Macroeconomic Effects of Financial Integration, Demographic Aging and Automation Technology

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch die

Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von

Katja Verena Mann

aus Soest

Bonn 2018

Dekan: Prof. Dr. Daniel Zimmer, LL.M.
Erstreferent: Prof. Dr. Jürgen von Hagen
Zweitreferent: Prof. Dr. Keith Kuester

Tag der mündlichen Prüfung 6.7.2018

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn
http://hss.ulb.unibonn.de/diss-online elektronisch publiziert.

Acknowledgements

First of all, I would like to thank my first supervisor Jürgen von Hagen, who has been supportive throughout all stages of my PhD. I not only benefited from his vast expertise in international macroeconomics, but also learnt a lot about critical thinking, which has shaped my researcher identity. I would like to thank my second supervisor Keith Kuester for his advice on anything from vague research ideas to concrete technical questions, and for his extraordinary support on the job market.

I am thankful to my co-authors Margaret Davenport and Lukas Püttmann for many productive (and some not so productive) hours spent together, for suffering through every drawback and celebrating every success. I will remember our discussions as some of the most delightful moments of my PhD.

My work has benefited tremendously from the open and collegial atmosphere at the Bonn Graduate School of Economics. For their advice on Chapters 1 and 2, I would in particular like to thank Christian Bayer, Thomas Hintermaier, Fabian Kindermann and Jörg Breitung. Chapter 3 has benefited from numerous discussions with Michael Böhm, Simon Jäger and Hans-Martin von Gaudecker, who I would like to thank not just for their feedback but also for believing in this project at all times. I am further grateful to Michael Evers for his advice on many issues related to research and academia.

I would like to thank Pierre-Olivier Gourinchas for inviting me to UC Berkeley and for deepening my understanding about international macroeconomics. My stay at UC Berkeley has inspired my research well beyond this dissertation. I would like to thank Cédric Tille for his advice on Chapter 2 and for supporting my career.

I would not have been able to complete my PhD without my fellow students at the Bonn Graduate School of Economics. I am particularly grateful to Jens Herold and Lukas Püttmann for being supportive and entertaining office mates, to the BGSE women's group for moral support and an exchange of experiences, to Alisa Spiegel for our joint first-year experience and to my fellow job marketeers Nina Bobkova, Yao Chen and Felix Ward.

Last but not least, I would like to thank my parents Gisela and Ortwin Mann for – in the words of Goethe – giving me roots and wings, and for their constant support. Thanks also to many wonderful friends, whose love and support over these last five years has been invaluable.

Contents

List of Tables

List of Figures

Introduction

This thesis combines work on three important long-run trends and their macroeconomic implications: Financial integration, demographic aging and the use of automation technology in the production process. The first chapter looks at the effects of financial integration – a country's accumulation of external assets and liabilities – on the allocation of capital across economic sectors. It shows how international capital flows are driven by differences in the development of countries' financial systems. An alternative explanation for international capital flows is provided in the second chapter. Regional differences in the age structure of the population are shown to generate cross-country differences in the demand for safe and risky assets and – in a financially integrated world – a risk asymmetry in external asset positions. Chapter three focuses on a recent technological trend: advances in automation technology. It assesses the labor market effects of automation by means of a novel patent-based measure and finds overall employment gains. All chapters have in common that the phenomena they study are not just important today, but will likely become even more so over the next decades. Therefore, this thesis offers not just relevant policy advice, but also a research agenda for analyzing the future of capital and labor markets. Next to covering a broad range of topics, this thesis also applies various methodologies, in particular structural and analytical modeling and panel data econometrics. The following paragraphs describe the chapters in more detail.

CHAPTER 1: In "Does Foreign Capital Go Where the Returns are? Financial Integration and Capital Allocation Efficiency" I ask whether, and under which conditions, financial integration improves the allocation of capital within an economy. I link this question to observed international differences in the price level of consumption relative to investment, which serve as a proxy measure of cross-sectoral capital allocations. In contrast to the existing literature, which explains international differences in relative prices primarily by different productivity levels, my explanation centers on cross-country differences in financial development and I show how relative prices are affected by international capital flows.

Frictions in a country's financial system have a more distortive effect on the investment goods sector than on the consumption goods sector, as this sector is more dependent on external capital. Therefore, in a country with low financial development, an insufficient amount of capