new automation technologies

result, the average wage drops sharply and remains depressed throughout the transition period.

Figure 2.24 shows that wages across the distribution decline sharply on impact, and recover only partially and towards the top of the initial distribution. The top panel illustrates the familiar pattern of the population moving up the asset spectrum. The small increase in average wages that we observe in the new steady state is entirely driven by workers saving more and thus moving along the wage schedule. The wage schedule itself is lower for all asset levels.

Finally, it is interesting to consider the behavior of the distribution of total income in this economy. Such analysis is particularly important given the "uneven growth" phenomenon in the United States and other advanced economies.[26] Figure 2.25 shows that the model is able to reproduce this qualitative pattern without appealing to the (undoubtedly important) skill heterogeneity observed in the data. The story here is that the asset rich benefit from the increase in return on their portfolios while the income of the asset poor is determined purely by wages which are weaker.

2.6 Conclusions

This paper is a preliminary attempt to capture the richness of macroeconomic adjustment to labor-replacing technologies. The general equilibrium model developed here features uninsurable unemployment risk for individual consumers, which automation can endogenously affect

[26] See, for example, https://www.census.gov/library/visualizations/2015/demo/real-household-income-at-selected-percentiles–1967-to-2014.html.

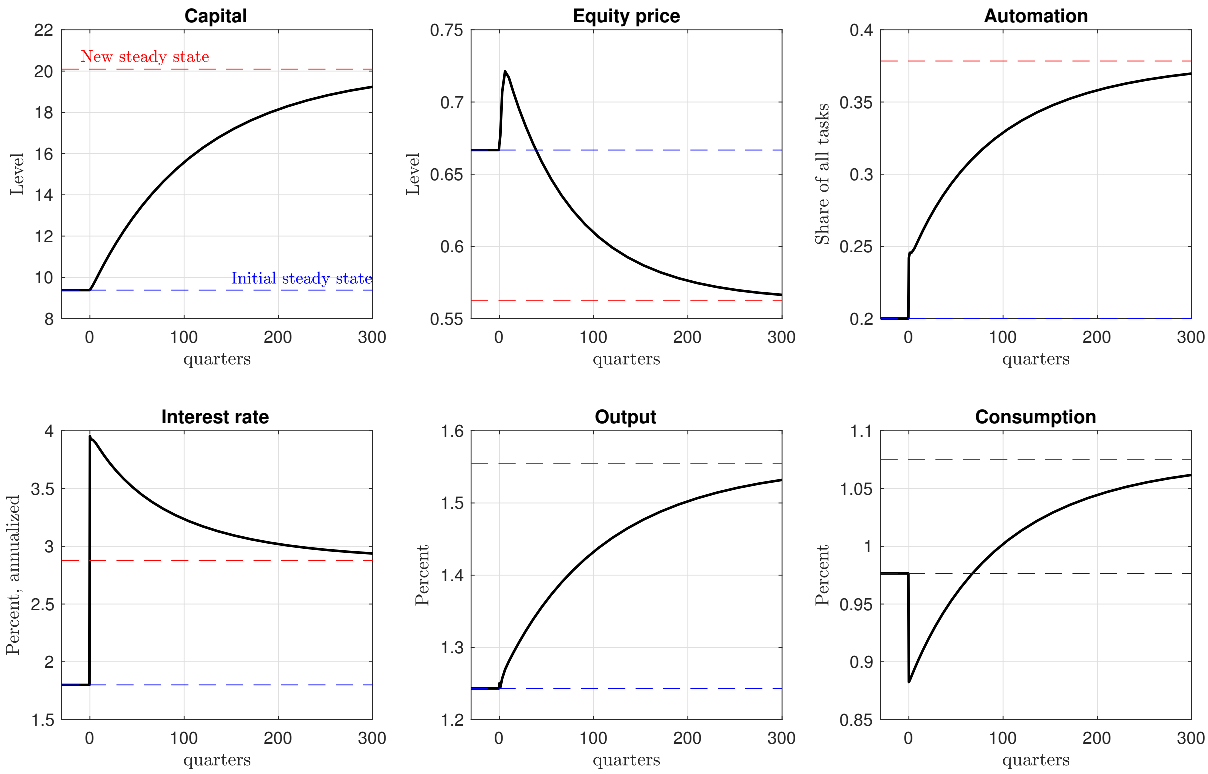

Figure 2.23: Transition path of the aggregate variables following an increase in μ

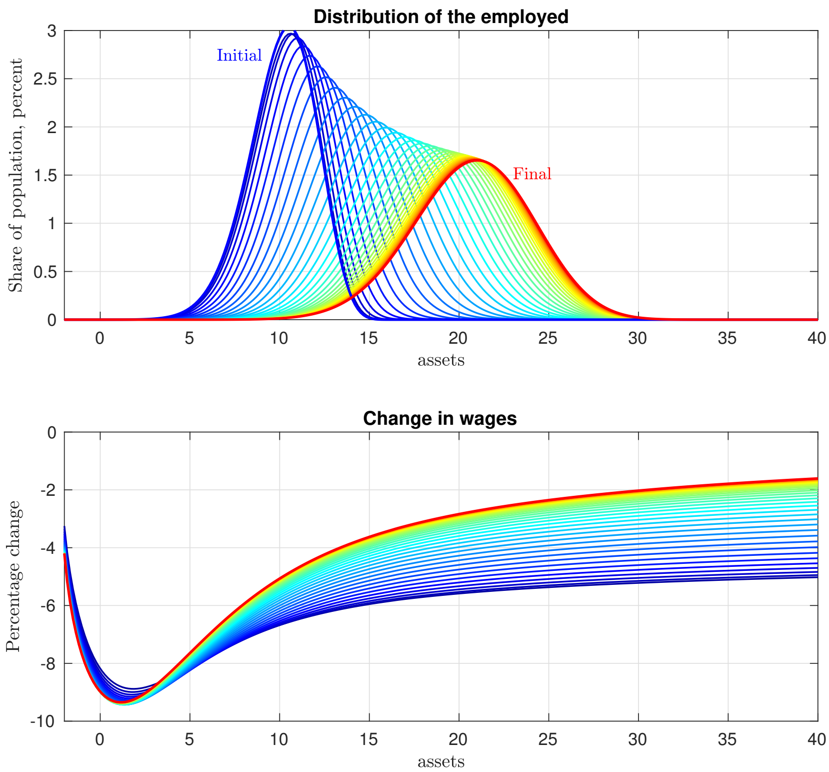

Figure 2.24: Dynamic path of wages across the distribution after an increase in μ

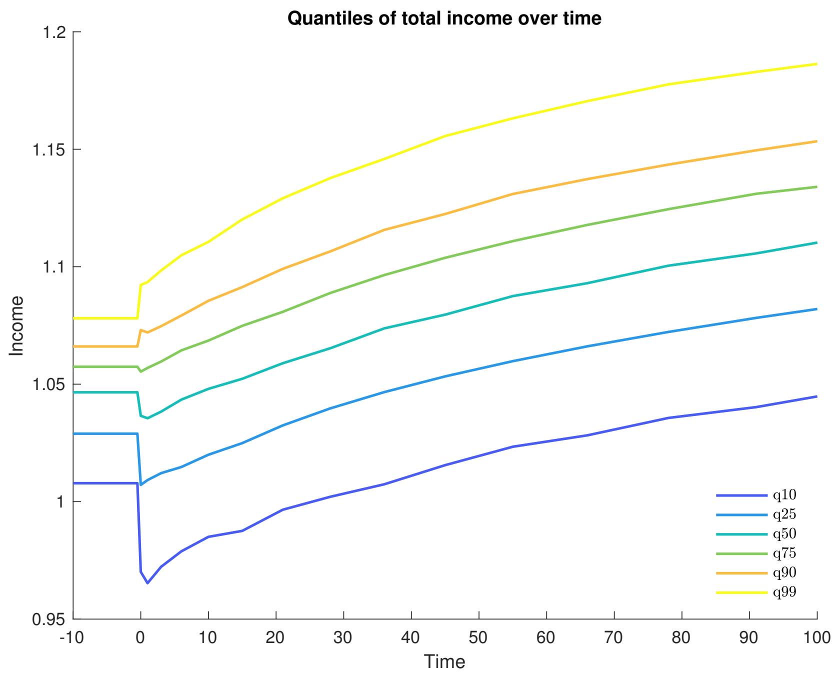

Figure 2.25: Time evolution of total income quintiles

through its impact on the labor market. The model offers a rich laboratory to investigate an economy's adjustment to automation shocks, both those that have occurred in the past, as well as those that may arrive in the future.

The analysis highlights two key results. First, labor replacing technologies may lead to a fall in wages that is large enough to boost demand for labor and job creation, leading to lower equilibrium unemployment. Labor-replacing technologies are thus plausible candidate explanations for the pattern observed across many advanced economies, namely the juxtaposition of low unemployment and subdued wage growth.

Second, the framework highlights the importance of general equilibrium effects. Greater uncertainty and higher wage dispersion increase the desire to save. These forces result in higher capital and lower interest rates – and, by the same token, in further adoption of automation technologies.

Going forward I plan to build on the framework developed here and expand the analysis in several important directions. First, I plan to refine the link between individual asset holdings and earnings. The current mechanism – wage bargaining between individual workers and their employers – is simple and easy to relate to the existing labor market literature. However, the slope of the wage schedule is ultimately an empirical matter (Appendix 2.B). Given the rising prominence of the incomplete markets literature in macroeconomics, bringing evidence to bear on the issue of how the labor market outcomes are affected is a priority. A different but potentially complementary direction is to investigate a different labor market structure, where instead of wage bargaining workers face wage offers posted by prospective employers. In such a framework individual assets will determine the optimal reservation wage and search effort, which might be somewhat easier to bring to the data. Relatedly, labor market dynamics could be cast in the context of a job-ladder framework, with workers affected by automation required to pay a fixed cost to re-skill. In this setting, heightened risk of automation may again lead to higher saving propensities and similar general equilibrium feedbacks.

Second, the model can be calibrated more carefully to the data and used to quantitatively assess the impact of labor-replacing technologies on the aggregates and the distributions. Such model would likely need to feature different matching efficiencies following a loss of employment driven by automation, and a life-cycle component to align the equilibrium wealth distribution more closely with what we observe in the real world.

Third, the model can be used to explore various policy options, from taxation of machines through minimum wages to universal basic income. Market incompleteness in the present framework leads to an aggregate externality, whereby individual attempts to self-insure against the risk of technological automation make this risk larger by increasing equilibrium adoption. This opens up an intriguing possibility that policy intervention may improve outcomes. Analysis of this possibility is an exciting avenue that I continue to pursue in ongoing work.

Chapter 3

On falling neutral real rates, fiscal policy, and the risk of secular stagnation

3.1 Introduction

What is the interest rate that is consistent with stable macroeconomic performance of a modern, developed economy? Few questions can rival this one in its difficulty and importance. Equilibrium real interest rates are unobservable; they are affected by a wide swathe of macroeconomic forces, both domestic and global; and are not invariant to policy regimes. All those factors make the assessment difficult and the answers uncertain. And yet a good handle on the equilibrium interest rate is fundamental to our ability to correctly assess the state of the economy, predict the future trends, and set policy appropriately.

The secular stagnation debate refocused attention on this important issue by pointing towards the downward trend in long-term interest rates across the developed world over the past four decades (Summers (2015b), Teulings and Baldwin (2014)). This decline, which started well before the financial crisis, is broad-based, both across countries and across asset classes. Yields on long-maturity inflation-protected government securities have been trending down since at least the early 1990s, and have hovered around their lowest levels on record over the past decade (left panel of Figure 3.1). The 5-year/5-year forward swap rates, which are less likely to be driven by the time-varying liquidity premia, are close to 0% in real terms (right panel of Figure 3.1). The decline in risk-free rates has, to a large extent, been mirrored in risky asset returns, such as rates of return on corporate bonds and on equites.

Policymakers have taken notice. Federal Reserve Chairman Jerome Powell's recent remark that the nominal federal funds rate – at the time set at between 2-2.25% – was "*just below the broad range of estimates of the level that would be neutral for the economy*" puts the level of the real neutral rate in the United States at around 0.5% (Powell (2018)). In Japan, faced with very low neutral rates for a long time, the central bank has engaged in aggressive monetary easing including directly targeting long-term interest rates (Kuroda (2016)). Similarly, European policymakers highlighted the equilibrium rate of interest as the key policy variable (Constâncio (2016) and Draghi (2016)), while the recent ECB paper (Brand, Bielecki, and Penalver (2018)) concluded that *“most of estimates of R* for the euro area have been negative regardless of the type of model used”* .

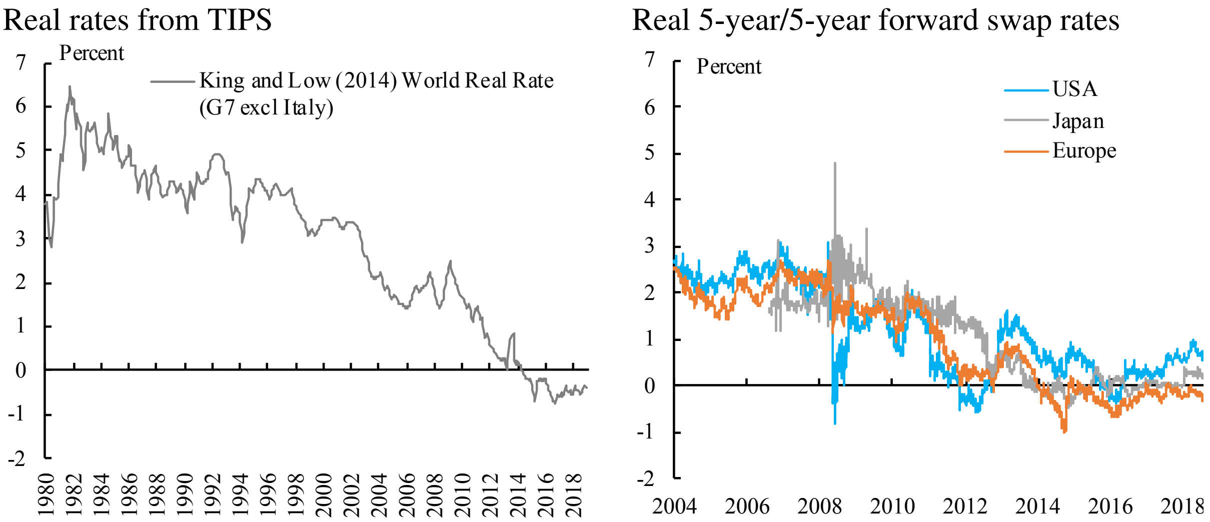

Notes: The world real rate is calculated following the methodology in King and Low (2014): it is the average of interest rates on inflation-protected government debt securities across the G7 excluding Italy. Data are from DataStream and form an unbalanced panel. In particular, the Figure relies on the UK inflation-indexed gilts in the early part of the sample. The US TIPS yield is the yield on a constant maturity 10-year Treasury Inflation-Indexed Security, retrieved from FRED, Federal Reserve Bank of St. Louis (code DFII10). Swaps data are from Bloomberg.

Figure 3.1: Real interest rates estimated from the inflation-linked bonds and forward swaps

Research to date has largely focused on the analysis of neutral rates for individual countries.[1] In this literature, the decline in the safe interest rates in advanced economies is a robust finding across studies that employ a wide range of tools and methods. For example, Marco Del Negro, Marc Giannoni, Domenico Giannone and Andrea Tambalotti (2017) established that a time-series model and a structural general equilibrium model both detect a downward trend in the equilibrium rate of interest in the United States since the mid-1990s. Internationally, Kathryn Holston, Thomas Laubach and John Williams (2017) applied the seminal methodology of Laubach and Williams (2003) to four advanced economies (United States, Canada, Euro Area and the United Kingdom), and found that the decline in the real rate of interest is present in each of them. As we explain below, estimating neutral real rates for individual open economies is a questionable procedure, since it implicitly takes as given an endogenous variable—the trade surplus or deficit.

We therefore estimate the neutral real rate using aggregated data for all of the advanced economies (as if they formed a single, fully integrated economy). Our results suggest that real rate for the advanced economy block – what we call AE R* for brevity – declined by around 3pp over the past 40 years, and is currently only slightly above zero in real terms.

Researchers have explored a wide range of potential drivers behind the decline in real in-

[1] A recent exception is the paper by Del Negro, Giannoni, Giannone, and Tambalotti (2018) who study trends in world interest rates in an econometric framework.

terest rates.[2] Etienne Gagnon, Benjamin Johannsen and David Lopez-Salido (2016), Carlos Carvalho, Andrea Ferrero and Fernanda Nechio (2015), Noemie Lisack, Rana Sajedi and Gregory Thwaites (2017) and Gauti Eggertsson, Neil Mehrotra and Jacob Robbins (2019) all used macroeconomic models with an overlapping generations structure to show that demographic trends can act as powerful forces driving the intertemporal preferences and hence intertemporal prices. Adrien Auclert and Matthew Rognlie (2016) as well as Ludwig Straub (2017) explored channels linking income inequality and real interest rates in general equilibrium models. The work of Barry Eichengreen (2015) stressed the importance of investment-specific technological change and the resulting decline in the price of capital goods, while the recent study by Emmanuel Farhi and Francois Gourio (2018) considered other supply-side forces such as intangible capital and market power as the drivers of the real interest rate decline.

While the literature has covered a lot of ground in terms of possible private sector explanations for declines in real rates, it has not highlighted an important issue. One would have expected a period of substantially enlarged government deficits and debt and substantially enhanced pay-as-you-go public pensions and increased health insurance to *ceteris paribus* have substantially raised neutral real rates.[3] That this has not occurred suggests that the economic forces operating to reduce neutral real rates have been stronger than has been generally recognized.

Public policies affect the interest rate through a range of channels. We review these mechanisms through the lens of several theoretical paradigms, concentrating in particular on the role of government borrowing, which is the main focus of both theoretical and empirical literatures in macroeconomics. We then survey the existing empirical estimates of the impact of government debt on interest rates. Simple calculations using observed estimates of the impact of deficits on interest rates suggest that the increase from 18% to 68% in the public debt-to-GDP ratio of the advanced economies should *ceteris paribus* have raised real rates by between 1.5 and 2 percentage points over the last four decades. This effect is quantitatively important but is a presumptive underestimate of the impact of fiscal policy changes given the rise in pay-as-you-go government pensions and other social insurance programs. This analysis leads to the conclusion that the fall in real long-term interest rate observed in the data masks an even more dramatic decline in the equilibrium “private sector” real rate.

To incorporate other policies such as the increase in social security and healthcare spending into our analysis, to build further understanding of the mechanisms involved, and to cross-check the magnitudes of these effects, we study these phenomena in a dynamic general equilibrium framework. We construct two tractable models, each one designed to capture different channels through which policies play out in equilibrium.

[2]An older literature considered why interest rates were so high in the 1980s. Blanchard and Summers (1984) conjectured that high interest rates in that period were driven in part by higher profitability of investment. Their conjecture was subsequently supported by Barro and Sala-i-Martin (1990) who identified a strong role for stock prices – a proxy of anticipated investment profitability – in affecting world interest rates.

[3]One recent study, developed independently and in parallel to ours, which points out the role of government debt is that of Eggertsson, Mehrotra, and Robbins (2019). In comparison to their paper, we focus on a wider set of policies (including social security and old-age healthcare), we consider empirical estimates of the link between debt and R*, and we analyze a fuller extent of theoretical channels through which this link operates.

Building on the work of Mark Gertler (1999), the first model captures the life-cycle behavior, with workers saving for retirement and retirees decumulating their wealth. The resulting heterogeneity in marginal propensities to consume and differences in the implicit discount rates across agents mean that Ricardian Equivalence – the proposition that government borrowing decisions are neutral in equilibrium – does not hold in our model, making the effects of a range of government policies on real rates non-trivial.[4] We simulate the model with the profiles of government debt, government spending, social security and old-age healthcare expenditures that match the experience of developed economies over the past 40 years. These simulations suggest that shifts in these policies pushed equilibrium real rates up by around 3.5pp between the early 1970s and today.

Our second model is of the Bewley-Huggett-Aiyagari tradition. It focuses on idiosyncratic risks and precautionary behavior, channels that are absent from the life-cycle model. When markets are incomplete, government debt is an asset which households can hold to self-insure. Supply of government bonds thus determines the total supply of investable assets, and hence the interest rate in equilibrium. We calibrate this model in a parallel fashion to the life-cycle model, ensuring that the degree of income inequality in equilibrium matches what we observe in the data. Our explorations suggest that the increase in the supply of government bonds has pushed interest rates up by about 50-70bps through this precautionary 'supply of safe assets' channel. Overall, then, we find that public policies may have pushed interest rates up by around 4pp.

Our final contribution is to use the two models to consider a wider range of secular trends in a coherent and unified way. This is motivated by the fact that much of the literature, including the studies cited above, maintain a narrow focus on one secular trend at a time. Some of the exceptions, such as cross-cutting studies of Davide Furceri and Andrea Pescatori (2014) and Łukasz Rachel and Thomas Smith (2015), used a simple reduced-form saving-investment framework to aggregate the different influences and thus suffered from a potential consistency problems – for instance, it was impossible to detect any non-linearities or prevent double-counting.[5] We can speak to this issue because, despite their rich structure, our models are highly tractable and well-suited for an internally consistent analysis of several factors widely regarded as instrumental in explaining the safe rate trends. We show how to use the models to quantitatively assess the impact of the slowdown in productivity growth, the demographic shifts, and the rise in income inequality. As a result, we arrive at a quantitative decomposition of the decline in the real interest rate in advanced economies which takes into account both private and public sector forces (Figure 3.2). Taken together, in our central calibration, these selected factors under-explain the decline in advanced economies' neutral real rate that we estimated, pointing to other forces which we have left out of our analysis. We hope that future research will enrich our framework and incorporate those trends.

[4]Following a change in government finances, there is some Ricardian offset, but unlike in the representative agent model, this offset is incomplete.

[5]Eggertsson, Mehrotra, and Robbins (2019) construct a large quantitative macro model and use it to consider several hypotheses.

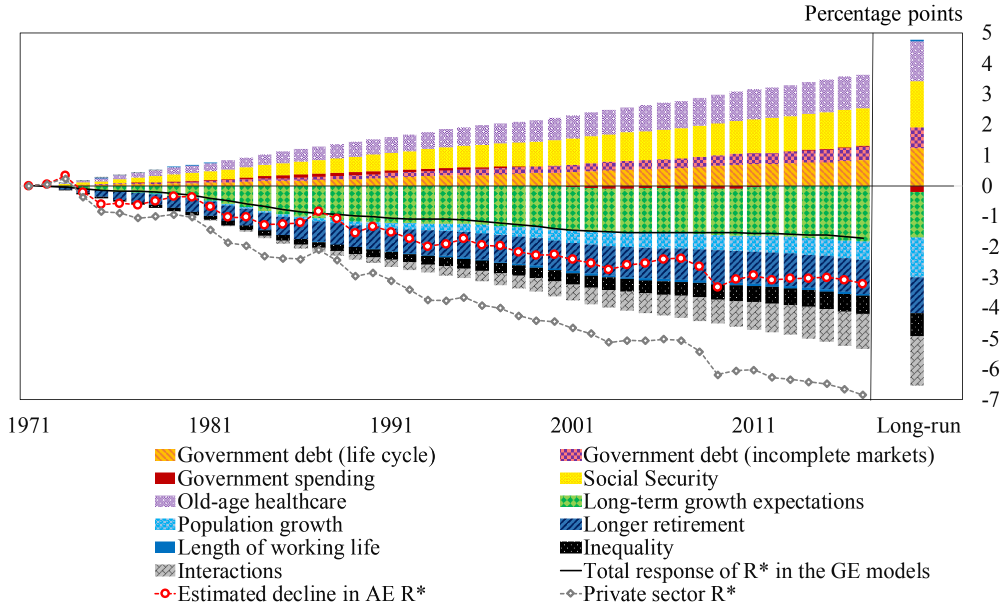

Figure 3.2: Changes in the equilibrium real interest rate as a result of policy, demographic and technological shifts

Our findings suggest that the private sector forces dragging down on interest rates are more powerful than previously anticipated, and that on average across the business cycle, equilibration of private-sector saving and private-sector investment may indeed require very low real rate of interest in advanced economies in the years to come. This conclusion is consistent with findings of Oscar Jorda, Katharina Knoll, Dmitry Kuvshinov, Moritz Schularick and Alan Taylor (2017), who established that the current low levels of interest rates are not unusual in historical terms.[6] It is also consistent with the Japanese experience.[7] Our findings raise the possibility that the developed world is at risk of mirroring the experience of Japan, whereby the very low equilibrium rate of interest appears to be a semi-permanent feature of the economic landscape.

The remainder of this paper is structured as follows. Section 3.2 discusses two methodological issues underlying our analysis. Section 3.3 contains the results of the estimation of the long-term equilibrium real interest rate for advanced economies. Section 3.4 starts with a discussion of the channels through which government policy influences the equilibrium rate; it then summarizes the results from the existing empirical literature which estimates the size of these effects; and finally it uses these elasticities to calculate some back-of-the-envelope measures of how government borrowing affected AE R*. In Section 3.5 we set up the two general equilibrium models and use them to study the impact of government policies. Section 3.6 contains further simulations, arriving at the full decomposition of the decline in the natural rate, including the impact of secular demographic changes, slowdown in technology and the rise in inequality. Section 3.7 concludes. Throughout the main text we focus on the results and keep

[6]Theoretical work focused on the possibility that the real rate remains depressed for a long periods of time or even indefinitely. Alejandro Justiniano and Giorgio Primiceri (2010), Gauti Eggertsson and Neil Mehrotra (2014, 2016, 2019), Bob Hall (2017), as well as Ricardo Caballero, Emmanuel Farhi and Pierre-Olivier Gourinchas (2016, 2017) provided formal models in which this possibility arises.

[7]For studies of the natural rate in the context of Japan, see Fujiwara et al. (2016), Okazaki and Sudo (2018) and Wynne and Zhang (2018).

the technical analysis to the minimum, delegating the details to online Appendices.

3.2 Methodological issues

We begin with a discussion of two methodological issues that permeate the analysis in this paper. The first is our treatment of advanced economies as a block. The second is our view that the decline in neutral real interest rates can be understood through the balance between desired saving and investment. This view leads us to focus on the macroeconomic forces affecting a broad range of returns, rather than on factors driving spreads or premia on particular financial instruments.

3.2.1 Advanced economies as a block

Our analysis assumes that the advanced economies block is fully integrated. In practice, we simply use aggregated data for all the developed countries (members of the OECD, whenever data are available), 'as if' the block was a single economic entity. Moreover, we treat this block as a large, closed economy. Obviously, these assumptions are a gross simplification to reality. But we believe that they are reasonable and that they constitute a significant improvement relative to the existing literature, which largely treats countries as separate economic entities.

These assumptions are justified by large capital flows across developed economies, strong commonality in trends in long-term real rates observed in the data, and the fact that net saving at the level of the block has been an order of magnitude smaller than the current account gaps of individual countries, fluctuating by less than 1.5 percent of GDP from peak to trough and trending upward over the recent decade (Figure 3.3).

More importantly, our approach avoids the erroneous assumption implicit in much of the country-level analysis that the economies under consideration are closed. Current account balances for individual economies are large and variable; they are endogenous outcomes of the saving and investment propensities within each economy relative to the global average. A country for example that runs a chronic trade surplus will be found to have a neutral real rate at a level where domestic demand is short of potential output and the reverse will be true for a country running a chronic trade deficit. External balances should therefore be taken into account in such country-level analyses. We instead posit that developed economies taken together experience structural excess saving, reflected in the trend-decline in real interest rates without a discernible trend in their current account. At this level of aggregation, the country-level differences wash-out, and econometric and theoretical analysis based on a closed-economy assumption is thus more credible.

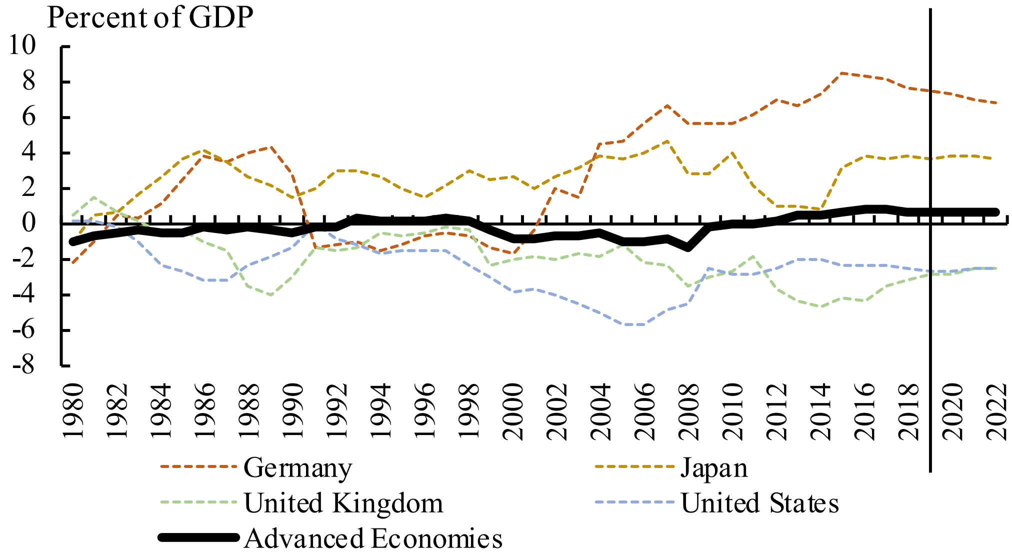

Note: The black line shows the current account for advanced economies as defined by the IMF. The OECD data are available only since 2004. Over the period when both data series are available they are very close to each other.

Figure 3.3: Current account in advanced economy block and in selected individual economies

3.2.2 Desired saving and investment framework, and the role of safety and liquidity premium

We carry out our analysis on the basis of the premise that, for analyzing long-term trend movements in neutral real rates, it is appropriate to focus on factors relating to saving and investment propensities rather than issues of liquidity or risk. Consequently, our analysis abstracts from aggregate uncertainty and differing levels of liquidity of various assets. Instead, in the empirical part of the paper we focus on the yield on practically risk-free,[8] highly liquid government bonds, and our theoretical framework makes no distinction across asset classes (meaning there is only one interest rate in our models).

Several facts support our approach. First, the decline in rates on highly liquid securities track declines in yields on relatively illiquid government indexed bonds and real swaps (Figure 3.1), suggesting that the liquidity characteristics of government bonds play only a secondary role. Second, even in the US, there has been little trend movement in spreads between Treasury securities and corporate securities in given rating classes, and while the pick-up in equity risk premia has been somewhat more pronounced, it is nonetheless small relative to the decline in real interest rates over the decades (the left panel in Figure 3.4). In any case, it is not clear whether one should interpret any changes in spreads as driven by changes in risk preferences or rather a result of changes in how risky the underlying assets are perceived to be. For instance, the recent global financial crisis has likely led to a reassessment of what it means that an asset

[8]Government bonds in advanced economies are free of default risk. However, holdings these assets is risky, as the overall return is determined not only by the coupon but also by the changes in the valuation of the bond. For example, the standard deviation of the total annual return on a 10-year US Treasury bond is 9.5pp over the period 1970-2018.

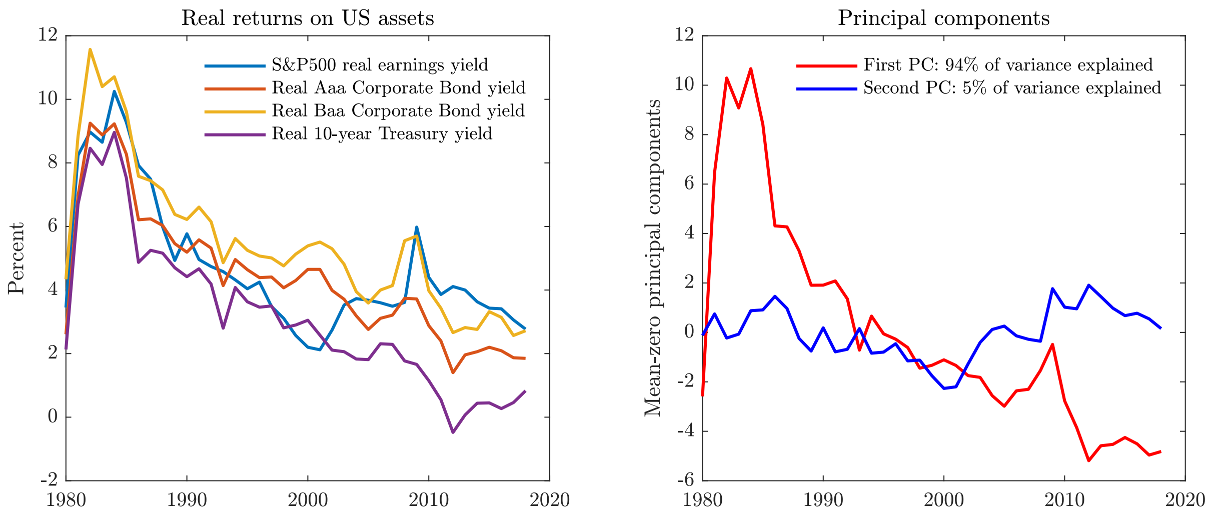

Source: FRED, Robert Shiller and author's calculations.
Notes: 10 year Treasury real yield post-2004 is the yield on 10-Year Treasury Inflation-Indexed Security (code: DFII10). Before 2004 it is the nominal yield (GS10) minus inflation expectations measure from the Michigan survey (code: MICH). The real corporate yields are nominal yields (codes: AAA, BAA) minus inflation expectations from TIPS (post-2004) or from Michigan survey (pre-2004). Real earnings yield is the inverse of the cyclically adjusted total return price real earnings ratio from Robert Shiller (available at http://www.econ.yale.edu/~shiller/data.htm).

Figure 3.4: Real returns in the United States: data and principal components

is triple-A rated; while the dot-com bubble appears to have had a lasting impact on pricing of equities.

To gauge a sense of the importance of the trend decline in real returns versus changes in the dispersion between them, we summarize the patterns in the US data using principal components analysis (Panel B in Figure 3.4). The first principal component picks up the downward trend visible in all returns, explaining 94% of the total variance in real returns in the US. The second principal component, which appears to be related to the increase in the "convenience yield", explains only 5%. These figures provide a useful perspective on the relative importance of the two narratives.

This evidence is consistent with the finance literature that investigates the decline in the neutral real interest rates in presence of term and liquidity premia. Jens Christensen and Glenn Rudebusch (2017) construct dynamic term structure models that account for time-varying term and liquidity risk premiums, and using data on inflation-linked bonds in the US, they estimate that the longer-run equilibrium real rate has fallen about 2 percentage points since the late 1990s. Stefania D'Amico, Don Kim and Min Wei use no-arbitrage term structure models to show that liquidity premium embedded in US TIPS has generally declined.

The "safe asset" literature finds somewhat larger role for the convenience yield, but even there the magnitudes are generally rather small relative to the large trend decline in real rates. Using very different approaches, Del Negro, Giannoni, Giannone, and Tambalotti (2018) and Rachel and Smith (2017) concluded that the rise in the spread between risky and risk-free rates accounted for about 70bps of the decline in risk-free rates. Another way to cross check the importance of the spread is to consider the relationship between government debt and corporate bond spreads, as in the classic work of Arvind Krishnamurthy and Annette Vissing-Jorgensen

(2012). Using data up to the global financial crisis, they showed that there has been a stable relationship between public debt and the spread in the United States, one that traces out the demand schedule for US Treasury debt. An update of their analysis that includes the data covering the past decade shows that the recent outturns defied the pre-crisis relationship. One interpretation of this finding is that the demand for US government debt has indeed shifted outwards. But the difference between the realized spread and the spread that one may expect based on the historical relationship is in the region of roughly 60 basis points, supporting the conclusion that changing safety and liquidity premia account for a minority of the total decline over the past several decades.

We also believe that our conclusions are robust to the lessons coming out of the "safe asset" literature. The key focus of our paper is on supply of government assets, which we argue has increased dramatically as government debt has risen. The point we make is simple: had the supply of safe assets been scarcer, the safe and liquid rate would have declined further. As far as this conclusion is concerned, we do not see any real tension here with the existing literature. One economic mechanism that is lacking in our exercise is that, when the economy is at the zero-lower bound and the equilibrium is achieved through lower levels of output and investment (when the economy finds itself in a "topsy-turvy" equilibrium as in Eggertsson and Krugman (2012)), a higher supply of safe assets can increase private investment. Instead, our models focus on the long-run equilibrium where the long-term rate adjusts freely, so that formally an increase in supply of safe assets crowds out private capital. But this observation would only further strengthen our policy conclusions which we discuss at the end of the paper, that given the low interest rate environment, government policy must play an active role in reinvigorating demand. On a more conceptual level, it seems plausible that some of the real-economy forces we discuss in Section 3.6 of the paper may have acted to put pressure on safe rates relative to risky rates, therefore potentially accounting for some of the 60bps increase in the convenience yield. For example, people may have preference for saving in safe assets for retirement, or they may demand safe assets to insure against idiosyncratic uncertainty. In short, we do not think that there necessarily is any tension between the "excess saving" explanation (one present in this paper) and "excess demand for safe assets" one.

3.3 Estimating the AE equilibrium real interest rate

We estimate the natural rate of interest for advanced economies adopting what is perhaps the most celebrated applied empirical model designed for this purpose, originally due to Laubach and Williams (2003) and recently re-applied internationally by Holston, Laubach, and Williams (2017b). Conceptually, this approach draws on two strands on the literature. By following Wicksell's (1989) definition of the natural rate as the rate consistent with stable inflation and output remaining at equilibrium ("potential") level, it is well aligned with the modern monetary theory, as in Walsh (1998), Woodford (2003) and Gali (2008). That literature is primarily concerned with fluctuations at the business-cycle frequency, where shocks move the economy around a stable steady state. In addition to those business-cycle shocks, the framework employed here is flexible enough to capture secular forces that affect the steady state.

3.3.1 Sketch of the model and the estimation procedure

Our approach to estimating the LW model is deliberately off-the-shelf: we use exactly the same procedures as the recent papers in that literature. Out contribution is solely to perform this exercise on the block of advanced economies as a whole. As such, we do not take a stance on the performance of the model, although we discuss some of the issues below.

The philosophy of the LW method is that the natural rate of interest is an endogenous object determined in general equilibrium, and as such it will depend on a host of socio-economic forces, such as trends in preferences, technology, demography, policies and policy frameworks, and so on. It is impossible to know and measure all of the relevant factors. At the same time, a robust prediction of most workhorse macroeconomic models is that the natural rate should vary together with the (expected future) trend growth rate of the economy.[9] To reflect the dependence on growth and on a range of (possibly unknown) other factors, the LW model assumes that the natural rate, denoted r_t^*, depends on the estimated trend growth rate of potential output g_t and a time-varying unobserved component z_t that captures the effects of other unspecified influences:

$$r_t^* = g_t + z_t. \tag{3.1}$$

The model further assumes that both the growth rate g_t and the unobserved component z_t are random walk processes:

$$g_t = g_{t-1} + \epsilon_{g,t} \quad \epsilon_g \sim N(0, \sigma_g^2) \tag{3.2}$$

$$z_t = z_{t-1} + \epsilon_{z,t} \quad \epsilon_z \sim N(0, \sigma_z^2) \tag{3.3}$$

The model specification also admits shocks to the level of potential output. Denoting by y_t^* the natural logarithm of potential output at time t:

$$y_t^* = y_{t-1}^* + g_{t-1} + \epsilon_{y*,t} \quad \epsilon_{y*} \sim N(0, \sigma_{y*}^2). \tag{3.4}$$

In short, the LW model views the natural rate as the sum of two independent random walks. To achieve identification, LW add two further equations to the model. First, they specify a simple reduced-form equation relating output gap to its own lags, a moving average of the lagged real funds rate gap, and a serially uncorrelated error:

$$y_t = y_t^* + a_1(y_{t-1} - y_{t-1}^*) + a_2(y_{t-2} - y_{t-2}^*) + \frac{a_r}{2}\sum_{j=1}^{2}(r_{t-j} - r_{t-j}^*) + \epsilon_{y,t} \quad \epsilon_y \sim N(0, \sigma_y^2) \tag{3.5}$$

The key in this estimated IS relation is the a_r coefficient, which we expect to be negative.

[9] We discuss the rationale for this link in some detail in Section 5.

Parameter point estimates (t-statistics in parentheses)							
a_1	a_2	a_r	b_π	b_y	σ_y	σ_g	σ_z
1.71	-0.79	-0.04	0.90	0.09	0.25	1.03	0.31
(21.65)	(10.28)	(2.13)	(17.78)	(2.06)	(5.30)	(29.63)	(9.38)
Average standard errors around the estimates							
		$y*$	$r*$	g			
		1.19	3.12	0.16			

Table 3.1: State-space model parameter estimates

Second, LW add the reduced-form Phillips curve to the model, linking current inflation π_t to lagged inflation and the output gap:

$$\pi_t = b_\pi \pi_{t-1} + (1 - b_\pi)\pi_{t-2,4} + b_y(y_{t-1} - y^*_{t-1}) + \epsilon_{\pi,t} \quad \epsilon_\pi \sim N(0, \sigma^2_\pi), \tag{3.6}$$

where the standard theory would suggest that coefficient b_y is positive.

The system above can be written in a state-space form, and the Kalman Filter can be used to estimate the unobservable states. To estimate the model, we use data for advanced economies as a block. The data comprise of (log) quarterly real GDP, core inflation and long-term interest rates over 1971Q1:2017Q4 for the aggregated sample of OECD countries. The interest rate series is the arithmetic average of long-term nominal interest rates across an unbalanced panel of 36 OECD economies.[10] To calculate real rates, we subtract from nominal rates a simple measure of expected inflation, constructed as the moving average of past core inflation rates, in line with Holston, Laubach, and Williams (2017b). See Appendix 3.A for further details on the data and the estimation procedure.

3.3.2 Results

Table 3.1 shows the coefficients of the estimated model. Point estimates are all significantly different from zero and have expected signs. In particular, a positive interest rate gap reduces the output gap, while a positive output gap raises inflation. Table 3.1 also shows the standard errors around the estimated trends, which are large, especially those around the estimates of equilibrium real rate. These wide standard error bands are not specific to our results – indeed, they are a norm in the literature. For instance, Holston, Laubach, and Williams (2017b) report similarly large errors for individual economies. These errors are, to an extent, an artifact of the long-sample, as they reflect the cumulative uncertainty of the underlying drivers of equilibrium rates. Nonetheless, these large error bands should act as a reminder of the high uncertainty surrounding the econometric estimates of equilibrium interest rates.

Figure 3.5 contains the key results. According to our estimates, AE R* declined steadily from 1980s onwards, and fell sharply during the crisis.[11] It then stabilized at low levels ($\approx$

[10]The results are robust to using weighted average or median of the interest rates across countries. Given the strong comovement, these interest rate series are close to each other.

[11]Estimates for the first decade should be taken with a grain of salt, as the model is less accurate during the first few years of the sample while the initial conditions play a larger role.

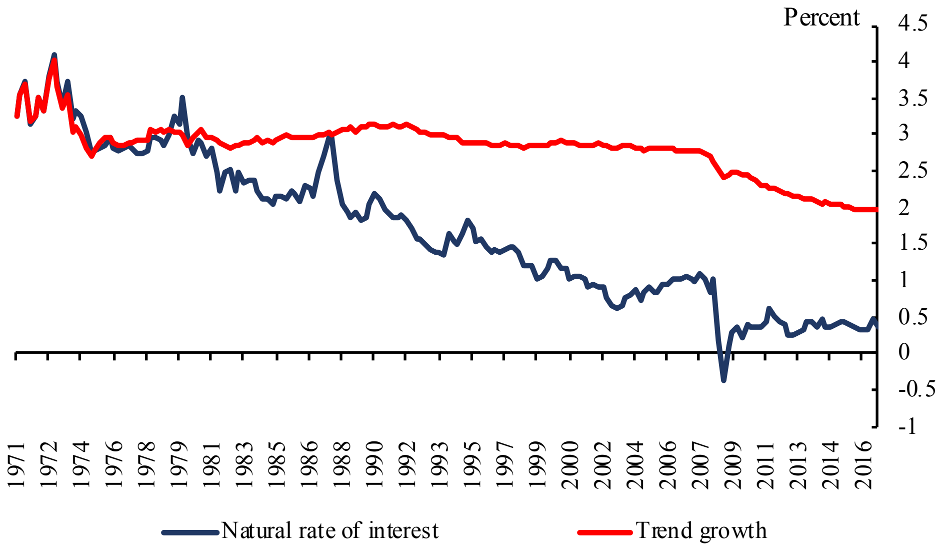

Figure 3.5: AEs R* and trend growth

0.5%). The estimated growth rate of potential output has been broadly stable up until the crisis, and declined during the crisis by about 1pp. Thus the model suggests that a bulk of the decline in real interest rates is due to factors other than trend GDP growth. This is consistent with the literature that finds only a loose connection between actual GDP growth and interest rates in historical data (Hamilton, Harris, Hatzius, and West (2016)).

These results corroborate other existing findings in the literature. In particular, Holston, Laubach, and Williams (2017b) estimated that the declines in real rates for US, Canada, Euro Area and the UK of around 2.3pp between 1990 and 2017; for comparison, the decline over this period for AEs as a whole that we estimate here is around 2pp.

Overall, despite large uncertainty surrounding the point-estimates of these trends, we interpret the results of this exercise as broadly in line with the country-level findings in the literature. Indeed, given the high level of aggregation, we find it encouraging that the estimated unobservables do well at picking up the main events, such as the global financial crisis, during which our estimate of AE R* declines very sharply.

To further illustrate the magnitude of the decline in real rates, we calculate the counterfactual private sector savings-investment gap under a scenario of no decrease in AE R*. The left panel of Figure 3.6 shows the realized PPP-weighted private savings and investment ratios in the OECD, in proportion to the aggregate OECD GDP. The right panel shows that, under reasonable parameter values, the gap between desired saving and desired investment could have been well over 10% of GDP in the most recent period had interest rate not declined. This illustrates that the decline in AE R* had counteracted a large excess of private sector saving.

We now turn to the discussion of the forces driving the evolution of the real neutral rate in advanced economies.

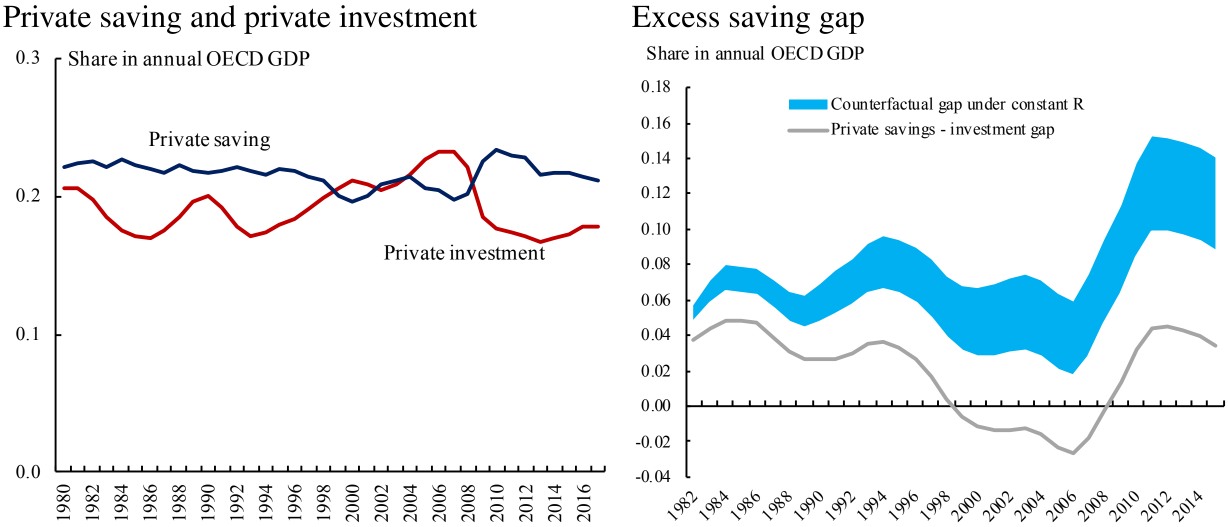

Source: IMF, OECD and authors' calculations.
Notes: The figure in the left panel shows PPP-weighted gross private saving and gross private fixed capital formation across the following countries: Australia, Austria, Belgium, Canada, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Iceland, Ireland, Israel, Italy, Japan, Korea, Latvia, Lithuania, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Kingdom, United States. The figure on the right shows the result of a simple counterfactual exercise where we calculate the private savings-investment gap under a scenario of no decline in the long-term interest rate since the 1980s. The swathe contains the counterfactual for values of responsiveness of S/I ratios to interest rates between 2 and 4.

Figure 3.6: Private saving and private investment: the level and gap

3.4 Government policy and the equilibrium interest rate

Over the past several decades, government policy in the developed world has shifted significantly in at least four respects (Figure 3.7). First, government debt has risen, from around 20% of GDP to around 70% (government consumption – excluding healthcare – remained relatively stable). Second, old-age payments administered through the social-security and healthcare systems have gone up, from 4% to 7% and from 2.5% to 4% of GDP, respectively, accounting for the lion share of the increase in total social spending (Figure 3.8). Third, significant changes to tax policies have taken place. The effective corporate tax rates in the rich economies have fallen, from around 32% at the turn of the century to 24% more recently.[12] Wealth taxes, operational in 12 OECD countries in 1990, remain in place only in 4 counties today (OECD (2018)). And, as documented by Thomas Piketty and Emmanuel Saez (2007), the overall progressivity of the tax system has decreased in some jurisdictions – notably in the United States and the United Kingdom. Finally, in some countries, the public provision of credit has increased: in the United States, for example, such provision amounts to around 7% of government debt.

These shifts are likely to have had a profound impact on the economy in general, and on the equilibrium rate of interest in particular. Specifically, all of these shifts – perhaps with the exception of tax changes[13] – are likely to have pushed interest rates *higher* over the past 30

[12] See http://www.oecd.org/tax/tax-policy/fig7-avg-statutory-tax-rates-by-region-large.png for details.

[13] Smaller incidence of wealth taxation likely reinforced the incentives to accumulate assets, particularly at the top of the wealth distribution. The magnitude of this effect will be governed by the elasticity of capital supply. The nascent but active literature puts these elasticities in the range of -3 to -20 (a one percentage point increase in wealth tax leads to a 3-20% change in wealth). See Zoutman (2015), Jakobsen, Jakobsen, Kleven, and Zucman

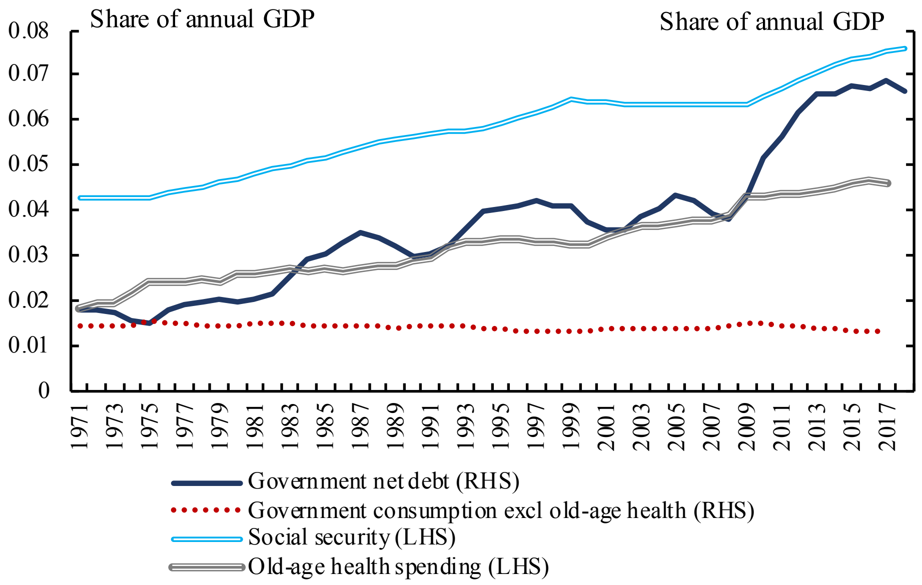

Notes: The Figure shows OECD aggregates in proportion to total OECD GDP. All data are from the OECD. Government net debt line measures the general government net financial liabilities, from the OECD Economic Outlook Database. It includes net government debt held by the public and also other net liabilities of the government. For example, in the United States in 2017, the net financial liabilities as reported by the OECD were 80% of GDP, while net debt held by the public was 75% of GDP. Government consumption figures represent the general government final consumption expenditure, adjusted by subtracting the old-age health spending (note that this series excludes the social security transfers by default). The old-age health spending is calculated as the aggregate health spending on ages 65+. The overall health spending figures are from the OECD/WHO statistics on sources of funding for healthcare. They include healthcare financed directly by the government and from the compulsory schemes. The old-age share is then calculated under the assumption that 60% of total health spending is directed at the older demographic groups, consistent with the evidence available for a handful of OECD countries.

Figure 3.7: Advanced economies government policy ratios (in proportion to GDP)

years. In this and the next section we turn to the analysis of the impact of these policy shifts on the natural rate, with the ultimate goal to inform the counterfactual ‘pure’ R* that would prevail without government intervention.

We focus on government debt, social security and healthcare spending, leaving the formal analysis of the impact of tax changes for future work. We find that shifts in government policy have likely pushed equilibrium rates of interest up by a significant amount over the period in question. As a rough rule of thumb on the magnitudes involved, our analysis suggests that the tripling of the government debt over the past half century raised rates by 1.5pp, while the expansion of social spending of around 5% of GDP added a further 2.5pp. While the precise magnitudes of these multipliers are subject to substantial model and statistical uncertainty, the qualitative conclusion is clear: had the public policy not responded, the advanced world’s equilibrium rate would likely be deeply negative.[14]

(2018), Brülhart, Gruber, Krapf, and Schmidheiny (2016) and Seim (2017).

[14] A corollary of this link between government debt and interest rates is that a higher value of public debt, compared to market expectations, is likely to raise the natural interest rates. For analysis of this argument, see Kocherlakota (2015).

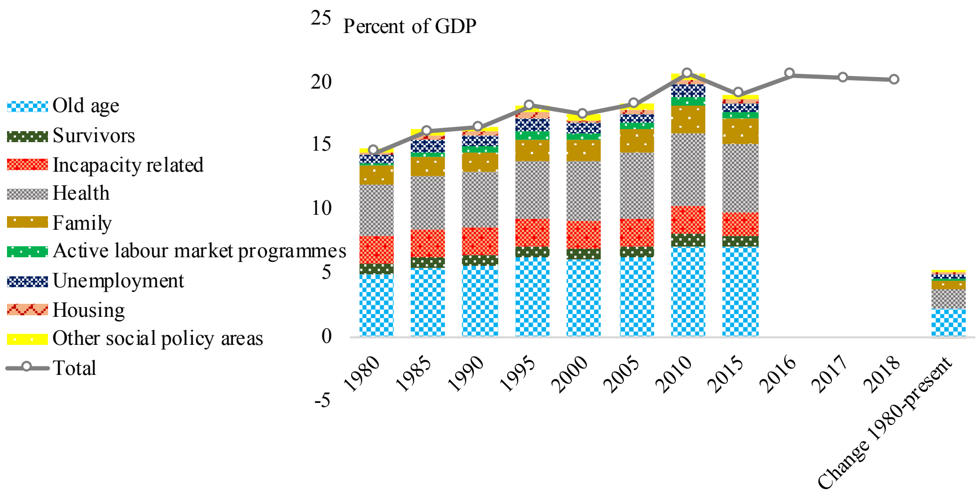

Notes: Data are from the OECD Social Expenditure Database (SOCX). The data for 2016-2018 are preliminary and the breakdown by function is not yet available.

Figure 3.8: Public social spending in the OECD

3.4.1 A brief review of the theoretical arguments

We begin by reviewing the effects of government policy on the equilibrium interest rate, focusing on government borrowing, as this has been the main subject of the large literature in macroeconomics which we can draw on. We next describe the key takeaways from three classes of models: the flow-based IS/LM, the neoclassical growth model, and a range of incomplete markets / heterogenous agents models.

In the IS/LM model the interest rate is determined in the flow equilibrium of goods and financial markets. It is thus the flow of government spending and the flow of budget deficits that matter for the determination of the interest rate. A higher deficit – operating through higher government spending or lower net taxes – shifts the IS curve to the right, raising interest rates in the short run. To the extent that in the medium-to-long-run the economy faces a vertical aggregate supply curve, the model predicts that the price level will tend to drift upwards as the economy operates above its potential. This reduces the real money stock (shifting the LM curve to the left) and thus further raises the interest rate until the long-run equilibrium level of output is restored.[15] This model thus predicts that in the new equilibrium (with *permanently* higher deficits) the real interest rate is higher. However, a *temporary* increase in deficits will only have a temporary effect on real rates.

In the canonical neoclassical model with complete markets and infinitely-lived agents, Ricardian Equivalence holds and neither deficit nor debt are relevant, as the representative household can fully offset the changes to government's borrowing policy through its saving decisions. Thus independent shocks to government borrowing alone have no effect on the equilibrium interest rate. The neoclassical model instead emphasizes the link between the stock of capital and the interest rate: in equilibrium $r = f'(k) - \delta$.[16] Thus government policies affect the interest

[15] In a richer model, or indeed in the real world, this may be brought about by tightening of monetary policy.

[16] On the balanced growth path, the level of effective capital stock adjusts such that the interest rate simultaneously satisfies the balanced growth version of the representative household's Euler Equation: $r = \frac{1}{IES} \cdot g + \theta$

rate only to the extent that they impact on the stock of private capital.

A change in government *spending* can affect private sector accumulation of capital in the neoclassical economy. But the model suggests that such policy will raise the real interest rate only temporarily (a classic reference is Baxter and King (1993)). An increase in government spending constitutes a negative wealth effect for the representative consumer: it tightens the resource constraint. The consumer responds by reducing consumption and leisure and by raising labor supply. The marginal product of capital increases, driving up investment. In a new steady state, both capital and labor supply are higher, but their ratio – and hence the interest rate – are unchanged. The interest rate is higher only along the transition path. The model suggests that this can be quite persistent, with half-life of about 6 years in the classic calibration of Baxter and King (1993).

Finally, in the micro-founded modern macroeconomic models that depart from the representative agent and complete markets assumptions, Ricardian Equivalence does not hold, and government transfer policies affect the equilibrium allocations through several distinct channels.

First, the intertemporal transfers – that is, *redistribution across time* – matters if peoples' planning horizons are finite. This could be because of finite lives coupled with less-than-perfect bequest motive, as in the seminal models of Peter Diamond (1965) and Olivier Blanchard (1985), or perhaps due to time-dependent preferences and myopic behavior pioneered by David Laibson (1997). The reason is intuitive: with finite planning horizon, agents currently alive expect to shoulder only a part of the financing burden that comes with today's transfer; the rest is to be serviced by future generations. Such transfers thus affect agents' wealth and their consumption and saving plans.

Second, *transfers across agents* can affect aggregate consumption and saving (and hence the interest rate) if agents have different marginal propensities to consume (MPCs). Differences in MPCs could arise because of several distinct features of the economic environment. They could be a result of uninsurable risks and binding borrowing constraints, as in the works of Rao Aiyagari and Ellen McGrattan (1993, 1998) and the model of Hyunseung Oh and Ricardo Reis (2012). They could emerge because some agents have little to no liquid wealth, preventing them from adjusting their consumption, as in the paper by Greg Kaplan, Giovanni Violante and Justin Weidner (2014). Another reason may be the life-cycle: propensity to consume may differs between workers and retirees, as in Gertler (1999), or may vary with age as in Gagnon, Johannsen, and Lopez-Salido (2016) and Eggertsson, Mehrotra, and Robbins (2019). Heterogenous MPCs and distortionary taxes deliver this result in the savers-spenders model of Gregory Mankiw (2000).[17] In all those models, government transfers from a low-MPC agent to a high-MPC agent will boost the aggregate desire to consume and lower desired savings, thereby raising the interest rate.

The third way in which government policy affects interest rates is what may be called a *pre-*

where IES is the intertemporal elasticity of substitution and θ is the rate of time preference.

[17] Interestingly, in the savers-spenders model of Mankiw (2000), if taxes are levied lump-sum, a deficit-financed transfer that permanently increases the level of debt does not affect the stock of capital or the interest rate in the long run. The reason is that the interest rate is pinned down by the savers, who are infinitely lived and Ricardian.

cautionary saving channel. One facet of this channel is that government policies can directly reduce the risks faced by the agents. The mechanism is close to the one analyzed by Engen and Gruber (2001). Under imperfect insurance, agents who face some idiosyncratic risks – for example those related to health or unemployment – attempt to self-insure through saving. This precautionary saving motive acts to push the interest rate below the rate that would prevail in a complete markets economy (where all risks are insurable and so do not affect the agents' behavior). Government policies such as social insurance will affect the importance of precautionary saving: a stronger social safety net or higher unemployment and disability benefits curtail the associated risks, curbing the desire to save. Conversely, lack of social insurance means that agents need to rely on their own resources when experiencing hardship, making personal saving a priority. However, as illustrated in Figure 3.8, the overall size of the social safety net across the OECD has changed little over the period in question. We do not attempt to model it here, but leave it as an important direction to be explored in future work.

The other facet of the precautionary saving channel – and the one we focus on in this paper – works through the provision of assets and liquidity which agents use to insure themselves against shocks. This mechanism is at the heart of Aiyagari and McGrattan (1998) and has recently been discussed in the context of secular stagnation in Caballero, Farhi, and Gourinchas (2016) and Caballero and Farhi (2014). The intuition we have in mind is simple: a rise in government debt raises the overall supply of assets in the economy, which, all else equal, pushes interest rates up. Indeed, there is evidence in the data that government debt constitutes a non-trivial proportion of the total investable financial assets in the developed world, so that this channel can have a quantitative bite. The estimates of the share of government bonds in total financial assets range from one-third in the US to two-thirds in Japan (Kay (2015)).

An important issue in the context of these channels is the quantitative easing (QE) policy pursued by advanced economy central banks over the crisis period. It is useful to distinguish between economically two distinct kinds of QE. The first kind encompasses policies that swap risky assets for safe assets and includes policies such as QE1, LTRO, and many other lender of last resort central bank interventions. The second kind is a policy whereby the central bank issues reserves to buy risk-free debt. The policies that belong to the first kind can indeed alleviate safe asset shortage during a crisis, as shown in the formal model of the "safety trap" by Ricardo Caballero and Emmanuel Farhi (2018). Such policies would thus raise the *short-run* natural rate of interest. This logic captures an important transmission channel of these unconventional tools. It is not, however, the focus of our paper, which focuses on the *long-run* real neutral interest rate – that is, the interest rate that prevails once cyclical conditions and policies have washed out.[18] Indeed, the composition of the Fed balance sheet shows that, in the United States, this first kind of QE was very short-lived, with the amount of risky assets on the Fed balance sheet scaled back quickly as the worst period of the crisis has passed. The second kind of QE proved much more persistent, with central banks balance sheets still elevated and close to their peaks today. But QE of the second kind constitutes primarily a maturity transformation of government

[18]The reader may want to think of the short-run equilibrium real rate as fluctuating cyclically around the long-term real rate.

debt, rather than change in the total availability of investable assets. Consequently, through the lens of our model (or the model of Caballero and Farhi cited above) such policy would not have any effect. It is possible to write down richer frameworks, notably ones where various assets have differing degree of liquidity. In those frameworks QE can indeed have (potentially large) real effects through changing the liquidity composition of households asset holdings. But the literature on this is still in its infancy (Cui and Sterk, 2018).

In summary, macroeconomic theory developed over the past couple of decades enriched the basic model of Frank Ramsey and Robert Barro (Barro (1974)) with several channels that make the government policy a relevant determinant of the long-term interest rate. We now turn to the empirical evidence that has been accumulated in parallel to these theoretical advances.

3.4.2 Empirical evidence on the link between government policy and long-term interest rates

The main challenge when estimating the effect of government borrowing on interest rates is the large number of potentially confounding factors which may make simple regressions of interest rates on debt spurious and uninformative. Shifts in both desired saving and desired investment brought about by secular trends unrelated to fiscal policy will affect the interest rate and will likely be correlated with headline measures of government borrowing. Unaccounted for, these factors will introduce an omitted variable bias into econometric estimates.

We shall not attempt a full-blown empirical assessment in this paper, and instead present a summary of the empirical estimates in the literature. For an interested reader, Appendix 3.B illustrates several challenges of estimating the causal relationship between equilibrium interest rates and government debt through a simple empirical exercise for the US, Canada, Euro Area and the UK. These challenges include the presence of international capital flows and of endogenous responsiveness of policy to excess of private saving over private investment, both of which are likely to attenuate the individual-country estimates of the impact of deficits on interest rates. With that caveat, we now present the estimates from a broad literature that attempted to deal with these and other confounding factors in finding the link between government finances and real rates.

Several key studies in the empirical literature focused on the United States. In a chapter of the *Handbook of Macroeconomics* at the turn of the century, Douglas Elmendorf and Gregory Mankiw (1999) reviewed the theoretical and empirical literature on the Ricardian Equivalence proposition, concluding that, while the studies that attempted to estimate the impact of government finances on interest rates cannot reject the null hypothesis of zero impact, they suffer from lack of statistical power.[19] More recent work appears more conclusive. In their literature review of this topic, William Gale and Peter Orszag (2002) conclude that the effect of

[19] They write of the literature that tends to find close to zero effect of government deficit on rates: *"Our view is that this literature [...] is ultimately not very informative. [...] Plosser (1987) and Evans (1987) generally cannot reject the hypothesis that government spending, budget deficits, and monetary policy each have no effect on interest rates. Plosser (1987) also reports that expected inflation has no significant effect on nominal interest rates. These findings suggest that this framework has little power to measure the true effects of policy."*

government deficit on the real rates is positive and economically significant: a 1pp increase in the deficit-to-GDP ratio tends to raise interest rates by around 50-100bps. And the two most authoritative contributions on the topic suggest estimates that are significant, albeit somewhat smaller. Thomas Laubach (2009) studies how forward rates on government securities react to news in CBO's fiscal forecasts. The identifying assumption in his work is that long-term rates and forecasts are not contaminated by current events and shocks at the business cycle frequency. According to his estimates, a rise in government deficit of 1pp of GDP raises interest rates by about 20-30bps; an equal increase in debt/GDP ratio results in a rise of about 3-4bps. He asserts that these flow- and stock-multipliers are broadly consistent, because of the autocorrelation of the deficits observed in the data.[20] Another important contribution to this literature is that of Eric Engen and Glenn Hubbard (2004), who consider a host of specifications linking interest rates or changes in interest rates to government debt or to the deficit, both contemporaneously and in a forward looking setting. Their results suggest that a 1pp rise in government debt / GDP pushes interest rates up by about 3bp, broadly in line with Laubach's findings.[21]

Further evidence is available for advanced economies beyond the United States. In an international setting, Anne-Marie Brook (2003) documents that the range of estimates of the effect of a 1pp increase in government debt/GDP ratio on interest rates is 1-6bps, with the corresponding range for a 1pp increase in deficits in the region of 20-40bps.[22] In an important study of the Euro Area, Riccardo Faini (2006) finds that a 1pp rise in deficits at the Euro Area level raises long-term rates by around 40bps, close to – and if anything, higher than – the US multipliers. Considering an even wider panel of 19 OECD economies spanning 1971-2004, Noriaki Kinoshita (2006) finds that the effect of a 1pp rise in government debt-to-GDP ratio is to raise interest rates by 4-5bps.

A complementary way to assess the size of these effects is to consider simulations from large-scale models used for quantitative analysis in policy institutions. Because these models are carefully estimated using real-world data, they should be able to provide a steer as to the size of the effects. A well known example is the FRB/US model, used and maintained by researchers at the Federal Reserve Board (Laforte and Roberts, 2014). In a recent speech, Stanley Fischer (2016) uses this model to estimate the impact of a persistent increase in deficit on real rates, and finds that a 1pp increase in deficit raises the equilibrium rate by between 40 and 50bps, depending on whether the deficit increased because of a tax cut (smaller effect) or a rise in government spending (larger effect). These figures are thus slightly larger than the empirical estimates cited above.

In summary, the estimates in the literature paint a fairly consistent picture: a 1pp rise in deficit tends to raise interest rates by around 30-40bps; while a 1pp rise in debt/GDP ratio results

[20]Specifically, he estimates the autocorrelation of 0.83, implying that the 1pp rise in the deficit should have $\frac{1}{1-0.83} = 6$ times the effect of a 1pp rise in debt – broadly in line with what he finds.

[21]The results vary across different specifications, highlighting that the precise econometric details matter for the conclusions of this line of empirical research.

[22]For the Euro Area, several older papers focused on the impact of fiscal policy on government bond spreads rather than interest rates. These papers tend to find smaller effects: on average, 1pp increase in deficit-GDP ratio raises spreads by around 10bps.

Study	Country / region	1pp increase in deficit/GDP	1pp increase in debt/GDP
Gale and Orszag (2002)	US	50-100bps	-
Laubach (2009)	US	20-30bps	3-4bps
Engen and Hubbard (2004)	US	18bps	3bps
FRB/US model	US	40-50bps	-
Faini (2006)	Euro Area	40bps	-
Brook (2003)	Advanced economies	20-40bps	1-6bps
Kinoshita (2006)	19 OECD economies	-	4-5bps
Average		*38bps*	*3.5bps*

Table 3.2: Impact of government borrowing on the interest rate: summary of the literature

in an increase of about 3-5bps (Table 3.2). We suspect this figure is an underestimate of the impact of an exogenous increase in budget deficits on real rates because fiscal expectations are measured with error, because any one country can import capital and so attenuate rate increases when budget deficits increase, and because there will be a tendency – as fiscal policy is used to stabilize the economy – for periods of low neutral real rates to coincide with periods of expansionary fiscal policy.

3.4.3 The historical impact of government borrowing on R*

The elasticities identified in the empirical work combined with the historical path of government borrowing give simple back-of-the-envelope estimates of the historical influence of fiscal policy on real interest rates. Over the past 40 years, the increase in government debt in the OECD has likely pushed interest rates higher, perhaps by as much as 1.5pp. The measure of R* that excludes the impact of public debt hovered around zero since the early 2000s, and remains negative at the moment (Figure 3.9).

The low levels of the private sector equilibrium rates identified here are consistent with some other estimates, for example with the recent work by Pierre-Olivier Gourinchas and Helene Rey (2016), who study the behaviour of the *private* consumption-to-wealth ratio. They conclude that real rates may be around -2% over the next decade. This is broadly in line with our adjusted measure in Figure 3.9. Our estimates are also closely aligned with those of Eggertsson, Mehrotra, and Robbins (2019).

To develop further intuition and to consider other mechanisms through which public policy may have affected the interest rate, we now turn to a complementary approach: a general equilibrium modeling framework.

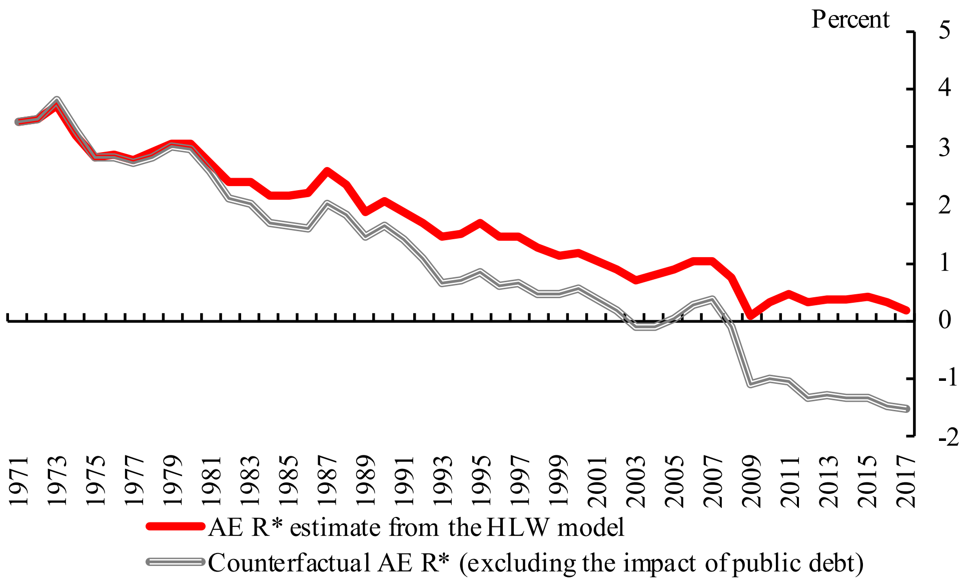

Note: The figure shows the estimated equilibrium real interest rate in advanced economies, and an adjusted measure that subtracts the impact of government borrowing using the average elasticities reported in Table 2.

Figure 3.9: Advanced economies R* adjusted for the impact of government debt

3.5 Government policy and R*: a model-based assessment

3.5.1 Two general equilibrium models

In Section 3.4.1 we outlined various channels through which government debt may affect the equilibrium real interest rate; our goal in this Section is to illustrate their quantitative importance within a general equilibrium framework. We want our approach to be simple and transparent, providing a credible complement and a cross-check to the empirical analysis above.

To achieve these goals, we build not one but two general equilibrium models: one capturing the finiteness of life and life-cycle heterogeneity, and another which focuses on precautionary behavior. We judge that two simpler models may be better than one complex one. Our approach allows for a coherent assessment of the importance of different channels while remaining clear and easy to understand and replicate.[23]

The first model, which builds closely on Gertler (1999), highlights life-cycle heterogeneity. In this economy, ex-ante identical individuals are at different points in their lives: some are working, some are already retired. This drives the differences in their consumption and saving behavior. The framework is similar to that of Blanchard (1985) and Yaari (1965) – individuals face constant probability of death and so their horizons are finite – but, in addition to their model, workers retire and finance consumption with savings until death.

The second model is a Bewley-Huggett-Aiyagari economy with incomplete markets and uninsurable income risk at the level of an individual household. A similar model was considered by Aiyagari and McGrattan (1998) who also studied the role of government debt on the equilibrium allocation in presence of idiosyncratic risk. The main differences between ours and

[23]We leave combining the two models for future work.

their approach are that: (i) we calibrate the risk component of the income process to deliver a realistic dose of uncertainty, which implies that distributions of income and assets in the model broadly match distributions observed in developed economies such as the United States;[24] (ii) we cast the model in continuous time, taking full advantage of the recent analytical and computational discoveries in macroeconomics.

Here we sketch the main workings of the two models and develop the intuition; a more detailed description of the models is available in Appendix 3.C for the life-cycle model and Appendix 3.D for the incomplete market model.

3.5.2 Model of finite lives and life-cycle heterogeneity

3.5.2.1 Demographics and preferences

There are two stages of life, work and retirement, with exogenous transition probabilities. That is, each worker faces a given probability of retirement $1-\omega$, and, once a retiree, a given probability of death $1-\gamma$. Population grows at a gross rate $1+n$.

There is no aggregate risk; the only sources of uncertainty facing an individual are the risk of retirement while a worker (associated with a loss of labor income) and a risk of death while a retiree. Left unchecked, these sources of risk would affect agents behavior. This would make aggregation problematic, and, more importantly, it would be unrealistic: timing of retirement is, for the most part, known. To deal with this unrealistic feature, we assume that there are perfect annuity markets for the retirees (neutralizing the influence of the risk of death on their behavior), and that workers' preferences have a certainty equivalence property (such that the *risk* of retirement does not affect workers' behavior in equilibrium).[25] These two assumptions are both realistic and convenient, in that they allow for the derivation of the aggregate consumption function, as we illustrate momentarily.

Specifically, we assume that agents have recursive Epstein-Zin preferences defined as follows:

$$V_t^z = [(C_t)^\rho + \beta^z \mathbb{E}_t\{V_{t+1} \mid z\}^\rho]^{1/\rho} \tag{3.7}$$

where C_t denotes consumption, V_t^z and β^z stand for agent's $z \in \{w, r\}$ value function and the discount factor respectively, and $\sigma = \frac{1}{1-\rho}$ is the intertemporal elasticity of substitution.

Retirees and workers differ in two crucial respects. First, they have different discount factors. Because of the positive probability of death facing any retiree, their discount factor is the time preference parameter β multiplied by the probability of surviving into the next period:

$$\beta^w = \beta$$

[24] We match the degree of income inequality in the data, but fall short of matching the extreme degree of wealth inequality observed in the real world. We discuss the (standard and well-known) reasons why this is so below.

[25] In particular, workers are assumed to have recursive Epstein and Zin (1991) preferences that generate certainty equivalent decision rules in the presence of income risk.

$$\beta^r = \beta \cdot \gamma.$$

Second, the expectation of the value function next period differs between a worker and a retiree. In particular, a worker takes into account the possibility of retiring, so that her expectation of the value function next period is a probability-weighted sum of the values in the two states:

$$E_t\{V_{t+1} \mid w\} = \omega V^w_{t+1} + (1-\omega) V^r_{t+1},$$

while the expectation of the value function of a retiree is simply given by

$$E_t\{V_{t+1} \mid r\} = V^r_{t+1}.$$

We now outline the problems of the two types of agents.

3.5.2.2 Retirees

Retirees consume out of savings and social security payments. Each period, some retirees die. We make the assumption – standard in the literature – that those who survive receive the proportional share of the proceeds. This means that the effective return faced by individual retirees is R_t/γ, higher than the ongoing interest rate R_t.[26]

Because probability of death is independent of age and the government does not discriminate across retirees in its social security transfer policy, each retiree (irrespective of age) solves an identical problem, which is:

$$V^r_t = \max_{C^r_t} [(C^r_t)^\rho + \beta\gamma \mathbb{E}_t\{V^r_{t+1}\}^\rho]^{1/\rho}$$

subject to the flow budget constraint:

$$A^r_{t+1} = (R_t/\gamma) A^r_t - C^r_t + E^r_t,$$

where A^r_t stands for retiree's assets, C^r_t is her consumption expenditure, and E^r_t is the social security and healthcare cost transfer.[27]

3.5.2.3 Workers

Individuals are born workers and have no assets at the start of life. They consume out of asset wealth and their labour income net of taxes. Because of the demographic structure (in particular the assumption that probability of retirement is independent of age[28]), worker's problem is

[26] For retirees *as a group* wealth accumulates at the interest rate R_t, as the higher individual return cancels out with some retirees dying.

[27] Our modeling of healthcare provision is very simple – we treat old-age healthcare cost as a lump-sum transfer, subsumed in the variable E.

[28] Of course this is an unrealistic assumption. But, as explained above, the effect of this assumption on workers' behavior is neutralized through the structure of preferences which exhibit a certainty equivalence property. The role of this assumption is thus only to simplify the model and achieve aggregation, with little cost to the economics.

effectively the same no matter the age. Each worker solves:

$$V_t^w = \max_{C_t^w}\{(C_t^w)^\rho + \beta[\omega V_{t+1}^w + (1-\omega)V_{t+1}^r]^\rho\}^{1/\rho}$$

subject to:

$$A_{t+1}^w = R_t A_t^w + W_t - T_t - C_t^w,$$

where T_t are lump-sum taxes levied by the government.[29]

3.5.2.4 Firms

The supply side of the model is extremely simple. Market are competitive. Production is carried out by firms employing capital and labor. The aggregate production function is

$$Y_t = K_t^\alpha (X_t N_t)^{1-\alpha}$$

where N_t is the number of workers in the economy. There is exogenous technological progress and population growth, that is $X_{t+1} = (1+x)X_t$ and $N_{t+1} = (1+n)N_t$. Perfect competition in factor markets means that the wage and the rental rate are equated to the marginal products of the factors: $W_t = \alpha\frac{Y_t}{N_t}$ and $R_t = (1-\alpha)\frac{Y_t}{K_t} + (1-\delta)$. Capital evolves according to the standard law of motion: $K_{t+1} = Y_t - C_t - G_t + (1-\delta)K_t$.

3.5.2.5 Government

The government consumes G_t each period, and pays retirees a total of E_t in social security and healthcare benefits. To finance its expenditures the government levies a lump sum tax T_t on the workers. It can also issue one period government bonds B_{t+1}. The government flow budget constraint is:

$$B_{t+1} + T_t = R_t B_t + G_t + E_t.$$

Iterating forward gives the intertemporal budget constraint of the government:

$$R_t B_t = \sum_{v=0}^{\infty} \frac{T_{t+v}}{\prod_{z=1}^{v} R_{t+z}} - \sum_{v=0}^{\infty} \frac{G_{t+v}}{\prod_{z=1}^{v} R_{t+z}} - \sum_{v=0}^{\infty} \frac{E_{t+v}}{\prod_{z=1}^{v} R_{t+z}}.$$

That is, the difference between the present discounted valued of government revenue and spending must be exactly equal to the current value of the outstanding debt.

[29] There are two key channels through which life-cycle considerations affect workers' behaviour. First, a worker takes into account the fact that with probability $1-\omega$ she becomes a retiree. This means that, relative to the representative agent case, she discounts the future stream of wages by more: effectively, this is the *saving for retirement* effect. Mechanically, a larger discount rate reduces the value of human wealth in the consumption function, thus leading to lower consumption and higher saving. Second, a worker discounts the future stream of wealth more because she anticipates that inevitably there will come a time when she becomes a retiree, facing the sad truth that her life is finite. With finite life, wealth can be smoothed out across fewer periods, so its marginal utility value is lower. This effect shows up as a higher effective discount rate applied to future wealth.

Government policy is exogenous. In particular, it is characterized by the four ratios, $\bar{g}_t, \bar{b}_t, \bar{e}_t, \bar{h}_t$, of government consumption, debt, social security and healthcare spending to GDP, respectively:

$$G_t = \bar{g}_t Y_t$$

$$B_t = \bar{b}_t Y_t$$

$$E_t = (\bar{e}_t + \bar{h}_t) Y_t$$

Given the paths of G_t, E_t and B_t, taxes adjust to satisfy the intertemporal budget constraint.

3.5.2.6 Equilibrium

In this economy, markets are competitive and agents take prices as given. Formally, a competitive equilibrium is a sequence of quantities and prices such that (i) households maximize utility subject to their budget constraints, (ii) firms maximize profits subject to their technology constraints, (iii) the government chooses a path for taxes, compatible with intertemporal solvency, to finance debt, spending and transfers, (iv) all markets clear.

Appendix 3.C contains the details of the derivation of the equilibrium conditions of the model. The individual policy functions within the two groups – workers and retirees – aggregate up nicely. Aggregating the two consumption levels, we derive the aggregate consumption function:

$$C_t = C_t^w + C_t^r = \pi_t \left\{ (1 - \lambda_t) R_t A_t + H_t + S_t^w + \epsilon_t (\lambda_t R_t A_t + S_t^r) \right\}.$$

In this consumption function, π_t denotes each worker's marginal propensity to consume out of wealth, and $\pi_t \epsilon_t$ is the MPC of each retiree. These MPCs multiply the total wealth of each group of consumers (with a slight abuse of notation, A_t now denotes aggregate financial wealth, H_t is aggregate human wealth (the net present value of future wages), and S_t stands for the aggregate value of social security and healthcare payments). Compared to a standard model, the only additional state variable is the share of wealth held by retirees, λ_t, which fully captures the heterogeneity in the economy.

The total supply of assets is the sum of capital stock K_t and government debt B_t so that the equilibrium requires:

$$A_t = A_t^w + A_t^r = K_t + B_t,$$

i.e. households asset demand equals the asset supply.

3.5.2.7 Calibration and the initial steady state of the life-cycle model

Despite the richness of the economics, the model is parsimonious and relatively straightforward to calibrate. We set the preferences and technology parameters at the standard values in the macro literature (Table 3.3). The growth rate of technological change, the demographics parameters and the government policy ratios are all calibrated to match the data in advanced

Parameter	Description	Calibration
Preferences and technology		
β	Discount factor	0.98
σ	Intertemporal elasticity of substitution	0.5
α	Capital share	1/3
δ	Depreciation rate	0.1
x	Rate of technological change	1.51%
Demographics		
n	Gross population growth rate	1.35%
$\frac{1}{1-\omega}$	Average length of working life (years)	47.6
$\frac{1}{1-\gamma}$	Average length of retirment (years)	10.5
Government		
$\bar{b}$	Government debt / GDP	0.18
$\bar{g}$	Government consumption / GDP	0.14
$\bar{e}$	Social security spending / GDP	0.04
$\bar{h}$	Old-age healthcare spending / GDP	0.02

Table 3.3: Calibration

economies in 1970.

Because there is population growth and technological progress in this economy, the steady state equilibrium takes the form of a balanced growth path where all variables grow at a constant gross rate equal to $(1+n)(1+x)$. We can characterize the equilibrium by expressing all variables as ratios in units of effective labor (defining, for any variable Z_t, $z_t \equiv \frac{Z_t}{X_t N_t}$).

Table 3.4 shows the key variables along the initial (early-1970s) balanced growth path. The interest rate is 4.5%.[30] As we pointed out above, the key feature of this economy is the heterogeneity in marginal propensities to consume between workers and retirees. Indeed, the endogenous MPC of retirees is over twice that of the workers'. The additional state variable λ – the ratio of retirees' wealth in total wealth – takes a plausible value of 16%. Ratios of aggregate consumption, investment, capital and assets to output also match the stylized facts from the data well.

3.5.2.8 The simulation exercise

We now explore how the model economy reacts to changes in government policy. We study four policy levers: government debt, government spending, old-age social security and healthcare transfers.

We carry out the following experiment. Starting the economy in the initial 1970s steady state, we feed the model with the policy profiles depicted in Figure 3.7. Once announced, the profile of these shifts is fully anticipated by the agents. Beyond the current date, we assume that

[30]With growth rate of 2.9% per annum, the economy is dynamically efficient.

Variable	Description	Value
ψ	Ratio of retirees to workers	0.19
R	Real gross interest rate	1.045
ϵ	Ratio of retirees' to workers' MPCs	2.01
π_w	Workers' MPC	0.06
π_r	Retirees' MPC	0.13
λ	Share of retirees' wealth in total wealth	0.17
y	Output	1.50
Ratios (in proportion to output):		
c	Consumption	0.57
c_r	Consumption of retirees	0.11
c_w	Consumption of workers	0.45
a	Assets	2.42
a_r	Assets of retirees	0.40
a_w	Assets of workers	2.03
h	Human capital	4.23
i	Investment	0.27
k	Capital	2.25
τ	Taxes	0.21
s	Social security wealth of the retireees	0.50
s_w	Social security wealth of the workers	0.91

Table 3.4: The 1970s steady state

future policy ratios remain constant at their 2017 values.[31] We then compute the transition path towards this new steady state.

Our focus is on the response of the interest rate to these policy shifts. Figure 3.10 contains the main result of this Section: the total response of the interest rate to the policy changes discussed above. This response is quantitatively large: according to the model, government policies pushed up on the equilibrium interest rate by around 3.5pp over the past 50 years. Moreover, the model suggests that further upward pressure is to be expected as the economy settles at the new steady state. All of the policies except government spending – which did not change much – play an important role. The final set of bars, labelled "Interactions", is the additional effect on the interest rate from the (non-linear) synergies between the three different policies.[32]

3.5.3 A model of precautionary savings

We now turn to the model of precautionary behavior, which is a continuous time version of the Aiyagari and McGrattan (1998) economy. Population consists of a large number of infinitely-lived individuals of measure 1. Every individual is ex-ante identical, but people face shocks

[31] This is a conservative assumption, as one may reasonably expect the upwards drift in both debt and social security spending to continue, at least for some time.

[32] More precisely, the interaction effect exists because the final steady state is a non-linear system of equations. These non-linearities make the overall effect of several exogenous changes different, in general, from the sum of the parts.

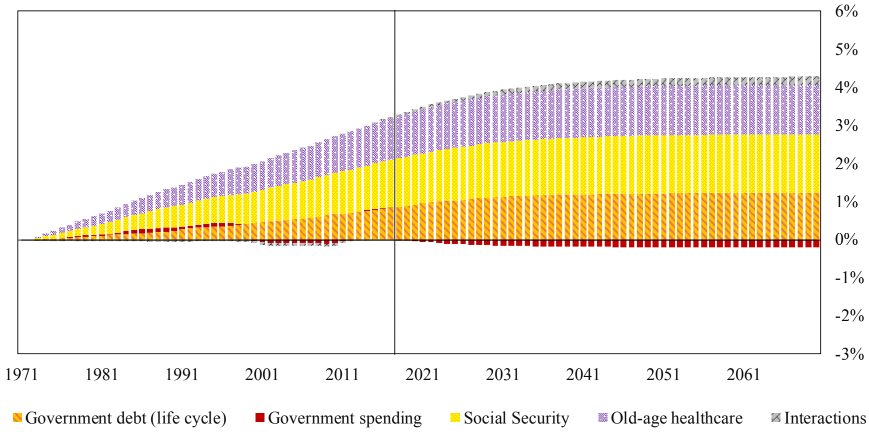

Notes: The figure shows how the equilibrium real interest rate adjusts to the exogenously given paths of government debt, spending, and old-age social security transfers depicted in Figure 3.7. The 2017 values in Figure 3.7 are assumed to be the new steady state values.

Figure 3.10: Simulated impact of government policies on the equilibrium real interest rate in the life-cycle model

to their income which they cannot fully insure against: markets are incomplete. As a result of this idiosyncratic risk, individuals experience different income histories and thus accumulate different levels of wealth. All the risk is at the individual level: for simplicity, we abstract from aggregate uncertainty.

Our goal here is to assess quantitatively the influence government debt has on precautionary behavior. In other words, how different is the prevailing interest rate when government debt-to-GDP ratio is 18% vs. when it is 68%?[33]

3.5.3.1 Brief outline of the model

An individual chooses consumption and asset holdings to maximize her expected utility, subject to the flow budget constraint, the consumption non-negativity constraint, the borrowing constraint and a realization of the idiosyncratic income shock:

$$\max_{\{c_t\}_{t\geqslant 0}} \mathbb{E} \int_0^\infty e^{-\rho t} u(c_t) dt$$

[33] Our model is highly stylized and abstracts from important features present in more advanced and larger models in the literature. We view our model here as an early attempt to quantify the precautionary saving channel of government debt. Richer features may usefully be incorporated in future attempts to answer this question. For analysis of saving rates across the distribution, see Straub (2017) and Fagereng, Holm, Moll, and Natvik, (forthcoming). For evidence on the differential rates of return, see Fagereng, Guiso, Malacrino, and Pistaferri (2016). For models with multiple assets or a more careful analysis of the constraint – both of which contribute to a better match to the empirical distribution around the borrowing constraint, see Kaplan, Violante, and Weidner (2014), Kaplan, Moll, and Violante (2018) and Achdou et al. (2017). For the state-of-the-art calibration of the income process, see Guvenen, Karahan, Ozkan, and Song (2015). We conjecture that a richer model with some of the above features would likely predict larger effects of government policy.

subject to

$$\begin{aligned}
\dot{a}_t &= (1-\tau)w_t e_t + (1-\tau)r_t a_t - c_t \\
c_t &\geqslant 0 \\
a_t &\geqslant \underline{a} \\
e_t &\in \{z_1, ..., z_n\}
\end{aligned}$$

where c_t is individual consumption, a_t are individual asset holdings (and $\dot{a}_t$ denotes the time derivative, i.e. saving), r_t is the real (net) interest rate, w_t is the wage, e_t is the idiosyncratic shock to household's productivity. The household cannot insure against that idiosyncratic uncertainty. The government levies a proportional tax rate τ on both labor and capital income.[34]

The supply side is identical to that in the previous model: production function is Cobb-Douglas and there is perfect competition in all markets. Government issues bonds and collects taxes to finance its consumption and transfers. The government budget constraint is

$$\dot{B}_t = G_t + r_t B_t - \tau(w_t + rA_t), \tag{3.8}$$

which says that the change in government debt is equal to the government funding gap: government consumption G_t plus interest payments $r_t B_t$ minus the tax revenue.

Appendix 3.D presents the definition and solution of the equilibrium of this economy.

3.5.3.2 Parametrization

We choose the values of the parameters in the precautionary savings model to match the typical values in the literature and to be broadly consistent with the life-cycle model above (Table 3.3). We set the capital share at $\frac{1}{3}$, the rate of time preference at 0.04, the depreciation rate at 10% and the IES at $\frac{1}{2}$.

The crucial new aspect of the calibration is the idiosyncratic income process. Intuitively, size and persistence of income shocks will determine the strength of the precautionary savings motive, the degree of inequality, and the proportion of households close to or at the borrowing constraint. These outcomes will in turn determine the potency of government financing policy.

In the real world, individual income varies over time for a host of reasons. We do not model these causes here. Instead, we make sure that the income process in our model reflects these uncertainties. Specifically, we follow Ana Castaneda, Javier Diaz-Gimenez and Jose-Victor Rios-Rull (2003) and Winter (2016) and calibrate the income process to match aggregate income inequality in the OECD. There are four productivity and income states:

$$e \in \{0.20,\ 0.55,\ 0.80,\ 5.43\}.$$

[34]The assumption of a proportional tax rate is natural in a model with income and wealth heterogeneity. With lump-sum taxation, the poorest households would find themselves unable to pay the tax bill. Note that even though the tax is proportional it does not distort the labor supply decisions as the labor supply is inelastic.

Note that to generate the high skewness of the income and wealth distributions observed in the data, individuals in the top state have income an order of magnitude larger than the rest. The corresponding matrix of Poisson intensities is

$$P = \begin{pmatrix} 0.07^{-} & 0.04 & 0.02 & 0.001 \\ 0.03 & 0.13^{-} & 0.01 & 0.001 \\ 0.001 & 0.08 & 0.09^{-} & 0.011 \\ 0.1 & 0.02 & 0.06 & 0.17^{-} \end{pmatrix},$$

where the values on the main diagonal marked with superscript $^{-}$ indicate the intensity of leaving the current state. For example, the first row indicates that individuals currently in the poorest state z_1 leave that state with intensity 0.07, and enter one of the other three states z_2, z_3, z_4 with respective intensities of 0.04, 0.02 and 0.001. Thus states 1 and 3 are the most 'absorbing', in the sense that the out-intensities are the lowest for these states.

Given this income process, the distributional outcomes in the equilibrium of our model are broadly in line with those observed in the data. In particular, the income Gini coefficient in the model is 0.32, close to the OECD average.[35] A question that remains is whether our calibration generates an individual income profile that is realistic. Castañeda, Díaz-Giménez, and Ríos-Rull (2003) compare the across-the-income-distribution mobility statistics implied by their model with those observed in the data, and claim success: the simple model does reasonably well in capturing the persistence moments. In particular, both in the model and in the data, the income process is highly persistent. Based on that result, we have some confidence that the simple income process does a decent job reflecting the individual income process and transition rates between different points in the income distribution.

3.5.3.3 Results

We now compare the two stationary equilibria of the model, one with the government debt / GDP ratio set at 18%, and another at 68%.[36] Here we only present the results most relevant to our focus on the government policy and the interest rate; for a further characterization of the two equilibria, in terms of policy functions and the distributions, see Appendix 3.D.

The key quantitative result in Table 3.5 is that an increase in public debt / GDP ratio observed in the data implies a real interest rate that is 66bps higher in equilibrium. We conclude that the baseline heterogenous agents incomplete markets model suggests that the precautionary saving channel from government policy to the real interest rate is quantitatively important.

[35]The wealth Gini delivered by the model is somewhat lower than that in the data (0.55 in the model vs. 0.8 in the data in the US), and the fraction of households at the constraint is also relatively low, at around 9% vs. around a third in the data. These shortcomings are well-known in the literature. Some of the features of the real world that drive the extreme concentration of wealth holdings, such as the systematic differences in the saving rates and the rates of return on investments across the wealth distribution, are missing from this model. Furthermore, the asset market and the constraint are modeled in a crude way.

[36]Analysis of the transition path is not crucial for our purposes and is beyond the scope of this paper.

	Low Debt Equilibrium	High Debt Equilibrium
Government debt / GDP	0.18	0.68
Government consumption / GDP	0.14	0.14
Average tax rate	0.35	0.36
Real interest rate	4.50	5.16
Private capital / GDP	2.56	2.40
Income Gini	0.32	0.32
Fraction of individuals at the constraint	0.09	0.09

Table 3.5: Equilibria in the precautionary saving model

3.5.4 Summary and discussion

In summary of this Section, our analysis underscores the importance of secular public policy shifts in accounting for changes in the equilibrium interest rate. The natural corollary of our findings is that government intra- and intertemporal transfer policy is, in principle, an effective tool that can affect equilibrium interest rates in the economy. Similar policy implications have been discussed previously by Kocherlakota (2015) and Caballero and Farhi (2014).

One objection to our analysis might be that economic agents – consumers, investors, firms etc. – may in fact be more Ricardian than we currently assume. Our response to this is three-fold. First, in Section 3 we presented a broad range of empirical evidence that is inconsistent with the Ricardian Equivalence proposition. Second, in our framework, Ricardian Equivalence does not hold despite fully rational expectations and no information asymmetries: indeed, it would be irrational to be Ricardian in the economy we describe. Third, and relatedly, the assumptions that lead to rejection of Ricardian Equivalence are rather natural (i. people retire; ii. people die; iii. some people are credit constrained; iv. some people face risks they find hard to insure). All those considerations make us comfortable with our assumptions that the Ricardian offset is imperfect.

At this point it is also useful to highlight that wide uncertainty bands surround our point estimates, including those coming out of the models discussed above. Like all theory models, these tools are built upon a set of uncertain assumptions, and as such are only rough approximations of reality – this is especially true for models as minimalistic and transparent as ours. Even abstracting from model misspecification, there is a wide range of plausible parameter values with which to calibrate these models. A different combinations of parameters will produce quantitatively different results. We come back to the robustness of our analysis in Appendix 3.E. Having said that, the combination of a range of empirical studies together with directional guidance from the theory suggest that there are strong reasons to conclude that the government policies we scrutinized here have put significant upward pressure on safe neutral real rate over the past several decades.

	Growth of 20+ population	Retirement age	Years working	Years in retirement
1970	1.4	67.6	47.6	10.5
1975	1.3	66.6	46.6	12.3
1980	1.2	66.1	46.1	13.4
1985	1.1	65.1	45.1	15.0
1990	0.9	64.7	44.7	16.1
1995	0.8	63.8	43.8	17.5
2000	0.7	63.6	43.6	18.6
2005	0.8	64.1	44.1	18.9
2010	0.7	64.8	44.8	18.8
2015	0.4	65.5	45.5	18.7
Projection:				
2020	0.2	66.1	46.1	18.6
2025	0.2	66.8	46.8	18.4
2030	0.2	67.5	47.5	18.3

Sources: United Nations and OECD.

Table 3.6: Demographic transition in advanced economies

3.6 Underlying weakness in R*

Our simulation analysis concluded that the major shifts in governments' policies over the past 50 years facilitated a significant transfer of resources from low-MPC to high-MPC individuals and permitted households to self-insure against idiosyncratic shocks. All else equal, added together these shifts would have pushed interest rates in advanced world up by around 4pp. But of course all else was not equal. In this Section we use our models to complete the picture by quantitatively assessing the impact of other secular forces – demographics, trend productivity growth and income inequality – on the long-term interest rate.

Table 3.6 documents the major demographic transition that has been underway in advanced economies for the past 50 years. Population growth in the developed economies has fallen rapidly in past decades, from around 1.4% per annum in 1970s to less than 0.4% today. This trend is expected to continue; in fact, the latest UN projections suggest that population in advanced economies will start shrinking around 2050. As population growth decelerated, life expectancy has gone up significantly, and retirement ages did not keep up. As a result, the average length of retirement is nearly twice what it was in the 1970s. This positive development carries significant implications for life-cycle budgeting and thus for the balance of desired saving and investment.

Another force that our life-cycle model is well-equipped to capture is the slowdown in the pace of long-run growth. Our modeling framework inherits the property shared by essentially all dynamic macroeconomic models, namely that the long-run equilibrium interest rate is linked to the expected future consumption growth. This relationship – the Euler Equation or the dynamic IS curve – is the result of intertemporal optimization of households, who choose how much to consume today vs. tomorrow (and hence determining the growth rate of their consumption) based on the interest rate. In general equilibrium, the expectations of future consumption growth in the long-run coincide with the expectations of TFP growth. Hence the theory suggests that real rates and expected productivity growth ought to be linked.[37]

This prediction of the theory is, however, more tenuous in practice. In an early contribution to this topic, Christopher Carroll and Lawrence Summers (1991) established that, across countries, consumption growth and income growth are tightly linked and follow each other, and that households with more steeply rising income profiles tend to save more, not less. These findings – inconsistent with the standard permanent-income hypothesis / life cycle model – have been rationalized in the literature with buffer-stock models of savings (whereby households face uncertain income process, similarly to our second model) and introducing consumption habits in household preferences.[38] While our models attenuate the link between interest rates and consumption choices in line with these findings, nonetheless we urge a significant degree of caution when interpreting the results on the link between TFP and R*. Our preferred interpretation is that the low interest rates today are chiefly a symptom of a demand-side problem. We return to this issue in the final section of the paper where we discuss policy implications.

These caveats notwithstanding, numerous studies – for instance, Adler et al. (2017) – reach the conclusion that trend growth rates of both productivity and of TFP have declined significantly in advanced economies, first in the early 1980s when TFP growth halved from about 2% to 1% per annum, and then again in the mid-2000s, and the macroeconomic models we use do suggest that such deterioration should have dragged on neutral real interest rates.

The third force we consider is the rise in income inequality, which has increased in the United States and many other advanced economies (Figure 3.11). Our second model is well suited to give us an estimate of this shift on the real rate of interest. To trace out the effects of rising inequality in this model, we recalibrate the income process in such a way as to match the increase in income Gini coefficient in the OECD since the 1970s. Our calculations implicitly assume that ex-post inequality is driven by larger variance of individual income shocks, which constitutes a source of additional uncertainty for individual workers. An alternative view is that the increase in inequality is a consequence of shifts more tightly linked to heterogeneity across households that is known ex-ante. The distinction is important because only the former kind of shift would lead to an increase in precautionary behavior: being predictable, the latter shift is not associated with heightened risk. There is a long-standing debate about the merits of the two formulations in the literature.[39] The recent work by Fatih Guvenen and co-authors has established the large departures of log-normality in the individual income changes: in particular, earnings changes display strong negative skewness and extremely high kurtosis. Important for

[37] In a representative agent, infinite-horizon economy, the Euler Equation takes a particularly straightforward form, whereby long-run consumption growth rate and the interest rate are linked linearly, with the coefficient equal to the intertemporal elasticity of substitution. Within our framework, that link is still there, although it is attenuated by finite horizons and borrowing constraints: intuitively, the interest rate is relatively "less important" in driving consumption growth, as other factors (such as possibility of death or credit constraints) come into play. This implies that a given change in the expectations of future consumption growth – driven by news about TFP, say – will require a *larger* response of the interest rate to restore equilibrium.

[38] See Deaton (1991), Carroll (1997) and Carroll, Overland, and Weil (2000) and the literature that followed.

[39] Classic references include Lillard and Weiss (1979), MaCurdy (1982) and Guvenen (2009).

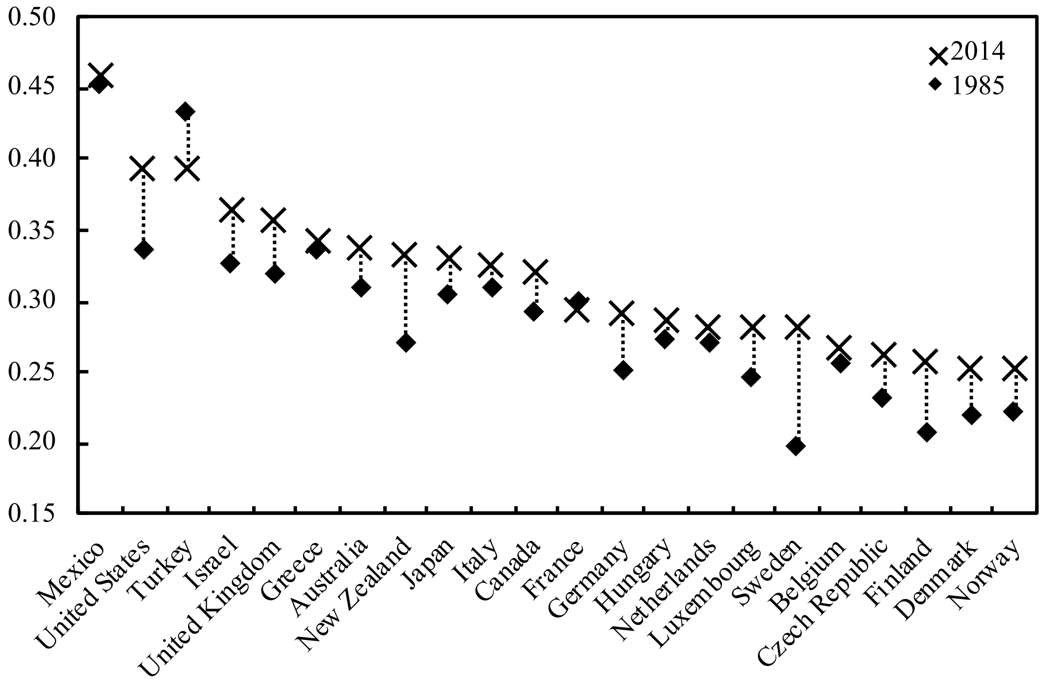

Source: OECD Database on Household Income Distribution and Poverty. Disposable income adjusted for household size.

Figure 3.11: Gini coefficient of disposable household income across the OECD

our interpretation is their finding that large shocks at the top of the income distribution tend to be very persistent. We view these results as supportive of the gist of our exercise, which interprets the increased disparity between the poor and the rich as going hand in hand with an increase in ex-ante uncertainty. Given the lack of consensus in the literature, it is possible that we overestimate the impact of inequality on real rates in this exercise. In any case, there likely are other powerful ways in which higher inequality has acted to depress rates, which we miss from our framework (and which we discuss momentarily).

To explore the implications of these trends for the equilibrium real interest rate, we perform the following exercise: in the life-cycle model, we calibrate the changes in demographic transition probabilities, ω and γ, to match the trends depicted in the final two columns of Table 3.6. We then feed in the series for population and TFP growth rates to match the evidence in the first column of Table 3.6 and Adler et al. (2017). We use the UN demographic projections to inform the path of demographics out to 2050, and assume that the terminal 2050 values are the steady state. We do not have a strong prior as to the path for future TFP growth, and we are well aware of the wide range of existing and plausible views. Aiming for a scenario that reflects the mode of these expectations, we assume that the TFP growth rate picks up from around zero in the latest available data to 0.7% in the long-run.[40] This pick-up in TFP growth is broadly in line with the CBO's assumption for the pickup of TFP growth in the United States (The U.S.

[40] There is very large uncertainty around any long-term forecast of the TFP growth rate. In particular, research has shown that current-decade growth of productivity holds little information as to the growth in the following decade. Perhaps naturally, the commentators are split on the prospects for innovation and productivity. See, for example, Brynjolfsson and McAfee (2014) and Gordon (2016) for two perspectives from the opposite ends of a spectrum.

	1970-2008	1970-2017	1970-2070
Estimated decline in AE R* (Sec 3)	**-2.8**	**-3.2**	
Public policies			
Government debt (life-cycle)	0.6	1.0	1.2
Government debt (incomplete markets)	0.3	0.5	0.7
Government spending	-0.1	-0.1	-0.2
Social Security	1.0	1.3	1.5
Old-age healthcare	0.9	1.2	1.3
Total impact of public policies	*2.7*	*3.9*	*4.5*
Implied decline in private sector R*	**-5.6**	**-6.8**	
Selected private sector forces			
TFP growth	-1.5	-1.7	-1.5
Population growth	-0.5	-0.8	-1.3
Longer retirement	-1.0	-1.1	-1.2
Length of working life	-0.1	0.0	0.0
Inequality	-0.5	-0.6	-0.8
Interactions	-0.8	-1.3	-1.6
Total private sector forces	*-4.4*	*-5.5*	*-6.4*

Note: all values are in percentage points.

Table 3.7: Decomposition of the decline in the neutral real interest rate in Advanced Economies

Congressional Budget Office (2019)).

In the precautionary saving model, we recalibrate the income process[41] and compare the steady states of the economy under the two calibrations.[42]

To reiterate, within each of the two models, we feed in the (model-specific) set of shocks all at the same time, thereby providing – within each model – an internally consistent laboratory to study this wide range of heterogenous trends. What we miss are the potential interactions across the two models. We assume that the comparable calibration across the two frameworks makes the results comparable and that simply adding the estimates of the impact on R* over the transition across the models results in a consistent picture. But ultimately, only the framework for analysis of all the forces that we consider – and perhaps further ones – in a single unifying setting would provide a definitive answer to these doubts. This avenue of inquiry is left for future work.

Table 3.7 summarizes the key results of the paper. First, in Section 3 we estimated that the neutral real rate has declined by over 3pp between 1970 and 2017. In Sections 4 and 5 we

[41] In particular, we change the income received in the highest income state. This is motivated by the fact that the increase in income inequality has been concentrated at the very top of the distribution, as documented by Piketty (2014) and others.

[42] We obtain the dynamic path by assuming that the effect builds steadily, including over the next decade. Our treatment of the dynamics is thus crude. We leave the analysis of the dynamic adjustment path for future work.

argued that public policies pushed rates up. Our models suggest that, together, the policies we considered have pushed rates up by nearly 4pp to date. This suggests that the private sector R* may have declined by around 7pp. The private sector forces we consider add up to a drag of 5.5pp, leaving over 1pp of the decline in private sector R* unaccounted for.

Our framework cannot account for the full extent of the decline in equilibrium rates, with over 1pp left unexplained in our preferred calibration. Our models miss some of the secular forces that likely pushed neutral rates lower over the past 40 years. One omission is the increasing concentration and the associated increase in market power of firms in the US and other advanced countries (Farhi and Gourio (2018)). Another force is driven by the finding that propensities to save are higher for those with high permanent income (Carroll (2000), Dynan, Skinner, and Zeldes (2004)). In light of these findings, our simulations likely understate the full impact of the increase in permanent income inequality. Using a model which captures this mechanism, Ludwig Straub (2017) estimates that the rise in inequality may have pushed down on the real equilibrium interest rate in the US by about 1pp through this channel. The decline in the price of capital goods may have contributed to lower investment propensities, further decreasing the neutral real rate (Sajedi and Thwaites (2016)). Finally, changes in the tax code – particularly the decline in overall tax progressivity in some jurisdictions – may have been a public sector force that depressed interest rates. We leave more detailed investigation of these forces for future work.

3.7 Conclusion

We draw three main conclusions from the analysis in this paper. First, the neutral real rate for the industrial world has trended downward for the last generation and this is best understood in terms of changes in private sector saving and investment propensities. In the face of neutral real rate estimates, past trends in indexed bond yields, and measures of real swap yields, this conclusion seems inescapable. It is also noteworthy that current real rates appear to be quite well predicted by pre-financial crisis trends. We believe that the these trends are best analyzed in terms of changes in saving and investment propensities or equivalently in terms of trends in desired wealth holdings by consumers and desired capital accumulation by producers. While factors involving liquidity, scarcity and risk no doubt bear on levels of real interest rates we find it highly implausible that they are the main factor accounting for trend movements. The movements are too large and too pervasive across assets and the fluctuations in spreads are too small and lacking in trend for these factors to account for the observed trends in the data.

Second, the neutral real rate would have declined substantially more over the last generation but for increases in government debt and expansions in social insurance programs. Both straightforward extrapolations of existing rules of thumb regarding debt and deficit impacts on interest rates and calculations using workhorse general equilibrium models suggest that fiscal policies have operated to raise real interest rates by several hundred basis points over the last generation. While this conclusion is dependent on our rejection of Ricardian equivalence, we see nothing that leads us to believe that increased government debt automatically calls for increased saving or that pay-as-you-go social security programs alter bequests for most families. The specific magnitudes are very uncertain, but open economy aspects and the possibility suggested by our analysis – that budget deficits emerge in response to excesses of private saving over private investment – lead us to think that we are more likely to understate than overstate the extent of fiscal support for real interest rates in recent years.

Third, the implication of our analysis that but for major increases in deficits debt and social insurance neutral real rates in the industrial world would be significantly negative by as much as several hundred basis points suggests substantial grounds for concern over secular stagnation. From the perspective of our analysis the private economy is prone to being caught in an underemployment equilibrium if real interest rates cannot fall far below zero. Full employment in recent years has been achieved where it has been achieved either through large budget deficits as in the United States or Japan or large trade surpluses as in Germany. It is worth considering that in the United States during the period prior to the financial crisis, negative real short term interest rates, a huge housing bubble, erosion of credit standards and expansionary fiscal policy were only sufficient to achieve moderate growth. Adequate growth in Europe was only maintained through what in retrospect appears to have been clearly unsustainable lending to the periphery.

What does our analysis say about stabilization policy? Most obviously it says that traditional levels of interest rates combined with balanced budgets or even stable debt-GDP ratios are a prescription for recession. Policymakers must, if they wish to avoid output being demand constrained, do some combination of accepting high and rising deficits and government debt levels, living with real interest rates very close to zero or negative, and finding structural policies that promote investment or reduce saving.

Olivier Blanchard (2019) makes the argument that traditional views about fiscal policy likely reflect excessive concern about debt when real interest rates are very low and likely to remain low for a long time to come. The sustainability of a given level of deficit or debt is greater when interest rates are low than when they are high. Nonetheless it has to be acknowledged that the US economy appears to be slowing to below potential growth despite projected primary deficits that will lead even on very favorable interest rate assumptions to steadily growing debt–to-GDP ratios that will ultimately set historical records. There is no guarantee that deficits sufficient to maintain positive neutral real rates will be associated with sustainable debt trajectories. Indeed, the Japanese experience suggests that this may not be the case.

Another possibility is the use of monetary policies that induce significantly negative real rates. This might be achieved through setting negative nominal rates, raising or adjusting inflation targets (e.g. through targeting average rate of inflation and thus "making up" for the past errors), or using unconventional monetary policies such as quantitative easing to achieve the equivalent of reductions in real rates. These approaches raise three issues. First, given that historically rates have been reduced by 500 basis points or more to mitigate recessions in industrial countries there is the question of whether enough room can be generated to stabilize the economy when the next downturn hits. Second, there are questions about whether starting at very low rates further rate reductions are actually stimulative. Eggertsson, Juelsrud, Summers, and Wold (2019) suggest that negative nominal rates actually may interfere with financial intermediation. Third, there is a range of concerns about the possible toxic effects of low rates, including suggestions that they make bubbles and over-leveraging more likely as they encourage risk taking, that they may lead to misallocation of capital by reducing loan payment levels and required rates of return, reinforce monopoly power, benefit the old at the expense of the young, and make the funding of insurance and pension obligations more difficult.

A final possibility is structural measures that reduce saving or promote investment. Clearly regulatory policies that encourage investment without sacrificing vital social objectives are desirable. The extent to which these are available is very much open to question. Investment incentives will also operate to raise demand. Policies that reduce the need for retirement saving, such as strengthening social security, or that improve social insurance, will increase aggregate demand even if operated on a balanced budget basis. So will policies that redistribute income from those with lower to those with higher propensities to consume.

It is tempting to suggest that any measure that increases productivity growth will operate to raise neutral real rates as consumers seek to spend more out of higher expected future incomes and firms increase their investment demand. Effects of this kind are indeed suggested by our formal model. We are not sure of their validity in practice. As Carroll and Summers (1991) point out, growth accelerations internationally have typically been associated with declining not increasing real rates and there is not much evidence that consumers are that forward looking, especially if the reforms are associated with transitional costs and heightened short-term uncertainties. Moreover in policy discussions central bankers usually cite stronger productivity as an antidote to inflation and therefore as a reason not to raise rates. Short-term productivity gains which reduce costs and inflation may act to elevate realized interest rates above the neutral rate, further worsening the demand imbalance.

All of this is to suggest that if secular stagnation is avoided in the years ahead it will not be be because it is somehow impossible in a free market economy, but instead because of policy choices. Our conclusions thus underscore the urgent priority for governments to find new sustainable ways of promoting investment to absorb the large supply of private savings and to devise novel long-term strategies to rekindle private demand.

Bibliography

Acemoglu, Daron (1998). "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality". In: *The Quarterly Journal of Economics* 113.4, pp. 1055–1089.

— (2002). "Directed Technical Change". In: *Review of Economic Studies* 69.4, pp. 781–809.

— (2007). "Equilibrium Bias of Technology". In: *Econometrica* 75.5, pp. 1371–1409.

Acemoglu, Daron and David Autor (2011). *Skills, tasks and technologies: Implications for employment and earnings*. Vol. 4. PART B, pp. 1043–1171.

Acemoglu, Daron and Veronica Guerrieri (2008). "Capital Deepening and Nonbalanced Economic Growth". In: *Journal of Political Economy* 116.3.

Acemoglu, Daron and Pascual Restrepo (2017). *Robots and Jobs: Evidence from US Labor Markets*. Tech. rep. Cambridge, MA: National Bureau of Economic Research, p. 43.

— (2018). *The race between man and machine: Implications of technology for growth, factor shares, and employment*.

— (2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor". Cambridge, MA.

Achdou, Yves et al. (2017). *Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach*. Tech. rep. Cambridge, MA: National Bureau of Economic Research, pp. 1–30.

Adler, Gustavo et al. (2017). "Gone with the Headwinds: Global Productivity".

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt (2014). "What Do We Learn From Schumpeterian Growth Theory?" In: *Handbook of Economic Growth* 2, pp. 515–563.

Aghion, Philippe, Antonin Bergeaud, et al. (2017). "Missing Growth from Creative Destruction". In: *Working Paper Series*.

Aghion, Philippe, Nick Bloom, et al. (2005). "Competition and innovation: An inverted-u relationship". In: *Quarterly Journal of Economics* 120.2, pp. 701–728.

Aghion, Philippe and Peter Howitt (1992a). "A Model of Growth Through Creative Destruction". In: *Econometrica* 60.2, pp. 323–351.

— (1992b). "A Model of Growth Through Creative Destruction". In: *Econometrica* 60.2, pp. 323–351.

— (2009). *Stages of Growth BT - The Economics of Growth*. 10. MIT Press, pp. 1–20.

Aghion, Philippe, Peter Howitt, and Giovanni L Violante (2002). "General purpose technology and wage inequality". In: *Journal of Economic Growth* 7.4, pp. 315–345.

Aghion, Philippe, Benjamin Jones, and Charles Jones (2017). “Artificial Intelligence and Economic Growth”. Cambridge, MA.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst (2017). *Leisure Luxuries and the Labor Supply of Young Men*. Tech. rep. Cambridge, MA: National Bureau of Economic Research.

Aguiar, Mark and Erik Hurst (2007a). “Measuring Trends in Leisure: The Allocation of Time Over Five Decades”. In: *The Quarterly Journal of Economics* 122.3, pp. 969–1006.

— (2007b). “Measuring trends in leisure: the allocation of time over five decades”. In: *The Quarterly Journal of Economics* 122.3, pp. 969–1006.

— (2016). “The Macroeconomics of Time Allocation”. In: *Handbook of Macroeconomics*. Vol. 2, pp. 1–44.

Aiyagari, Rao S (1993). “Uninsured Idiosyncratic Risk and Aggregate Saving”. In: *Source: The Quarterly Journal of Economics* 109184.3, pp. 659–684.

Aiyagari, S. R. (1994). “Uninsured Idiosyncratic Risk and Aggregate Saving”. In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Aiyagari, S.Rao and Ellen R McGrattan (1998). “The optimum quantity of debt”. In: *Journal of Monetary Economics* 42.3, pp. 447–469.

Alexopoulos, Michelle and Tricia Gladden (2006). “Wealth , Reservation Wages , and Labor Market Transitions in the U . S .: Evidence from the Survey of Income and Program”.

Anderson, Simon P and Régis Renault (2006). “Advertising Content”. In: *American Economic Review* 96 (1).

Armstrong, Mark (2006). “Competition in two-sided markets”. In: *RAND Journal of Economics* 00.3, pp. 668–691.

Arrieta Ibarra, Imanol et al. (2018). “Should We Treat Data as Labor?” In: *American Economic Review, Papers and Proceedings*.

Auclert, Adrien and Matthew Rognlie (2016). “Inequality and Aggregate Demand”.

Autor, David (2010). “The Polarization of Job Opportunities in the U.S. Labor Market”. In: *Community Investments* 23.April, pp. 360–361.

— (2015). “Why Are There Still So Many Jobs? The History and Future of Workplace Automation”. In: *Journal of Economic Perspectives*.

Autor, David H, Frank Levy, and Richard J Murnane (2003). “The Skill Content of Recent Technological Change: An Empirical Exploration”. In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333.

Bagwell, Kyle (2007). “The economic analysis of advertising”. In: *Handbook of Industrial Organisation*.

Bardóczy, Bence (2017). “Labor-Market Matching with Precautionary Savings”.

Barro, Robert J. (1974). “Are Government Bonds Net Wealth?” In: *Journal of Political Economy* 82.6, pp. 1095–1117.

Barro, Robert J and Xavier Sala-i-Martin (1990). “World Real Interest Rates”. In: *NBER Macroeconmics Annual* 5.January, pp. 15–74.

Baxter, Marianne and Robert G King (1993). *Fiscal Policy in General Equilibirium*.

Bean, Charles (2016). *Independent Review of UK Economic Statistics*. Tech. rep.

Becker, Gary S (1965). "A Theory of the Allocation of Time". In: *The Economic Journal* 75.299, pp. 493–517.

Blanchard, Olivier (2019). "Public Debt and Low Interest Rates". In: *American Economic Review, Papers and Proceedings*.

Blanchard, Olivier J (1985). "Debt, Deficits, and Finite Horizons". In: *The Journal of Political Economy* 93.2, pp. 223–247.

Blanchard, Olivier Jean and Lawrence H Summers (1984). "Perspectives on High World Real Interest Rates". In: *Brookings Papers on Economic Activity* 1984.2, pp. 273–324.

Bloemen, Hans and Elena Stancanelli (2002). "Individual Wealth, Reservation Wages, and Transitions into Employment". In: *Journal of Labor Economics* 19.2, pp. 400–439.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb (2017). "Are Ideas Getting Harder to Find?" In:

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb (2017). "Are Ideas Getting Harder to Find?" In:

Boik, Andre, Shane Greenstein, and Jeffrey Prince (2016). "The Empirical Economics of Online Attention". In:

Boppart, Timo and Per Krusell (2016). "Labor Supply in the Past, Present, and Future: A Balance-Growth Perspective".

Boppart, Timo and L Rachel Ngai (2016). "Rising inequality and trends in leisure *".

Brand, Claus, Marcin Bielecki, and Adrian Penalver (2018). "The natural rate of interest: estimates, drivers, and challenges to monetary policy".

Bridgman, Benjamin (2016). "Is Productivity on Vacation? The Impact of the Digital Economy on the Value of Leisure". In:

Broda, C. and D. E. Weinstein (2006). "Globalization and the Gains From Variety". In: *The Quarterly Journal of Economics* 121.2, pp. 541–585.

Brook, Anne-marie (2003). "Recent and Prospective Trends in Real Long-Term Interest Rates".

Brülhart, Marius, Jonathan Gruber, Matthias Krapf, and Kurt Schmidheiny (2016). "Taxing Wealth: Evidence from Switzerland". In: *NBER Working Paper Series*, p. 37.

Brynjolfsson, Erik, Felix Eggers, and Avinsash Gannamaneni (2017). "Using Massive Online Choice Experiments to Measure Changes in Well-being".

Brynjolfsson, Erik and Andrew McAfee (2014). *The second machine age : work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company, 2014, p. 306.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" In: *AEA Papers and Proceedings* 108, pp. 43–47.

Brynjolfsson, Erik and Joo Hee Oh (2012). "The attention economy: Measuring the value of free digital services on the internet". In: *International Conference on Information Systems, ICIS 2012* 4, pp. 3243–3261.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2017). “Artificial Intelligence and the Modern Productivity Paradox : A Clash of Expectations and Statistics”.

Byrne, David M, John G Fernald, and Marshall B Reinsdorf (2016). “Does the United States have a productivity slowdown or a measurement problem?” In: *Brookings Papers on Economic Activity*.

Caballero, Ricardo J and Emmanuel Farhi (2014). “The safety trap”. In: *NBER Working Paper*.

Caballero, Ricardo J. and Emmanuel Farhi (2018). “The safety trap”. In: *Review of Economic Studies* 85.1, pp. 223–274.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre Olivier Gourinchas (2016). “Safe asset scarcity and aggregate demand”. In: *American Economic Review*. Vol. 106. 5, pp. 513–518.

Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas (2017). “Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares”. In: *Mimeo*.

Carroll, Christopher D (1997). “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis”. In: *The Quarterly Journal of Economics* 112.1, pp. 1–55.

— (2000). “Why Do the Rich Save So Much?” In: *Does Atlas Shrug? The Economic Consequences of Taxing the Rich.* Pp. 465–84.

Carroll, Christopher D, Jody Overland, and David N Weil (2000). “Saving and Growth with Habit Formation”. In: *American Economic Review* 90.3, pp. 341–355.

Carroll, Christopher D and Lawrence H Summers (1991). “Consumption Growth Parallels Income Growth: Some New Evidence”. In: *National Saving and Economic Performance*, pp. 305–348.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio (2015). “Demographics and Real Interest Rates :” in: pp. 1–29.

Caselli, Francesco (1999). “Technological Revolutions”. In: *American Economic Review* 89.1, pp. 78–102.

Caselli, Francesco and Alan Manning (2018). “Robot Arithmetic: New Technology And Wages”. In: *American Economic Review: Insights*, pp. 1–15.

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull (2003). “Accounting for the U.S. Earnings and Wealth Inequality”. In: *Journal of Political Economy* 111.4, pp. 818–857.

Chamberlin, Edward (1933). *The theory of monopolistic competition*. Harvard University Press, p. 396.

Christensen, Jens H E and Glenn D Rudebusch (2017). “A New Normal for Interest Rates? Evidence from Inflation-Indexed Debt”. In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–39.

Christensen, Matthew A, Laura Bettencourt, et al. (2016). “Direct Measurements of Smartphone Screen- Time: Relationships with Demographics and Sleep”. In:

Constâncio, Vítor (2016). *The challenge of low real interest rates for monetary policy*. Tech. rep.

Cui, Wei and Vincent Sterk (2018). “Quantitative Easing”.

De Loecker, Jan and Jan Eeckhout (2017). “The Rise of Market Power and the Macroeconomic Implications”. In:

Deaton, Angus (1991). “Saving and Liquidity Constraints”. In: *Econometrica* 59.5, pp. 1221–1248.

Del Negro, Marco, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti (2017). “Safety, Liquidity, and the Natural Rate of Interest”. In: *Brookings Papers on Economic Activity* May 2017.812, pp. 1–81.

Del Negro, Marco, Marc P Giannoni, Domenico Giannone, and Andrea Tambalotti (2018). “Global Trends in Interest Rates”. In: *Federal Reserve Bank of New York Staff Reports* 2018.866.

DellaVigna, Stefano and Matthew Gentzkow (2010). “Persuasion: Empirical Evidence”. In: *Annual Review of Economics* 2.1, pp. 643–669.

Diamond, Peter A. (1965). “National Debt in a Neoclassical Growth Model”. In: *American Economic Review* 55.4, pp. 1126–1150.

Dixit, Avinash K and Joseph E Stiglitz (1977). “Monopolistic Competition and Optimum Product Diversity”. In: *American Economic Review* 67.3, pp. 297–308.

Draghi, Mario (2016). *The International Dimension of Monetary Policy*. Tech. rep.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes (2004). “Do the Rich Save More?” In: *Journal of Political Economy* 112.2, pp. 397–444.

Eggertsson, Gauti B. and Paul Krugman (2012). “Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach”. In: *Quarterly Journal of Economics* 127.3, pp. 1469–1513.

Eggertsson, Gauti B and Neil R Mehrotra (2014). “A Model of Secular Stagnation”.

Eggertsson, Gauti B, Neil R Mehrotra, and Jacob A Robbins (2019). “A Model of Secular Stagnation: Theory and Quantitative Evaluation”. In: *American Economic Journal: Macroeconomics* 11.1, pp. 1–48.

Eggertsson, Gauti B, Neil R Mehrotra, and Lawrence H Summers (2016). “Secular Stagnation in the Open Economy”.

Eggertsson, Gauti, Ragnar Juelsrud, Lawrence Summers, and Ella Getz Wold (2019). “Negative Nominal Interest Rates and the Bank Lending Channel”. Cambridge, MA.

Eichengreen, Barry (2015). “Secular Stagnation: The Long View”. In: *American Economic Review* 105.5, pp. 66–70.

Elmendorf, Douglas W and N. Gregory Mankiw (1999). “Government debt”. In: *Handbook of Macroeconomics*, pp. 1615–1669.

Engen, Eric M. and Jonathan Gruber (2001). “Unemployment insurance and precautionary saving”. In: *Journal of Monetary Economics* 47.3, pp. 545–579.

Engen, Eric and R Glenn Hubbard (2004). “Federal Government Debts and Interest Rates”.

Epstein, Larry G and Stanley E Zin (1991). “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis”. In: *Journal of Political Economy* 99.2, pp. 263–286.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2016). “Heterogeneity and Persistence in Returns to Wealth”.

Fagereng, Andreas, Martin Holm, Benjamin Moll, and Gisle Natvik (2018). “Saving Behavior across the Wealth Distribution”.

Faini, Riccardo (2006). “Fiscal policy and interest rates in Europe”. In: *Economic Policy* 21.47, pp. 444–489.

Farhi, Emmanuel and François Gourio (2018). “Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia”.

Fischer, Stanley (2016). *Why Are Interest Rates So Low? Causes and Implications*. Tech. rep.

Frey, Carl B and Michael A Osborne (2013). “the Future of Employment: How Susceptible Are Jobs To Computerisation? *”. In: pp. 1–72.

Fujiwara, Shigeaki et al. (2016). “Developments in the Natural Rate of Interest in Japan”. In: *Bank of Japan Review* 2016-E-12.

Furceri, Davide and Andrea Pescatori (2014). “Perspectives on global real interest rates”. In: *WEO: Recovery Strengthens, Remains Uneven* April, pp. 1–32.

Gagnon, Etienne, Benjamin Johannsen, and David Lopez-Salido (2016). “Understanding the New Normal: The Role of Demographics”. In: *Finance and Economics Discussion Series* 2016.080.

Gale, William G and Peter R Orszag (2002). “The Economic Effects of Long-Term Fiscal Discipline”. In: *Urban-Brookings Tax Policy Center Discussion Paper*.

Gali, Jordi (2008). *Monetary policy, inflation, and the business cycle : an introduction to the new Keynesian framework and its applications*, p. 279.

Gertler, Mark (1999). “Government debt and social security in a life-cycle economy”. In: *Carnegie-Rochester Conference Series on Public Policy* 50, pp. 61–110.

Giupponi, Giulia (2018). “When Income Effects are Large : Labor Supply Responses and the Value of Welfare Transfers”. In:

Goolsbee, Austan and Peter J Klenow (2006). “Valuing consumer products by the time spent using them: An application to the Internet”. In: *American Economic Review*. Vol. 96. 2, pp. 108–113.

Gordon, Robert J. (2016). “Perspectives on the rise and fall of American growth”. In: *American Economic Review*. Vol. 106. 5, pp. 72–76.

Gourinchas, Pierre-Olivier and Hélène Rey (2016). “Real Interest Rates, Imbalances and the Curse of Regional Safe Asset Providers at the Zero Lower Bound”.

Grossman, Gene M. and Elhanan Helpman (1991). “Quality Ladders in the Theory of Growth”. In: *The Review of Economic Studies* 58.1, p. 43.

Grossman, Gene M. and Esteban Rossi-Hansberg (2008). “Trading tasks: A simple theory of offshoring”. In: *American Economic Review* 98.5, pp. 1978–1997.

Guvenen, F., F. Karahan, Serdar Ozkan, and Jae Song (2015). “What do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?” In: *NBER Working Paper Series* January.wp29013.

Guvenen, Fatih (2009). “An empirical investigation of labor income processes”. In: *Review of Economic Dynamics* 12.1, pp. 58–79.

Hall, Robert E (2014). “What the Cyclical Response of Advertising Reveals about Markups and other Macroeconomic Wedges”. In: *Working Paper* March 2013, pp. 1–34.

— (2017). “The Role of the Growth of Risk-Averse Wealth in the Decline of the Safe Real Interest Rate”.

Hamilton, James D (1994). *Time series analysis*. Princeton, N.J.: Princeton University Press.

Hamilton, James D., Ethan S. Harris, Jan Hatzius, and Kenneth D. West (2016). “The Equilibrium Real Funds Rate: Past, Present, and Future”. In: *IMF Economic Review* 64.4, pp. 660–707.

Havranek, Tomas, Roman Horvath, Zuzana Irsova, and Marek Rusnak (2015). “Cross-country heterogeneity in intertemporal substitution”. In: *Journal of International Economics* 96.1, pp. 100–118.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante (2009). “Quantitative Macroeconomics with Heterogeneous Households”. In: *Annual Review of Economics* 1.1, pp. 319–354.

Hémous, David and Morten Olsen (2018). “The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality”.

Ho, Paul (2018). “Estimating the Effects of Demographics on Interest Rates: A Robust Bayesian Perspective”. In:

Holston, Kathryn, Thomas Laubach, and John C Williams (2017a). “Documentation of R Code and Data for Measuring the Natural Rate of Interest: International Trends are Determinants”.

Holston, Kathryn, Thomas Laubach, and John C. Williams (2017b). “Measuring the natural rate of interest: International trends and determinants”. In: *Journal of International Economics* 108, S59–S75.

Huberman, Michael and Chris Minns (2007). “The times they are not changin’: Days and hours of work in Old and New Worlds, 1870-2000”. In: *Explorations in Economic History* 44.4, pp. 538–567.

Huggett, Mark (1993). “The risk-free rate in heterogeneous-agent incomplete-insurance economies”. In: *Journal of Economic Dynamics and Control* 17.5-6, pp. 953–969.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman (2018). “Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark”. Cambridge, MA.

Jones, Charles I (1995). “R&D-Based Models of Economic Growth”. In: *Journal of Political Economy* 103.4, pp. 759–784.

Jordà, Òscar, Katharina Knoll, et al. (2017). “The Rate of Return on Everything”. In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 1–123.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2016). “Macrofinancial History and the New Business Cycle Facts”. In: *NBER Macroeconomics Annual 2016* 31.

Justiniano, Alejandro and Giorgio E Primiceri (2010). "Measuring the equilibrium real interest rate". In: *Economic Perspectives* Q1, pp. 14–27.

Kalai, Ehud (1977). "Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons". In: *Econometrica* 45.7, p. 1623.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante (2018). "Monetary Policy According to HANK". In: *American Economic Review* 108.3, pp. 697–743.

Kaplan, Greg, Giovanni L Violante, and Justin Weidner (2014). "The wealthy Hand-to-Mouth". In: *Brookings Papers on Economic Activity* Spring.

Katz, Lawrence F and Kevin M Murphy (1992). "Changes in Relative Wages, 1963-1987: Supply and Demand Factors". In: *The Quarterly Journal of Economics* 107.1, pp. 35–78.

Kay, John (2015). *Other People's Money: Masters of the Universe or Servants of the People?* Profile Books.

King, Mervyn and David Low (2014). "Measuring the "World" Real Interest Rate".

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo (1988). "Production, growth and business cycles". In: *Journal of Monetary Economics* 21.2-3, pp. 195–232.

— (2002). "Production, Growth and Business Cycles: Technical Appendix". In: *Computational Economics* 20.1/2, pp. 87–116.

Kinoshita, Noriaki (2006). "Government Debt and Long-Term Interest Rates". In: *IMF Working Papers* 06.63, p. 1.

Kocherlakota, Narayana (2015). "Public Debt and the Long-Run Neutral Real Interest Rate".

Kortum, Samuel (1997). "Research, Patenting, and Technological Change". In: *Econometrica* 65.6, pp. 1389–1419.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen (2012). "The Aggregate Demand for Treasury Debt". In: *Journal of Political Economy* 120.2, pp. 233–267.

Krueger, Alan B. and Andreas I. Mueller (2016). "A Contribution to the Empirics of Reservation Wages". In: *American Economic Journal: Economic Policy* 8.1, pp. 142–179.

Krusell, Per, Toshihiko Mukoyama, and Aysegul Sahin (2010). "Labour-market matching with precautionary savings and aggregate fluctuations". In: *Review of Economic Studies* 77.4, pp. 1477–1507.

Krusell, Per, Lee E Ohanian, Jose-Victor Rios-Rull, and Giovanni L Violante (2000). "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis". In: *Econometrica* 68.5, pp. 1029–1053.

Kuroda, Haruhiko (2016). *The battle against deflation - the evolution of monetary policy and Japan's experience*. Tech. rep. Speech at Columbia University.

Laforte, Jean-Philippe and John M. Roberts (2014). "November 2014 Update of the FRB/US Model". In: *FEDS Notes* 2014.0036.

Laibson, D. (1997). "Golden Eggs and Hyperbolic Discounting". In: *The Quarterly Journal of Economics* 112.2, pp. 443–478.

Lammers, Marloes (2014). "The effects of savings on reservation wages and search effort". In: *Labour Economics* 27, pp. 83–98.

Laubach, Thomas (2009). “New Evidence on the Interest Rate Effects of Budget Deficits and Debt”. In: *Source Journal of the European Economic Association* 7.4, pp. 858–885.

Laubach, Thomas and John C. Williams (2003). “Measuring the Natural Rate of Interest”. In: *Review of Economics and Statistics* 85.4, pp. 1063–1070.

Lentz, Rasmus (2009). “Optimal unemployment insurance in an estimated job search model with savings”. In: *Review of Economic Dynamics* 12, pp. 37–57.

Lewis, Randall A and Justin M Rao (2015). “The unfavorable economics of measuring the returns to advertising”. In: *Quarterly Journal of Economics* 130.4, pp. 1941–1973.

Lewis, Randall A and David H Reiley (2014). “Online ads and offline sales: Measuring the effect of retail advertising via a controlled experiment on Yahoo!” In: *Quantitative Marketing and Economics* 12.3, pp. 235–266.

L’Haridon, Olivier, Franck Malherbet, and Sébastien Pérez-Duarte (2013). “Does bargaining matter in the small firms matching model?” In: *Labour Economics* 21, pp. 42–58.

Lillard, Lee A and Yoram Weiss (1979). “Components of Variation in Panel Earnings Data: American Scientists 1960-70”. In: *Econometrica* 47.2, p. 437.

Lisack, Noemie, Rana Sajedi, and Gregory Thwaites (2017). “Demographic trends and the real interest rate”.

Lucas, Robert E and Benjamin Moll (2014). “Knowledge Growth and the Allocation of Time”. In: *Journal of Political Economy* 122.1.

MaCurdy, Thomas E. (1982). “The use of time series processes to model the error structure of earnings in a longitudinal data analysis”. In: *Journal of Econometrics* 18.1, pp. 83–114.

Mankiw, N Gregory (2000). “The Savers–Spenders Theory of Fiscal Policy”. In: *American Economic Review* 90.2, pp. 120–125.

Marshall, Alfred (1890). *Principles of Economics*. Eighth Edi. London: MacMillan and Co. Ltd.

McKinsey Global Institute (2017). *A future that works*. Tech. rep.

Moll, Benjamin, Lukasz Rachel, and Pascual Restrepo (2019). “Automation’s Impact on Income and Wealth Inequality”.

Nakamura, Leonard, Jon Samuels, and Rachel Soloveichik (2017). “Measuring the ”Free” Digital Economy within the GDP and Productivity Accounts”.

Ngai, L. Rachel and Christopher A. Pissarides (2008). “Trends in hours and economic growth”. In: *Review of Economic Dynamics* 11.2, pp. 239–256.

Nixon, Dan (2017). *Is the economy suffering from the crisis of attention?*

OECD (2018). *OECD Tax Policy Studies: The Role and Design of Net Wealth Taxes in the OECD*. Tech. rep.

Oh, Hyunseung and Ricardo Reis (2012). “Targeted transfers and the fiscal response to the great recession”. In: *Journal of Monetary Economics* 59.

Okazaki, Yosuke and Nao Sudo (2018). “Natural Rate of Interest in Japan-Measuring its size and identifying drivers based on a DSGE model”. In: *Bank of Japan Working Paper Series*.

Piketty, Thomas (2014). *Capital in the Twenty-First Century*.

Piketty, Thomas and Emmanuel Saez (2007). "How Progressive is the U.S. Federal Tax System? A Historical and International Perspective". In: *Journal of Economic Perspectives* 21.1, pp. 3–24.

Pissarides, Christopher A (1990). *Equilibrium Unemployment*. First Edit. Oxford: Blackwell Publishing Ltd.

Powell, Jerome H (2018). *The Federal Reserve's Framework for Monitoring Financial Stability*. Tech. rep.

Quadrini, Vincenzo and José-Víctor Ríos-Rull (2014). "Inequality in Macroeconomics". In: *Handbook of Income Distribution*.

Rachel, Lukasz and Thomas D Smith (2015). "Secular drivers of the global real interest rate".

— (2017). "Are low real interest rates here to stay?" In: *International Journal of Central Banking* 13.3, pp. 1–42.

Ramey, Valerie A. and Neville Francis (2009). "A century of work and leisure". In: *American Economic Journal: Macroeconomics* 1.2, pp. 189–224.

Rochet, Jean-Charles and Jean Tirole (2003). "Platform Competition in two-sided markets". In: *Journal of the european economic association 1 (4), 990-1029*.

Romer, Paul M. (1990). "Endogenous Technological Change". In: *Journal of Political Economy* 98.5, S71–S102.

Sajedi, Rana and Gregory Thwaites (2016). "Why Are Real Interest Rates So Low? The Role of the Relative Price of Investment Goods". In: *IMF Economic Review* 64.4, pp. 635–659.

Segerstrom, Paul S (1998). "Endogenous Growth without Scale Effects". In: *American Economic Review* 88.5, pp. 1290–1310.

Seim, David (2017). "Behavioral responses to wealth taxes: Evidence from Sweden". In: *American Economic Journal: Economic Policy* 9.4, pp. 395–421.

Shimer, Robert (2005). "The Cyclical Behavior of Equilibirum Unemployment and Vacancies". In: *The American Economic Review* 8.2004, pp. 169–235.

Smets, Frank and Rafael Wouters (2007). "Shocks and Firctions In US Business Cycle: A Bayesian DSGE Approach". In: *American Economic Review* 97, pp. 586–606.

Spence, Michael and Bruce Owen (1977). "Television Programming, Monopolistic Competition and Welfare". In: *Quarterly Journal of Economics* 91.1, pp. 103–126.

Stancanelli, Elena (1999). "Do the rich stay unemployed longer? An empirical study for the UK". In: *Oxford Bulletin of Economics and Statistics* 61.3, pp. 295–314.

Straub, Ludwig (2017). "Consumption, saving, and the distribution of permanent income". In:

Summers, Lawrence H (2015a). "Demand side secular stagnation". In: *American Economic Review* 105.5, pp. 60–65.

— (2015b). "Have we entered an age of secular stagnation". In: *IMF Economic Review* 63, pp. 277–280.

Syverson, Chad (2016). "Challenges to mismeasurement explanations for the U.S. productivity slowdown".

Terranova, Tatiana (2012). "Attention, economy and the brain".

Teulings, Coen and Richard Baldwin (2014). *Secular Stagnation: Facts, Causes and Cures — VOX, CEPR's Policy Portal*.

The U.S. Congressional Budget Office (2019). *The Budget and Economic Outlook*. Tech. rep.

Tirole, Jean (2017). *Economics for the Common Good*.

Varian, Hal (2017). "Artificial Intelligence, Economics, and Industrial Organization".

Verduyn, Philippe et al. (2017). "Do Social Network Sites Enhance or Undermine Subjective Well-Being? A Critical Review". In: *Social Issues and Policy Review* 11.1, pp. 274–302.

Wallsten, Scott (2013). "What are we not doing when were online".

Walsh, Carl E. (1998). *Monetary theory and policy*. MIT Press, p. 528.

Ward, Adrian F., Kristen Duke, Ayelet Gneezy, and Maarten W. Bos (2017). "Brain Drain: The Mere Presence of One's Own Smartphone Reduces Available Cognitive Capacity". In: *Journal of the Association for Consumer Research* 2.2, pp. 140–154.

Wicksell, Knut (1989). *Interest and Prices*, p. 252.

Winter, Christoph (2016). "The Impact of Government Debt on the Long-Run Natural Real Interest Rate - a Quantitative Evaluation". In:

Woodford, Michael (2003). *Interest and prices : foundations of a theory of monetary policy*. Princeton University Press, p. 785.

Wynne, Mark A. and Ren Zhang (2018). "Estimating the natural rate of interest in an open economy". In: *Empirical Economics* 55.3, pp. 1291–1318.

Yaari, Menahem E. (1965). "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer". In: *The Review of Economic Studies* 32.2, p. 137.

Zeira, Joseph (1998). "Workers, Machines, and Economic Growth". In: *The Quarterly Journal of Economics* 113.4, pp. 1091–1117.

Zoutman, Floris T (2015). "The Effect of Capital Taxation on Households' Portfolio Composition and Intertemporal Choice".

Appendices

Appendices to Chapter 1

Appendix 1.A: Trends in leisure-enhancing technologies

Material in this Appendix further motivates the focus of this paper and is a useful background to the analysis. I discuss the supply of leisure-enhancing innovations as well as the evidence on changing time allocation patterns.

Evidence of leisure-enhancing innovations

Figure 1.1 made the point that the availability of zero-price services aimed at extracting peoples' time and attention is not a new phenomenon; instead, leisure technologies have been around for at least a century. Nonetheless, the amount of free content grew significantly faster than GDP over the recent past, chiefly due to the rise of digital content. This recent trend can be observed across a wide range of indicators. For example, Figure 3.12 shows mobile phone apps available on the two main mobile platforms. There are now more than 3 million apps on either of the platforms, and the majority of those are available free of charge to the consumer (Figure 3.13).

Consistent with this rapid growth of new leisure services, the leisure sector appears to be an increasingly important driver of the overall R&D spending. No explicit measure for the R&D share of the leisure economy exists; but it is possible to construct proxies by looking at the industries most likely engaged in the leisure-enhancing innovations. Figure 3.14 shows that the share of R&D spending accounted for by the sectors in the leisure economy such as video and

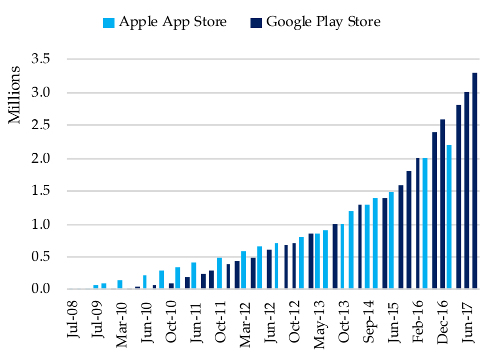

Figure 3.12: Number of apps worldwide
Source: Statista, Android, Google, Apple.

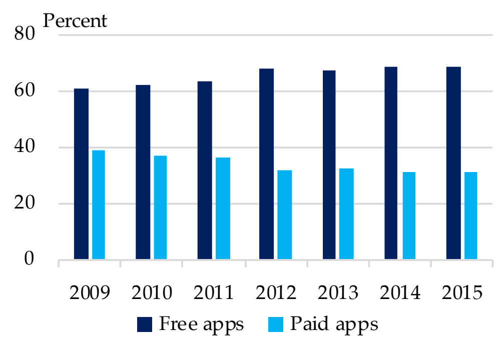

Figure 3.13: Free and paid apps offered in the Google Play Store
Source: Statista.

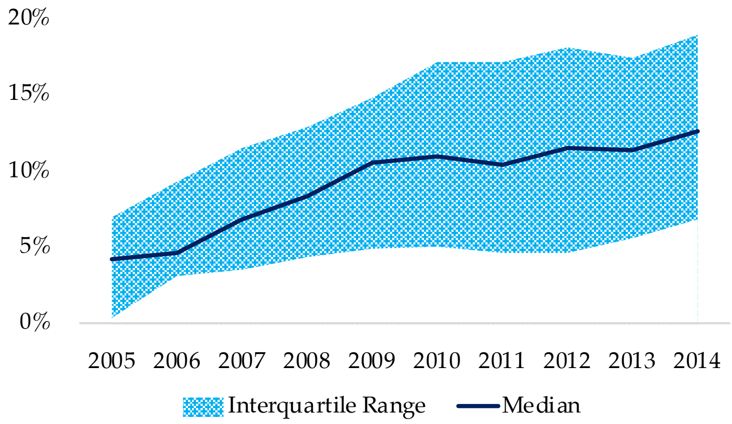

Figure 3.14: R&D expenditure share of the (proxy for the) leisure economy
Source: OECD. Includes data for an unbalanced panel of 39 countries. The figure shows the median and the interquartile range of the country-level share of R&D spending in the following sectors: publishing; motion picture, video and television program production; sound recording; programming and broadcasting activities; telecommunications services; computer programming, consultancy and related activities; information service activities; data processing, hosting and related activities; and web portals.

TV program production, sound recording, broadcasting and web portals has been increasing at rapid pace.

Recent changes in time allocation patterns

There is some indicative evidence that the increased leisure time is, in part, substituting for time spent at work. A recent study by Christensen, Bettencourt, et al. (2016) measured smartphone screen time over the course of an average day among a sample of 653 people (Figure 3.15). Time spent on the phone averaged 1 hour and 29 minutes per day, corroborating the order of magnitude in the Nielsen 2016 data. Crucially, the mobile phone usage appears to be uniformly distributed throughout the day, suggesting that leisure time is, in part, substituting for time spent working. In a different study, Wallsten (2013) uses time use surveys to estimate that each minute spent on the internet is associated with loss of work-time of about 20 seconds.

Indeed, one feature of the latest technology is that it allowed leisure to "compete" with work much more directly than has been the case in the past. While it may not have been possible to watch TV at work, online entertainment is available during the work hours. The theory in this paper focuses on the work-leisure margin, which may thus be particularly relevant for the digital technologies of the past decade and going forward. I come back to this issue when interpreting the results in Section 1.3.

In summary, the data presented above paint a fairly consistent picture of the increasing importance of the leisure sector of the economy. This increasing relevance motivates the rest of the paper, in which I develop a simple framework consistent with the observations above and use it to study the broader implications of these phenomena.

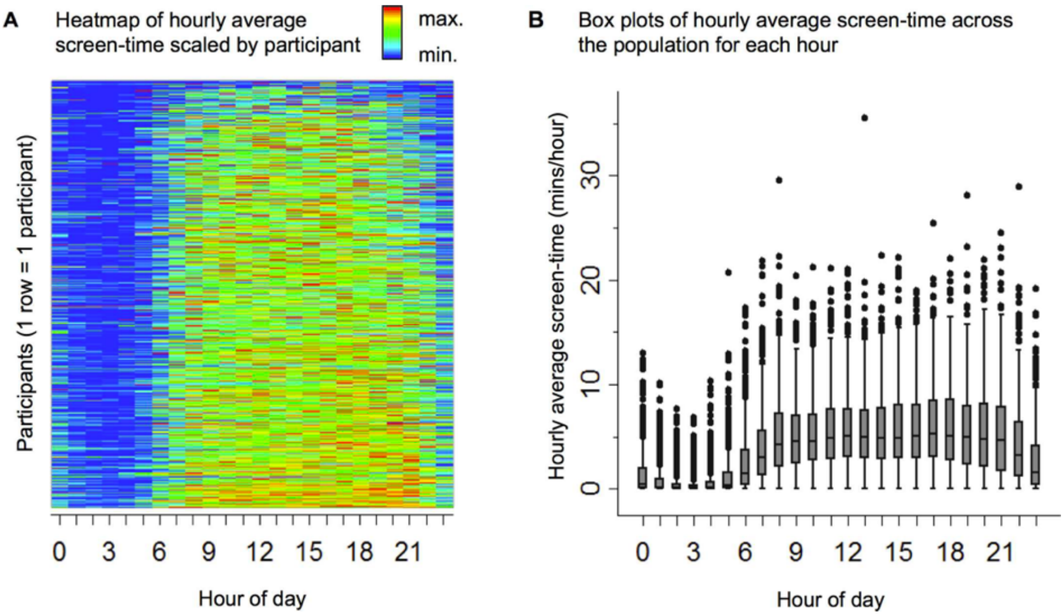

Figure 3.15: Mobile phone use over the course of an average day

Source: Christensen, Bettencourt, et al. (2016).

Appendix 1.B: Proofs and derivations

Proof of Lemma 1

Proof. To see why this Lemma holds, note that the assumption that ν lies between 0 and 1 implies that each activity exhibits a diminishing marginal productivity in leisure production, and so households like to diversify their use of leisure time. By definition, the total amount of marketable leisure time is given by $H_L \equiv \int_0^M h(j)dj$. Recall that for each activity j, household can access any amount of services $m(j)$ at no cost: these are free of charge, non-rival and non-excludable. Naturally then, availability of free services will never be constraining households leisure production, so that $m(j) \geqslant h(j) \forall j$. Consequently, we can simply write:

$$l = \left(\int_0^M h(j)^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}} dj. \tag{3.9}$$

Households choose how much time to devote to each individual activity to maximize production of leisure for any given amount of total leisure time. Because of the symmetry of this problem, it is immediate that the optimal choice is to spread free time evenly across the M available leisure options: $h(j) = h \; \forall j$; and therefore

$$H_L \equiv \int_0^M h(j)dj = Mh \Rightarrow h = \frac{H_L}{M}.$$

Plugging this back into the simplified production function for leisure (3.9) yields the result. ■

Verification of goods market clearing in the static model

Consider first the equilibrium output of each firm in the monopolistically competitive intermediate sector. The equilibrium output is:

$$x(i) = x = [\alpha^2 \left(\frac{b(i)^*}{\bar{B}^*}\right)^{\alpha\chi} L_Y^{1-\alpha}]^{\frac{1}{1-\alpha}} = \alpha^{\frac{2}{1-\alpha}} L_Y = \alpha^{\frac{2}{1-\alpha}} (\kappa A)^{\frac{\nu-1}{2-\nu}}$$

which is exactly the same as in the model without advertising. And as I showed above, the price that producers charge are unchanged from the standard model too. Thus the revenue they generate does not depend at all on the value of χ(which measures the prevailing importance of advertising). The only channel through which the introduction of advertising to the model affects the intermediate firms' profits is through the additional cost it imposes.

Gross production of the final good is thus:

$$Y = \int_0^A \left(\left(\frac{b(i)}{\bar{B}}\right)^{\chi} x(i)\right)^{\alpha} L_Y^{1-\alpha} di = Ax^{\alpha} L_Y^{1-\alpha} = \alpha^{\frac{2\alpha}{1-\alpha}} AL_Y = \alpha^{\frac{2\alpha}{1-\alpha}} (\kappa A)^{\frac{1}{2-\nu}} \tag{3.10}$$

How is this gross output split in equilibrium? Recall that output is used as consumption (of the capitalists and the workers), and as intermediate input into the consumer goods sector and in the marketing sector. Recall also that consumption of the capitalists is simply equal to the profits of all the firms in the economy:

$$C = (\pi + \pi_B) AL_Y$$

Workers consume their labour income, hence

$$c = H_W \cdot w = L_Y \cdot (1-\alpha)\frac{Y}{L_Y} = (1-\alpha)Y = (1-\alpha)\alpha^{\frac{2\alpha}{1-\alpha}} AL_Y$$

Intermediate inputs in the two sectors are:

$$X = \alpha^{\frac{2}{1-\alpha}} AL_Y$$

$$X_B = \kappa AL_Y$$

To verify that good markets clear, note that

$$\begin{aligned} X + X_B + C &= AL_Y \alpha^{\frac{2}{1-\alpha}} \left(1 + \frac{\alpha}{1-\alpha}\chi^2(1-\nu) + [\frac{1}{\alpha} - 1 - \chi] + \chi - \frac{\alpha}{1-\alpha}\chi^2(1-\nu)\right) \\ &= AL_Y \alpha^{\frac{1+\alpha}{1-\alpha}} = \alpha Y \end{aligned} \tag{3.11}$$

which together with the fact that $c = (1-\alpha)Y$ proves market clearing.

Details on the solution of the firms' and platform's problem

The optimal choice of quantities and prices by the monopolistically competitive firms results in equilibrium profits given by:

$$\Pi(i) = \left(\frac{b(i)}{\bar{B}}\right)^{\frac{\alpha}{1-\alpha}\chi} \left[\alpha^{\frac{2}{1-\alpha}} \left(\frac{1}{\alpha} - 1 - \chi\right)\right] L_Y. \tag{3.12}$$

Profits are positive if $\alpha < \frac{1}{1+\chi}$, a condition that restricts the importance of advertising. Intuitively, if advertising is perceived to be very effective (χ is high), firms would purchase so much of it that it would make them unprofitable. Firms would then exit the market. This is clearly not consistent with the notion of equilibrium or indeed with the observation of the real world: while marketing expenditure does raise overall costs of running a company, firms stop short of the point where ad spending eats up their entire profits. I therefore assume that the condition $\alpha < \frac{1}{1+\chi}$ is satisfied throughout the paper; the illustrative parametrization below shows that this assumption is not restrictive.

The assumption in the main text is that the marketing platform takes the average amount of advertising $\bar{B}$ as given.[43] There are two reasons for this assumption: first, it simplifies the algebra slightly, allowing for closed form solution without altering any of the substance that follows; and second, it "limits" the market power of the platform in the model. While in the real world these companies are large and they have significant market power, they are not true monopolists, so this assumption aligns the model more closely with the real world. An alternative assumption would be to let the platform internalize the fact that more advertising for a given firm imposes an externality on all other firms. Mathematically, the optimal aggregate advertising output would be given by $\tilde{B} = (\kappa(A-1)L_Y)^{1-\nu} < B$. The difference between the two solutions will be small for large A. Intuitively, if there are many varieties, the choice of advertising by any single firm has only a very small effect on the average level of advertising.

Households' intertemporal problem in the dynamic model

The representative household maximizes the stream of per-member utility discounted with the effective discount rate $\rho - n$:[44]

$$\max_{c, H_L} \int_0^\infty e^{-(\rho-n)t} \left(log(c) + l \cdot \bar{H}_L^{1-\nu}\right) dt \tag{3.13}$$

[43]Future work could consider platform competition, following the footsteps of Aghion, Bloom, et al. (2005) and contributing to the subequent literature on competition and innovation.

[44]The law of motion for total assets is $\dot{D} = rD + wL_Y + wL_A - C$, where D are assets in zero net supply, C is total household consumption, c is consumption per capita $c = \frac{C}{N}$, and $L_A$and L_Y are total hours worked. Total labour supply in the economy is given by total hours, which is average hours H_W times working age population N: $L = L_Y + L_A = H_W \cdot N$. Now defining per capita assets as $d = D/N$, dividing the law of motion by N, and noticing that $\dot{D} = \dot{\left(dN\right)} = \dot{d}N + d\dot{N}$ yields the law of motion of per capita assets.

subject to the law of motion for per-capita assets, the leisure production function, and the no-Ponzi condition:

$$\dot{d} = d(r-n) + w \cdot H_W - c \tag{3.14}$$

$$l = H_L M^{\frac{1}{\nu-1}}$$

$$\lim_{t\to\infty} \left[d(t) \cdot \exp\left(- \int_0^t (r(s)-n) ds \right) \right] \geqslant 0$$

To solve the problem, I set up the Hamiltonian and take the optimality conditions, utilizing the Maximum Principle. The Hamiltonian is:

$$\mathcal{H}(d, c, H_L, \mu) = \log(c) + l + \mu \left(d(r-n) + w(1-H_L) - c \right)$$

And the associated optimality conditions are:

$$H_c = 0 \Rightarrow \frac{1}{c} - \mu = 0$$

$$H_h = 0 \Rightarrow M^{\frac{1}{\nu-1}} - \mu w = 0$$

$$H_d = -\dot{\mu} + (\rho - n)\mu \Rightarrow \mu(r-n) = -\dot{\mu} + (\rho-n)\mu$$

$$\lim_{t\to\infty} \left[d(t) \cdot \exp\left(- \int_0^t (r(s)-n) ds \right) \right] = 0$$

Combining these equations yields the two optimality conditions in the main text.

Share of workers in the R&D sector

The share s can be calculated from the free-entry condition. Free entry to the R&D sector ensures that profits in the R&D sector must be zero in equilibrium. That is, in equilibrium we have:

$$wL_A = P_A \dot{A} \tag{3.15}$$

Plugging in for the wage from the FOC of the retailer, and for the price of a patent from the no-arbitrage condition (1.22), and noting that $\dot{A} = A\gamma$ gives:

$$(1-\alpha)\frac{Y}{L_Y} \cdot L_A = \frac{\Pi}{r-n} \cdot A\gamma$$

Define $s = L_A/L$, and plug in for profits from (3.12) to get:

$$\frac{s}{1-s}(1-\alpha)Y = \frac{\gamma}{r-n} \cdot \alpha(1-\alpha-\alpha\chi)Y$$

Solving for $s$gives:

$$s = \frac{1}{1 + \frac{1-\alpha}{\alpha(1-\alpha-\alpha\chi)} \left(\frac{r-n}{\gamma}\right)} \tag{3.16}$$

Social planner problem in a static model

It is convenient to first separately analyze the problem of choosing the optimal quantity of each differentiated product $x(i)$. After noting that, in the static model, the total labor supply is just the hours worked of the representative household ($L_Y = H_W$), the problem amounts to the following maximization:

$$\max_{x(i)} \int_0^A x(i)^\alpha H_W^{(1-\alpha)} di - \int_0^A x(i) di.$$

Optimal $x(i)$ thus satisfies:

$$x(i) = \alpha^{\frac{1}{1-\alpha}} H_W \; \forall i, \tag{3.17}$$

which gives the expression for gross output in the optimal allocation:[45]

$$Y = A\alpha^{\frac{\alpha}{1-\alpha}} H_W.$$

Substituting this back into the objective function yields the maximization in two variables only:

$$\max_{X_B, H_L} log\left((1-H_L)A\left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right) - X_B\right) + H_L X_B^{\frac{1}{\nu-1}}. \tag{3.18}$$

Let $\hat{\alpha} = \alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}$. The main difficulty of this problem is that it may well be optimal for the planner to choose a corner solution in either of the two choice variables. *If* the solution for X_B and H_L was interior, it *would* satisfy the two first order conditions:

$$H_L: \; \frac{1}{c}\left[A\hat{\alpha}\right] = X_B^{\frac{1}{\nu-1}}$$

$$X_B: \; \frac{1}{c} = \frac{1}{\nu-1} H_L X_B^{\frac{2-\nu}{\nu-1}}.$$

Dividing one by the other yields the result that, in any interior optimal solution, H_L and X_B are proportional to each other:

$$X_B = \frac{1}{\nu-1} A\hat{\alpha} H_L \tag{3.19}$$

[45]Note that in the decentralized setting, the ratio between $X \equiv Ax$ and Y is α^2 , while in the planner's problem it is α, meaning that the in the market allocation too little of each of the consumer good is produced. This is a well-known result due to market power in the the monopolistically competitive consumer goods sector.

Equation (3.19) is a ray through the origin in the (H_L, X_B) space. Thus the corner solution $(H_L, X_B) = (0, 0)$ also satisfies (3.19), implying that equation (3.19) is satisfied at all times, whether at a corner or in an interior solution. It is thus possible to further reduce the dimensionality of the problem, by plugging equation (3.19) into the objective function. This finally yields a maximization problem in a single choice variable, H_L:

$$\max_{H_L} \log\left(A\hat{\alpha}(1 - \frac{\nu}{\nu - 1}H_L)\right) + H_L^{\frac{\nu}{\nu-1}}\left(\frac{A\hat{\alpha}}{\nu - 1}\right)^{\frac{1}{\nu-1}}. \tag{3.20}$$

Figure (1.12) in the main text characterizes the solution to this problem.

Appendices to Chapter 2

Appendix 2.A: Numerical solution to the HJB equation in continuous time

Derivations: finite difference method

The stationary equilibrium version of the HJB equation of a household holding assets a and in state s is

$$\rho W(a,s) = \max_c \ \{u(c) + W'(a,s)\,[w(a) + ra - c] + \sigma\,[W(a,s') - W(a,s)]\}$$

The first order condition is

$$u'(c) = W'(a,s)$$

which implies that optimal consumption satisfies

$$c = u'^{-1}(W'(a,s))$$

and the optimal asset holdings equal

$$s = w(a) + ra - c.$$

The finite difference method approximates the derivative of the value function W_a with the discretized differences on the grid of assets $\{a_i\}_{i=1}^{I}$, with distance between grid points denoted by Δa. It is possible to construct a forward and a backward difference, and we apply one or the other depending on whether the household is increasing or decreasing their assets at any given grid point. Using the shorthand notation $W_i \equiv W(a_i, s)$ we have:

$$W'_B = \frac{W_i - W_{i-1}}{\Delta a}$$

$$W'_F = \frac{W_{i+1} - W_i}{\Delta a}$$

The former (backward) approximation is used when optimal saving is negative and the assets are declining; the latter is used when saving is positive. There is an inherent circularity in the algorithm, however: to calculate savings and consumption, we need to approximate W', and which of the backward or the forward difference approximation to use depends on the optimal savings. To overcome this problem, construct two candidate optimal saving decisions for each grid point using both backward and forward approximation. That is, construct:

$$c_{i,B} = u'^{-1}(W'_{i,B})$$

$$s_{i,B} = w(a_i) + ra_i - c_{i,B}$$

and

$$c_{i,F} = u'^{-1}(W'_{i,F})$$

$$s_{i,F} = w(a_i) + ra_i - c_{i,F}$$

For all grid points where $s_{i,B}$ is negative, use W'_B; conversely, at all grid points where $s_{i,F}$ is positive, use W'_F.

To implement this, denote by $s_F^+ = max\{0, s_F\}$ and $s_B^- = min\{0, s_B\}$. The HJB equation can thus be written as:

$$\rho W_i = u(c_i) + \frac{W_i - W_{i-1}}{\Delta a} s_{i,B}^- + \frac{W_{i+1} - W_i}{\Delta a} s_{i,F+}^+ + \sigma \left[W_i(s') - W_i(s)\right]$$

Let W_i^n denote the n-th iteration on the value function. The update procedure (in the "implicit method") is given by:

$$\frac{W_i^{n+1} - W_i^n}{\Delta} + \rho W_i^{n+1} = u(c_i^n) + \frac{W_i^{n+1} - W_{i-1}^{n+1}}{\Delta a} s_{i,B}^{n,-} + \frac{W_{i+1}^{n+1} - W_i^{n+1}}{\Delta a} s_{i,F}^{n,+} + \sigma \left[W_i^{n+1}(s') - W_i^{n+1}(s)\right]$$

The intuition is that, when the value function converges, the equation beyond the first term on the left side holds exactly. The update is done based on the error. For a given guess of the value function W^n, equation above is linear, because savings s^n depends on guess W^n but is multiplied by the unknown W^{n+1}. That is why it is useful to write this equation with consumption and saving decisions determined by the current guess but the derivative of the value function to be determined by the solution to this equation (which will give us another iteration of the guess, W^{n+1}).

It is useful to rewrite this equation in the following way: collect the W terms with the same subscripts together

$$\frac{W_i^{n+1} - W_i^n}{\Delta} + \rho W_i^{n+1} = u(c_i^n) + W_{i-1}^{n+1} x_i + W_i^{n+1} y_i + W_{i+1}^{n+1} z_i + \sigma \left[W_i^{n+1}(s') - W_i^{n+1}(s)\right]$$

where $x_i = -\frac{s_{i,B}^{n,-}}{\Delta a}$, $y_i = \frac{s_{i,F}^{n,+}}{\Delta a} - \frac{s_{i,B}^{n,-}}{\Delta a}$ and $z_i = \frac{s_{i,F}^{n,+}}{\Delta a}$.

There are I equations like this for each of the state s. If there are two states, s and s', e.g. employed and unemployed, then there are $2I$ of equations in the system. This system can be re-written, in a matrix form, in the following way:

$$\frac{1}{\Delta}(\mathbf{W^{n+1}} - \mathbf{W^{n}}) + \rho\mathbf{W^{n+1}} = \mathbf{u}(\mathbf{c^n}) + \mathbf{B}\cdot\mathbf{W^{n+1}}$$

where $\mathbf{B} = \begin{pmatrix} \mathbf{A}(s) & \sigma\mathbf{I}_{I\times I} \\ \sigma\mathbf{I}_{I\times I} & \mathbf{A}(s') \end{pmatrix}$ and $\mathbf{A}(s) = \begin{pmatrix} y_1 & z_1 & 0 & \cdots & 0 \\ x_2 & y_2 & z_2 & \ddots & 0 \\ 0 & x_3 & y_3 & z_3 & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & 0 & x_I & y_I \end{pmatrix}$. The system can be written in a familiar way $Ax = b$ as follows:

$$((\frac{1}{\Delta} + \rho)\mathbf{I}_{2I\times 2I} - \mathbf{B})\mathbf{W^{n+1}} = \mathbf{u}(\mathbf{c^n}) + \frac{1}{\Delta}\mathbf{W^n}.$$

This can be solved for $\mathbf{W^{n+1}}$. This is done easily and quickly on a computer, because the system is linear and inverting the matrix on the left hand side of this equation is easy, because the $\mathbf{B}$ matrix is very sparse.

The algorithm to find equilibrium

To solve for the equilibrium, I develop the following algorithm, which builds on the techniques developed by Bardóczy (2017).

1. Set up grid for assets.
2. Guess r. Calculate m^* and $\bar{p}$ from 2.8.
3. Guess θ. This gives job finding and vacancy filling rates q and f, the unemployment rate u and vacancies v.
4. Solve unemployed and employed workers HJB equation using the finite difference method. This step is based on techniques developed in Achdou et al. (2017), and the details of the approach are outlined in the Appendix. In the first iteration, use a guess for the wage schedule, $w_0(a)$. This gives $W(a, s)$ and policy functions $c(a, s)$ and $\dot{a}(a, s)$.
5. Solve Kolmogorov Forward equation using the policy functions. Gives the distribution of assets across employed and unemployed workers, $g(a, s)$.
6. Solve firm's HJB. This gives $J(a)$.
7. Solve the bargaining problem to obtain the wage schedule via 2.35.
8. Check whether the Free Entry condition 2.15 holds by computing the difference the residual, FE.

9. Check whether asset market clears:
   (a) Compute asset demand as an integral of $g(a, e)$.
   (b) Compute dividend from $d = \pi' g(e) da - \xi v$
   (c) Compute asset supply as $K + p^e$ where $p^e = \frac{d}{r}$.
   (d) Compute excess supply of assets, XA.
10. If both FE and XA are zero, stop. If not, update r (based on XA) and θ (based on FE) and go back to step 4.

Appendix 2.B: Some indicative evidence on the wage schedule

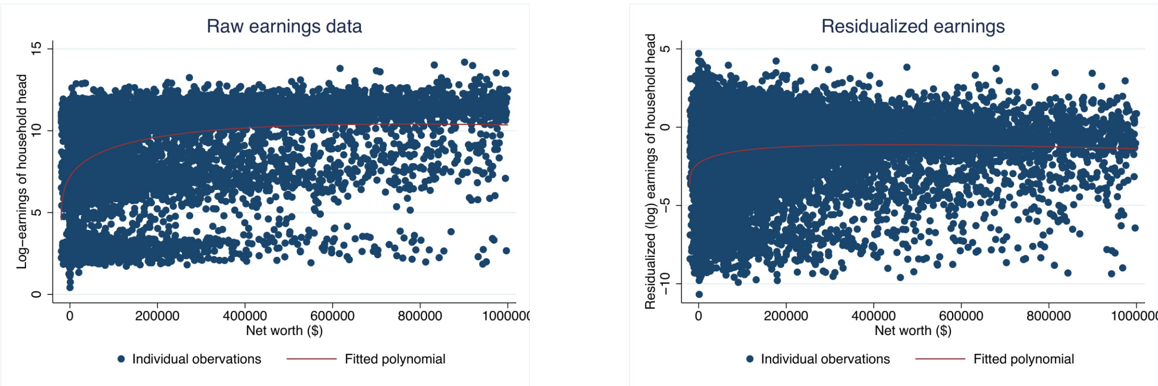

Figure 3.16: Wage schedule in the SCF

One source that allows for a simple and preliminary gauge at the shape of the wage schedule is the Survey of Consumer Finances, which contains micro data for individual asset holdings and individual earnings. The left panel of Figure (3.16) plots the individual net worth against (the log of) earnings using this source, and fits a polynomial function in this data. There is of course a lot of heterogeneity in the observed data, but an upward sloping and concave curve is clearly visible and identified. Positive slope of this curve persists into around the mean – 2.5 times the median, or \$170,000 – of net worth holdings, with the wage schedule flattening after that.

Perhaps a closer match with the theory is a version of this analysis based on residualized earnings that take out the variation related to some of the more obvious observed individual characteristics. The right hand side figure thus presents a figure where only the residuals of log-earnings are shown. The regression used is of the following form:

$$\log Earn = \alpha + \boldsymbol{\beta_1}[age,\ age^2,\ age^3] + \beta_2 sex + \beta_3 race + \beta_4 educ + \beta_5 jobtitle + \beta_6 industry + \epsilon$$

This regression soaks up 27% of the variation in log-earnings. Qualitatively, the wage schedule is robust to residualizing eanings this way. Quantitatively the slope appears to be a little shallower.

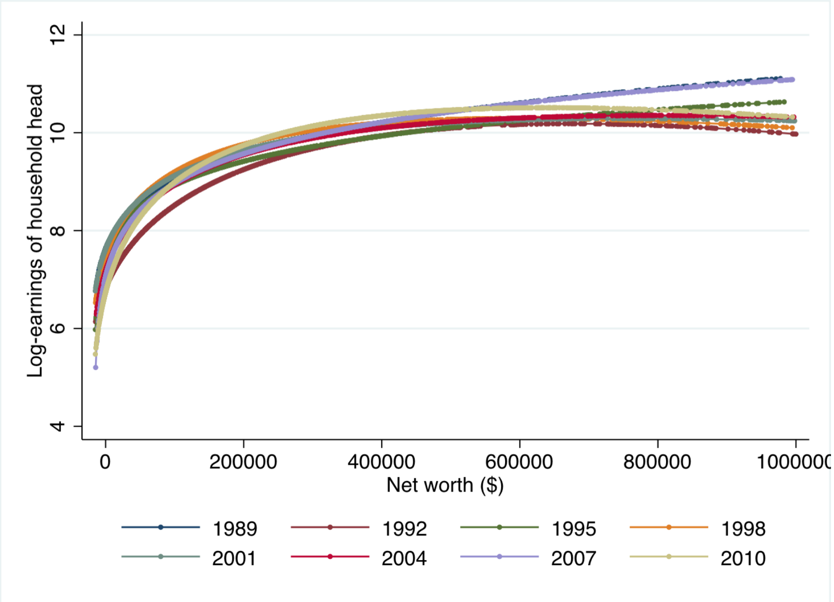

Figure 3.17: Fitted wage schedules across the years in the SCF

An alternative way to check the robustness of the finding that there is wage schedule in the data is a year-by-year analysis.

There are of course obvious limitations of this analysis. Assets are an endogenous result of earnings propensities and saving behavior, so the causality runs two ways.

Appendix 2.C: Transition path with the final steady state constrained by technology

When the technology threshold μ shifts to the right, the resulting long-run equilibrium can be one of a binding or non-binding constraint. The main text considered the case where in the final equilibrium the constraint was not binding. The two Figures in this Appendix plot the transition of the economy for when the new constraint is binding – see the top right panel in the second Figure in particular.

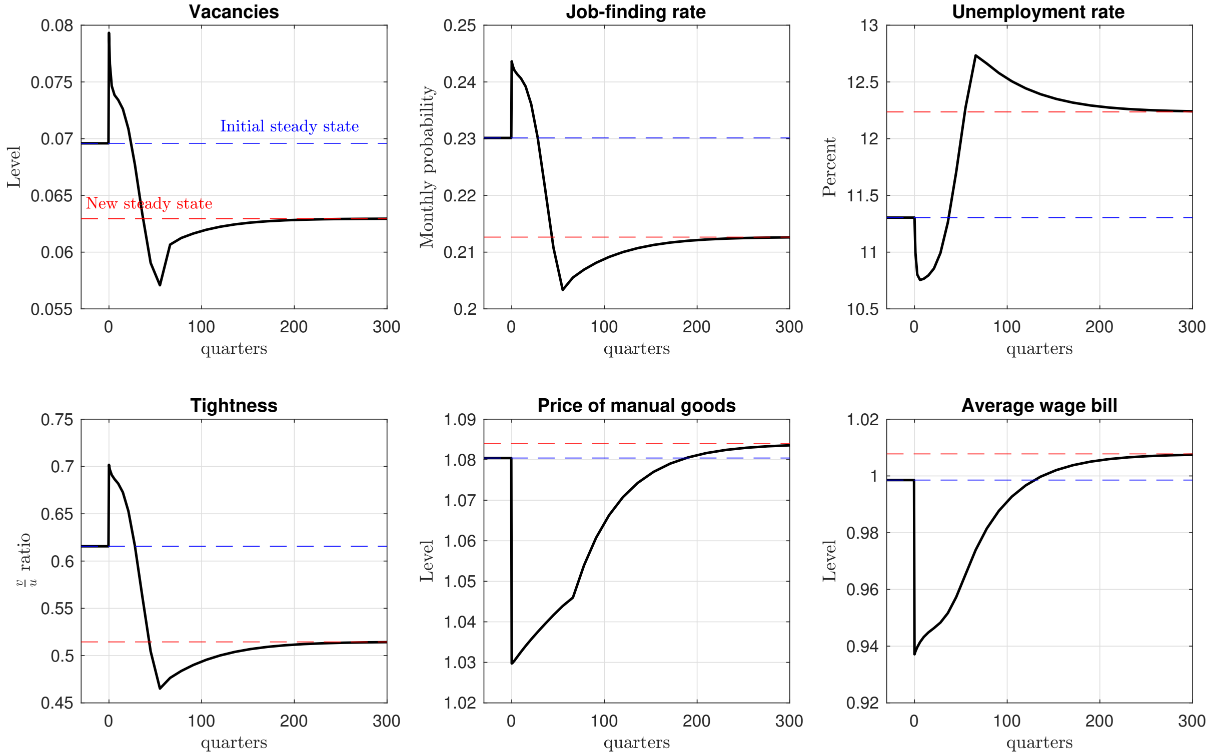

Figure 3.18: Transition path of the labor market variables following an increase in μ when the technological constraint is binding in the new equilibrium

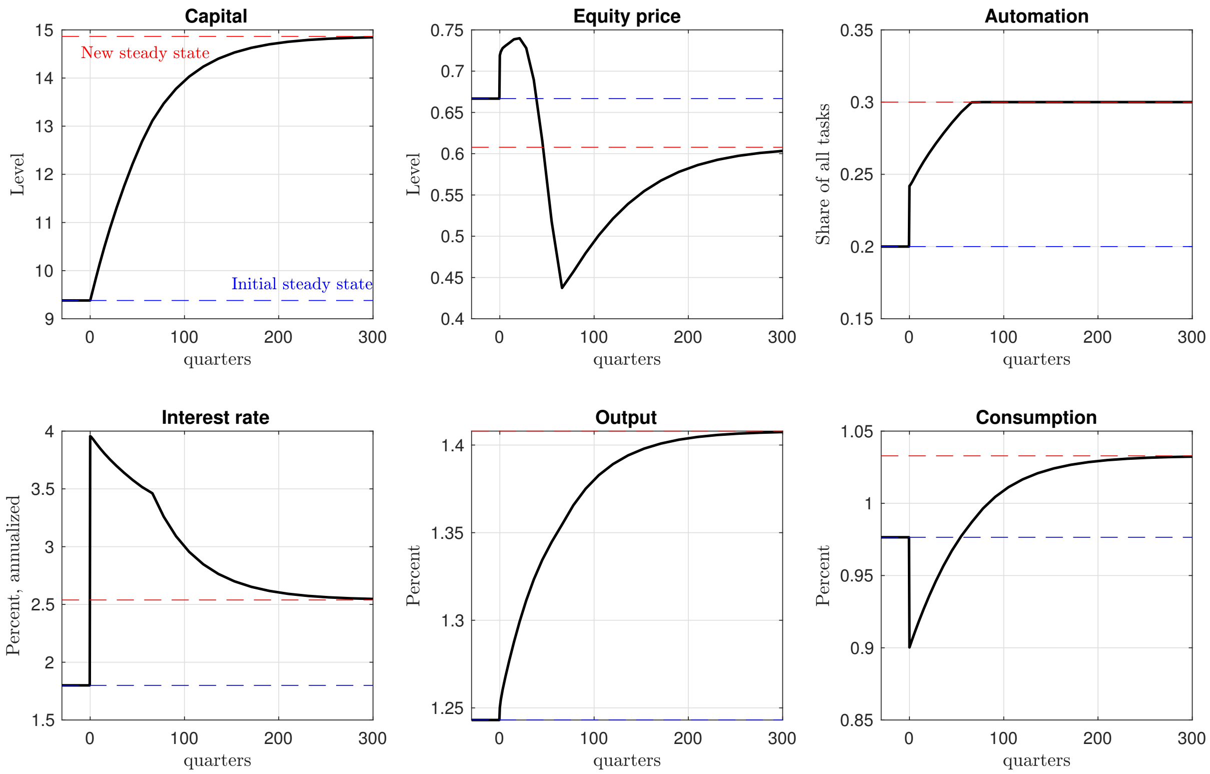

Figure 3.19: Transition path of the aggregate variables following an increase in μ when the technological constraint is binding in the new equilibrium

Appendices to Chapter 3

Appendix 3.A: Estimation of the state-space system

Data

All data are of quarterly frequency, spanning 1971Q1:2017Q4. All series used in this exercise are for a sample of OECD economies and are produced by the OECD and can be downloaded from the OECD website. For more details on the OECD data, see the OECD Economic Outlook data inventory.

To construct the observed advanced economy interest rate, we use long-term interest rates (10-year government bond yields; database here). As explained in the main text, we calculate the arithmetic average of these interest rates across the unbalanced panel of 36 OECD countries. Our results are unchanged if a median interest rate is used.

The inflation series is the non-food, non-energy consumer price index for OECD-total sample (database metadata available here). We construct the measure of inflation expectations as a moving average of observed core inflation rates over the past four quarters. This, of course, is not an ideal measure, but it follows past efforts in this literature and allows for estimation using the data from a large block of countries (alternative measures of inflation expectations for a large sample of countries going this far back are not available).

Finally, the GDP data cover the OECD-total sample. These are seasonally adjusted real (constant prices) measures, calculated using fixed 2010 PPPs, and expressed in 2010 US dollars. The series we use are calculated using the expenditure approach (OECD series code VPVOBARSA).

Estimation

The estimation is a recursive process that starts with a guess for the unconditional mean and variance of the unobservable state variables (such as the equilibrium real rate r^* or the level of potential output y^*) in the initial period. These are used to produce forecasts for the observables (such as inflation π or actual output y) next period. For every $t > 1$, the procedure consists of two steps. First, the update step changes the best guess for the unobservable states based on the forecast error on the set of observables. The direction of the revision is determined by the covariance of the observable and unobservable states, and the size of the update is determined by the variance of the observable (higher variance means there is more noise relative to signal on average, which reduces the size of the update). In the second step, this latest information is used to produce a new forecast for next period.

Following the classic notation of Hamilton (1994), we can write equations 3.5-3.3 in the state-space form as follows:

$$\mathbf{y}_t = \mathbf{A}'\mathbf{x}_t + \mathbf{H}'\xi_t + \mathbf{v}_t \tag{3.21}$$

$$\xi_t = \mathbf{F}\xi_{t-1} + \epsilon_t. \tag{3.22}$$

In this system, $\mathbf{x}_t$ is the vector of exogenous or lagged state variables; ξ_t is the vector of endogenous states; and $\mathbf{v}_t$ and ϵ_t are the vector of uncorrelated Gaussian disturbances. Specifically in this model, after substituting equation (3.1) into (3.5), (3.21) contains two observation equations (equations (3.5) and (3.6)) and (3.22) are the three state equations: (3.2), (3.3) and (3.4).

Following Holston, Laubach, and Williams (2017b) and as explained in detail in Holston, Laubach, and Williams (2017a), the estimation is carried out in three stages, building up to the full model. This is done to avoid the downward bias in the estimates of the standard deviations of the shocks to z and the trend growth rate. Instead of estimating these parameters directly, they are constructed from the first and second stages and imposed in the final estimation stage. For more details, see the referenced papers.

Appendix 3.B: Illustration of the difficulties estimating the link between government debt and R*

In this Appendix we showcase the difficulties of estimating the causal link between government debt and R* described in the main text. For the purpose of this illustration, we construct a panel of equilibrium interest rates and government debt / GDP ratios for four large developed economies: the US, Canada, Euro Area and the UK. We use the estimates of R* estimated by Holston, Laubach, and Williams (2017b), and study how these estimated equilibrium interest rates vary with the headline measure of government debt in these four economies. We run three regression specifications: the pooled regression, equivalent to OLS on the entire dataset, which effectively ignores the panel-structure of the data; the fixed effects or within regression, which controls for constant unobserved heterogeneity at the economy-level; and a 'between' regression, based on economy-level averages.[46]

Figures 3.20 and 3.21 contain the results. The difference between the two figures is the dependent variable: Figure 3.20 uses the estimate of R* from Holston, Laubach, and Williams (2017b) as a dependent variable, while Figure 3.21 uses their estimate of the unobserved component z (which excludes the effects of declining trend growth). The latter specification is motivated by the idea that it may be easier to deduce the relationship between R* and debt after accounting for trends in productivity. In each figure, the different-colored squares are the data points for the four economies, and the large blurbs denote the economy-level averages. The

[46] Our panel comprises of annual observations spanning 1961-2013 for the US, Canada and the UK, and 1972-2013 for the Euro Area. We proxy for the government debt in the EA using data for Germany.

Figure 3.20: Panel regressions of R* and government debt

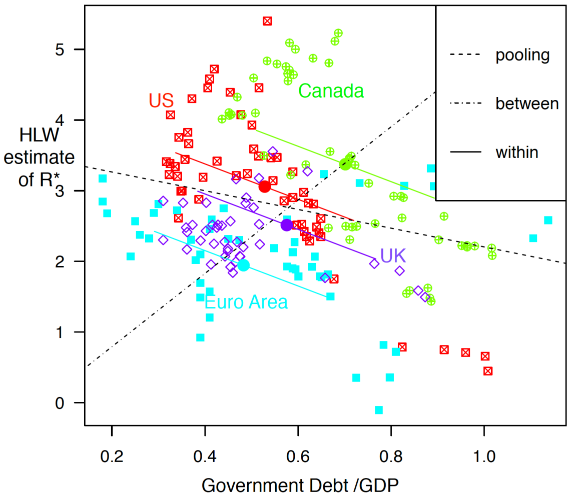

Source: Holston, Laubach, and Williams (2017b) and Jorda-Schularick-Taylor database Jordà, Schularick, and Taylor (2016).

Figure 3.21: Panel regressions of the z component of R* and government debt

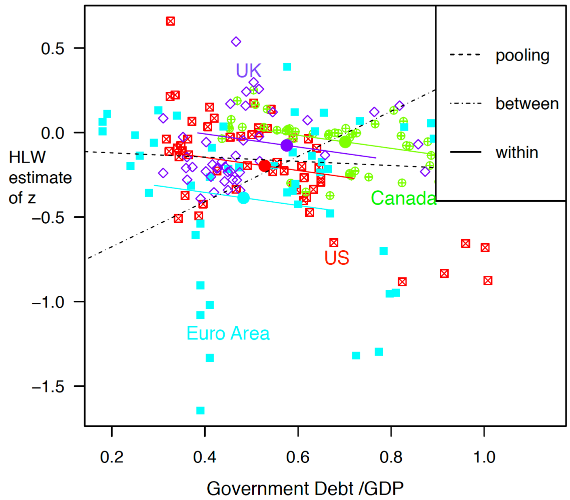

Source: Holston, Laubach, and Williams (2017b) and Jorda-Schularick-Taylor database Jordà, Schularick, and Taylor (2016).

three different kinds of lines show the estimated relationships: the broken line is the pooled regression, the solid colored lines show the fixed effects model, and the upward sloping dash/dot line shows the 'between' model. The notable result is that both in the pooled and fixed effects models the regression lines are downward sloping: the secular trend in interest rates, which coincided with increasing government debt, dominates these econometric estimates. Only the 'between' model detects a positive relationship between debt and R*, but the inference is extremely limited as the 'between' regression uses only four data points, each corresponding to one economy.[47]

Our simple exercise brings to the fore the difficult challenge that the empirical literature needs to overcome to uncover the link between debt and interest rates, namely that the secular fall in interest rates coincided with a rapid increase in advanced nations' public debt, making identification using measured macro data problematic. In the main text we discuss the papers that overcome the identification difficulties by including detailed measures of the output gap, inflation expectations and portfolio shifts in their regressions or use fiscal forecasts rather than the realized debt/GDP ratios to alleviate the concern that cyclical variation drives the results.

[47]Taken with an appropriately sized pinch of salt, the 'between' estimate suggests that a 1pp increase in government debt / GDP ratio raises the equilibrium rate of interest by about 5bps.

Appendix 3.C: The life-cycle model derivations

Retirees

The first order conditions of the retiree's problem yield the Euler Equation:

$$C_{t+1}^{rjk} = (R_{t+1}\beta)^{\sigma} C_t^{rjk}$$

Denoting by $\epsilon_t \pi_t$ the retiree's marginal propensity to consume out of wealth[48], we can write down retiree's consumption function as:

$$C_t^{rjk} = \epsilon_t \pi_t (R_t/\gamma) A_t^{rjk}$$

Plugging this expression into the Euler Equation yields the expression for the evolution of the retiree's MPC:

$$\epsilon_t \pi_t = 1 - (R_{t+1}^{\sigma-1}\beta^{\sigma}\gamma)\frac{\epsilon_t \pi_t}{\epsilon_{t+1}\pi_{t+1}}. \tag{3.23}$$

Workers

The Euler Equation from the worker's problem is:

$$\omega C_{t+1}^{rjk} + (1-\omega)\Lambda_{t+1} C_{t+1}^{rj(t+1)} = (R_{t+1}\Omega_{t+1}\beta)^{\sigma} C_t^{wj}$$

where Λ is the marginal rate of substitution across consumption while being a worker and a retiree, and Ω is a weighing factor which captures the fact that workers discount future more: $\Omega_{t+1} = \omega + (1-\omega)\epsilon_{t+1}^{\frac{1}{1-\sigma}}$.[49]

Denoting the MPC of the worker by π, and conjecturing that the consumption function takes the form:

$$C_t^{wj} = \pi_t (R_t A_t^{wj} + H_t^j + S_t^j)$$

(where H stands for human wealth and S is social security wealth, given respectively by $H_t^j = \sum_{\nu=0}^{\infty} \frac{W_{t+v} - T_{t+\nu}}{\prod_{z=1}^{v} R_{t+z}\Omega_{t+z}/\omega}$ and $S_t^j = \sum_{\nu=0}^{\infty} \frac{E_{t+\nu}^j}{\prod_{z=1}^{v} R_{t+z}\Omega_{t+z}/\omega}$), we obtain the time path of worker's MPC :

$$\pi_t = 1 - (R_{t+1}\Omega_{t+1})^{\sigma-1}\beta^{\sigma}\frac{\pi_t}{\pi_{t+1}}.$$

Aggregation

Marginal propensities to consume are the same across all retirees, so we can just add up their consumptions to get the aggregate consumption function. With slight abuse of notation, denoting now by C_t^r, A_t^r and S_t the aggregate variables, we have that:

$$C_t^r = \epsilon_t \pi_t (R_t A_t^r + S_t)$$

[48] Reasons for this notation will become clear momentarily.

[49] For the complete derivation of worker's Euler Equation, see the Appendix of Gertler (1999).

where social security wealth is given by the discounted sum of social security payments:

$$S_t = \sum_{\nu=0}^{\infty} \frac{E_{t+\nu}}{\prod_{z=1}^{\nu}(1+n)R_{t+z}/\gamma}.$$

The evolution of $\epsilon_t\pi_t$ is governed by the Euler Equation of the retiree, given by equation (3.23).

Marginal propensities to consume are the same across all workers, so we can add individual worker consumption across individuals to obtain their aggregate consumption function. This will depend not only on the aggregate asset wealth but also on the aggregate human wealth and aggregate social security wealth:

$$C_t^w = \pi_t(R_t A_t^w + H_t + S_t^w).$$

The aggregate human wealth is given by

$$H_t = \sum_{\nu=0}^{\infty} \frac{N_{t+\nu}W_{t+\nu} - T_{t+\nu}}{\prod_{z=1}^{\nu}(1+n)R_{t+z}\Omega_{t+z}/\omega}.$$

Human wealth is a discounted sum of the economy-wide net-of-tax wage bill. The discount rate that is applied to the aggregate wage bill is the product of the gross population growth rate and the rate at which individual workers discount their labor income. The importance of the generation currently alive declines over time, however – they get replaced by newly born generations. So from the point of view of the current generation, the human wealth is discounted more heavily than in the infinite horizon case – the gross population growth rate $(1+n)$ enters the discount factor. In total, therefore, there are three distinct factors in the life-cycle setting that raise the discount rate on future labor income (relative to the infinite horizon case). They are: (1) finite expected time spent working (reflected by the presence of ω in the discount rate); (2) greater discounting of the future owing to expected finiteness of life (reflected by the presence of Ω); and (3) growth of the labor force (reflected by the presence of $(1+n)$).

Social security wealth of the workers is:

$$S_t^w = \sum_{\nu=0}^{\infty} \frac{(1-\omega)\omega^{\nu} N_t \left(\frac{\epsilon_{t+\nu+1} \frac{S_{t+\nu+1}}{\psi N_{t+\nu+1}}}{R_{t+\nu}\Omega_{t+\nu}} \right)}{\prod_{z=1}^{\nu} R_{t+z}\Omega_{t+z}}.$$

The numerator of the sum on the right hand side is a time-$t+\nu$ capitalized value of the social security payments to all the individuals who were in the workforce at t and retire at $t+\nu+1$. The total social security wealth is just the infinite sum of the discounted value of these capitalized payments.

Denoting by λ the share of assets held by retirees, we can add the two aggregate consumptions above to get aggregate consumption in the main text:

$$C_t = C_t^w + C_t^r = \pi_t \left\{(1-\lambda_t)R_t A_t + H_t + S_t^w + \epsilon_t(\lambda_t R_t A_t + S_t^r)\right\}$$

The novel feature is the presence of λ. Because the MPC of retirees is higher than MPC of workers ($\epsilon > 1$), higher λ raises aggregate consumption. So transferring resources across the demographic groups changes overall demand.

The evolution of total wealth of retirees is the sum of return on their wealth from last period plus what the newly retired bring in:

$$\lambda_{t+1}A_{t+1} = \lambda_t R_t A_t - C_t^r + (1-\omega)[(1-\lambda_t)R_t A_t + W_t - C_t^w]$$

From this we get the explicit expression for the evolution of the retiree share:

$$\lambda_{t+1} = \omega(1-\epsilon_t \pi_t)\lambda_t R_t \frac{A_t}{A_{t+1}} + (1-\omega).$$

Appendix 3.D: The model of precautionary savings: derivations and equilibrium

Equilibrium in the asset market

Equilibrium in the asset market requires that asset demand (households' desired asset holdings) equals asset supply (firms' capital plus government bonds):

$$A_t = K_t + B_t.$$

Because of exogenous technological progress, the equilibrium in this economy will be characterized by a balanced growth path along which the aggregate variables – K_t, w_t and Y_t – grow at rate η. Below we show how to rewrite the model with variables normalized by GDP, thus making it stationary.

Transformation into a stationary model

Growth is exogenous, driven by increases in labor augmenting technology: $x_t = e^{\eta t}x_0$. In the balanced growth equilibrium w_t, Y_t and K_t will be growing at rate η whereas the interest rate will be constant.

Let $k_t = K_t/Y_t$, $\tilde{w}_t = w_t/Y_t$, $\tilde{c}_t = c_t/Y_t$, $\tilde{a}_t = a_t/Y_t$, $\tau_t = T_t/Y_t$, $b_t = B_t/Y_t$, $\bar{a}_t = A_t/Y_t$, $tran_t = \frac{TR_t}{Y_t}$ denote the normalized variables.

Households. We begin by rewriting the consumer problem. First, note that $c_t = Y_t\tilde{c}_t$ and $Y_t = e^{\eta t}Y_0$. We can rewrite the integral as:

$$\int_0^\infty e^{-\rho t}\frac{c_t^{1-\gamma}}{1-\gamma}dt = \int_0^\infty e^{-\rho t}\frac{(e^{\eta t}Y_0\tilde{c}_t)^{1-\gamma}}{1-\gamma}dt = \frac{Y_0^{1-\gamma}}{1-\gamma}\int_0^1 e^{-(\rho-(1-\gamma)\eta)t}\tilde{c}_t^{1-\gamma}dt.$$

The original budget constraint is:

$$\dot{a}_t = (1-\tau)w_t e_t + (1-\tau)r_t a_t - c_t.$$

Dividing through by Y_t:

$$\frac{\dot{a}_t}{Y_t} = \frac{(1-\tau)w_t e_t}{Y_t} + \frac{(1-\tau)r_t a_t}{Y_t} - \frac{c_t}{Y_t}.$$

Note that

$$\dot{a}_t = \frac{\partial a_t}{\partial t} = \frac{\partial \tilde{a}_t}{\partial t} Y_t + \tilde{a}_t \frac{\partial Y_t}{\partial t} = \dot{\tilde{a}}_t Y_t + \tilde{a}_t Y_0 \eta e^{\eta t}.$$

It follows that

$$\frac{\dot{a}_t}{Y_t} = \dot{\tilde{a}}_t + \tilde{a}_t \eta.$$

Thus the budget constraint in transformed variables is:

$$\dot{\tilde{a}}_t = (1-\tau)\tilde{w}_t e_t + ((1-\tau)r - \eta)\tilde{a}_t - \tilde{c}_t$$

And the transformed problem of the household is:

$$\max_{\{\tilde{c}_t\}} \mathbb{E}\, \frac{Y_0^{1-\gamma}}{1-\gamma} \int_0^\infty e^{-(\rho-(1-\gamma)\eta)t} \tilde{c}_t^{1-\gamma} dt$$

subject to

$$\begin{aligned}
&\dot{\tilde{a}}_t = (1-\tau)\tilde{w}_t e_t + ((1-\tau)r - \eta)\tilde{a}_t - \tilde{c}_t \\
&\tilde{c}_t \geqslant 0 \\
&\tilde{a}_t \geqslant 0.
\end{aligned}$$

This is a standard optimal control problem. Because the individual problem is recursive, its stationary version can be summarized with a Hamilton-Jacobi-Bellman (HJB) equation:[50]

$$(\rho-(1-\gamma)\eta)v_j(\tilde{a}) = \max_{\tilde{c}} \left\{ \frac{\tilde{c}^{1-\gamma}}{1-\gamma} + v_j'(\tilde{a})((1-\tau)\tilde{w}e_j + ((1-\tau)r-\eta)\tilde{a} - \tilde{c}) + \sum_{i\neq j} P_{j,i} v_i(\tilde{a}) - P_{j,j} v_j(\tilde{a}) \right\} \tag{3.24}$$

where the variables with a tilde are normalized by GDP, and $P_{j,i}$ is the Poisson intensity of a change from state $e = z_j$ to state $e = z_i$. This equation has a natural economic interpretation, related to the intuition from the asset pricing literature: the required return to an asset equals the dividend plus the change in value. The left hand side of the equation is the instantaneous required return to holding assets $\tilde{a}$ in state j: it is the effective discount rate (i.e. the return $\rho - (1-\gamma)\eta$) times the value function. The first term on the right is the 'dividend': the stream of consumption utility sustained by the given level of asset holdings. The remaining terms

[50] We focus our attention on stationary equilibria, so that the value function or any other variable in the HJB equation do not depend on time.

denote the instantaneous changes in value, due to asset accumulation and a possibility of a Poisson event that changes the state from j to i.

We solve the problem summarized in equation (3.24) by deriving the first order and the boundary conditions. The first order condition is obtained by differentiating (3.24) with respect to $\tilde{c}$:

$$\tilde{c_j}^{-\gamma} = v_j'(a). \tag{3.25}$$

One of the advantages of the continuous time formulation is that equation (3.25) holds at the constraint. We can use this to derive the boundary condition:[51]

$$v_j'(\underline{a}) \geqslant ((1-\tau)\tilde{w}e_j + ((1-\tau)r - \eta)\tilde{a})^{-\gamma}. \tag{3.26}$$

Using these two conditions we solve the HJB equation (3.24) utilizing the methods described in detail in Achdou et al. (2017).

Production. All markets are competitive. The aggregate production function is

$$Y_t = AK_t^{\alpha}(x_t L_t)^{1-\alpha}$$

where $x_t = e^{\eta t}x_0$ is the process for exogenous labor augmenting technical progress, which grows at rate η.[52] Because we normalized population to be of measure one, we have $L_t = 1$. Capital and labor demands are pinned down by the usual first order conditions:

$$\alpha A K_t^{\alpha-1} x_t^{1-\alpha} = r + \delta \tag{3.27}$$

$$(1-\alpha)AK_t^{\alpha} x_t^{1-\alpha} = w_t. \tag{3.28}$$

Capital demand is

$$K_t = \left(\frac{\alpha A}{r+\delta}\right)^{\frac{1}{1-\alpha}} x_t.$$

Dividing through by Y_t we get

$$k_t = \left(\frac{\alpha A}{r+\delta}\right)^{\frac{1}{1-\alpha}} \frac{x_t}{Y_t}.$$

Because x and Y both grow at rate η, this is the same as:

$$k_t = \left(\frac{\alpha A}{r+\delta}\right)^{\frac{1}{1-\alpha}} \frac{x_0}{Y_0}.$$

[51] To see this, note that equation (3.25) holding at the constraint implies that $\tilde{c_j}(\underline{a})^{-\gamma} = v_j'(\underline{a})$ where the notation $c_j(\underline{a})$ stands for the policy function in state e_j and assets $\underline{a}$. At the borrowing constraint saving cannot be negative (otherwise the constraint would be breached), so that we must have $\dot{\underline{a}} = (1-\tau)\tilde{w}e_j + ((1-\tau)r - \eta)\tilde{\underline{a}} - \tilde{\underline{c}} \geqslant 0$, which implies $(1-\tau)\tilde{w}e_j + ((1-\tau)r - \eta)\tilde{\underline{a}} \geqslant \tilde{c}(\tilde{\underline{a}})$. This, together with the concavity of utility function and equation (3.25), imply the boundary condition (3.26).

[52] To see this, note that $\frac{\dot{x}_t}{x_t} = \frac{\partial x_t}{\partial t}\frac{1}{x_t} = \eta$.

Similarly, labor demand is given by:

$$(1-\alpha)Ak_t^{\alpha}\left(\frac{x_0}{Y_0}\right)^{1-\alpha} = \tilde{w}_t.$$

When solving the model, we normalize the starting values x_0 and Y_0 to unity.

Government. Using the homogeneity of the production function, we can rewrite the government budget constraint (3.8) as follows. By Euler's Theorem we have:

$$w_t = Y_t - (r+\delta)K_t.$$

This and the asset market clearing $K + B = A$ together imply that (3.8) becomes:

$$\dot{B}_t = G_t + TR_t + rB_t - \tau(Y_t - (r+\delta)K_t + rK_t + rB_t).$$

Simplifying:

$$\dot{B}_t = G_t + TR_t + (1-\tau)r_t B_t - \tau(Y_t - \delta K_t).$$

Dividing through by Y_t and rearranging we get the transformed government budget constraint:[53]

$$g_t + tran_t + ((1-\tau)r_t - \eta)b_t - \dot{b}_t = \tau(1 - \delta k_t)$$

In steady state, government debt / GDP ratio is constant, so that for given values of r, b and g, the government budget constraint pins down the tax rate:

$$\tau = \frac{g + tran + b(r-\eta)}{1 - \delta k + rb}.$$

In the main text we set $TR_t = 0\ \forall t$.

Stationary Equilibrium

The following set of equations fully characterizes the stationary equilibrium:

- The HJB equation summarizing the household's problem:

$$(\rho - (1-\gamma)\eta)v_j(\tilde{a}) =$$

$$= \max_{\tilde{c}}\left\{\frac{\tilde{c}^{1-\gamma}}{1-\gamma} + v_j'(\tilde{a})((1-\tau)\tilde{w}e_j + ((1-\tau)r - \eta)\tilde{a} - \tilde{c}) + \sum_{i\neq j} P_{j,i}v_i(\tilde{a}) - P_{j,j}v_j(\tilde{a})\right\}$$

- The Kolmogorov Forward Equations which characterize the distributions of workers in

[53] Note that $\frac{\dot{B}_t}{Y_t} = \frac{\frac{\partial \dot{B}_t}{\partial t}}{Y_t} = \frac{\dot{b}_t Y_t + b_t \dot{Y}_t}{Y_t} = \dot{b}_t + b_t\eta.$

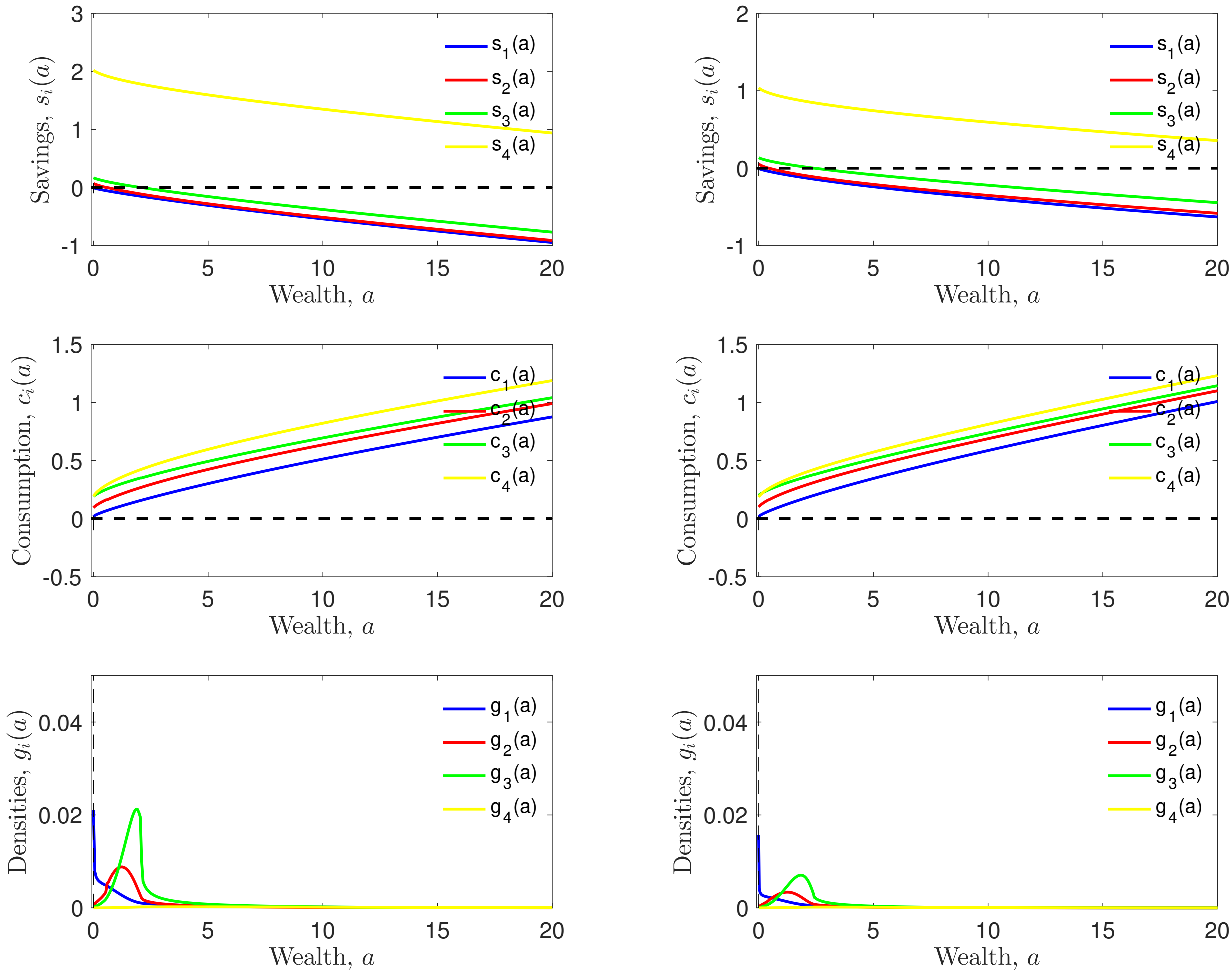

Figure 3.22: Stationary equilibrium with low (left) and high debt (right)

the four income states. In stationary equilibrium these take the following form:

$$0 = -\frac{d}{da}[s_j(a)g_j(a)] - \lambda_j g_j(a) + \sum_{i \neq j} \lambda_i g_i(a)$$

where $s_j(a)$ is the saving policy function from the HJB equation and $g_j(a)$ denotes the distribution (density) of type-j worker, so that

$$\int_{\underline{a}}^{\infty} (g_1(a) + g_2(a) + g_3(a) + g_4(a))\, da = 1, \quad g_j \geqslant 0\, \forall j.$$

- The asset market clearing condition (expressed using the transformed variables):

$$\bar{a} = \sum_j \int_{\underline{a}}^{\infty} a g_j(a) da = k + b.$$

Figure 3.22 shows the consumption and saving policy functions as well as the stationary distributions of agents across the asset space in the low- and high-debt equilibria described in the main text.

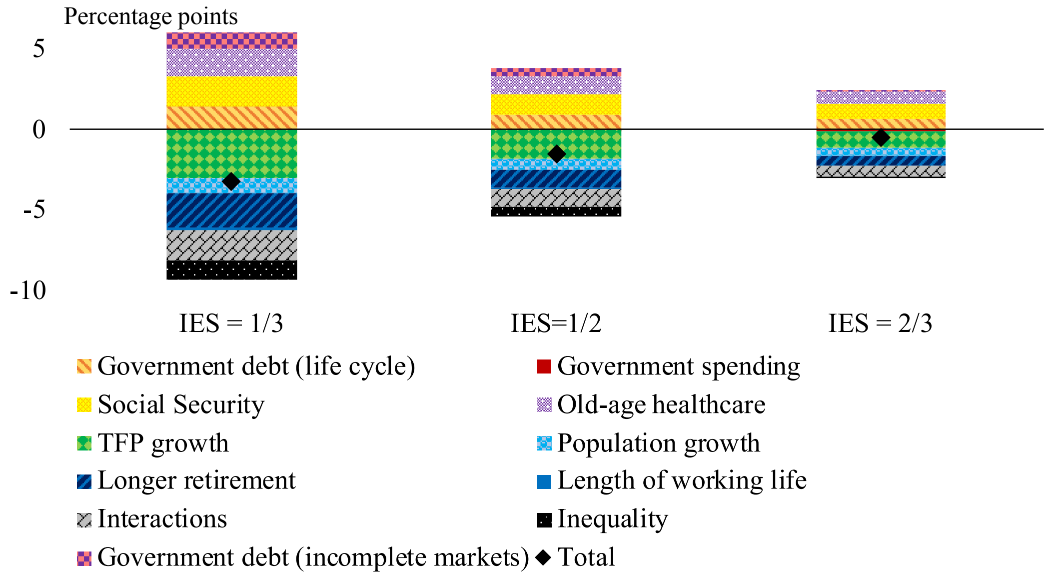

Figure 3.23: Factors driving the decline in the equilibrium real interest rate in the life-cycle model between 1970-2018, for different values of the IES.

Appendix 3.E: Sensitivity of the model-based results to alternative parametrization

The key parameter that determines the overall sensitivity of the interest rate to long-term fiscal stance as well as other secular trends, such as technological and demographic change, is the intertemporal elasticity of substitution, $\frac{1}{\gamma}$. In general, the higher this elasticity, the smaller the impact of a change in the macroeconomic environment on the interest rate. The intuition is that, if consumers are very willing to substitute consumption across time, smaller changes in the interest rate will be sufficient to induce them to do so. So less of a change in the interest rate will be required to restore equilibrium.

Our simulations reported in the main text are based on the parametrization in which this elasticity is set equal to $\frac{1}{2}$, which is a standard value used in many macroeconomic models. It is also the average value of the estimated elasticity across a large number of studies described in a comprehensive review by Havranek, Horvath, Irsova, and Rusnak (2015). Still, some models assume a different elasticity. For example, the calibration of the Smets and Wouters (2007) model assumes that $IES = \frac{2}{3}$, and calibrations with lower IES are also common. To illustrate the sensitivity of our results, we now present the results of a robustness exercise in which we simulate the models under these alternative parametrizations.

Figure 3.23 illustrates how much of the decline in real rates the models can account for depending on the calibration of the IES. As we highlighted in the main text, under the central calibration, the models under-explain the overall decline. However, the sensitivity analysis shows that relatively small changes to the value of the parameters can result in large changes in the contributions of the different explanations. The conclusion that parameter uncertainty plays an important role is in line with the results of the detailed study of these sensitivities presented in Ho (2018).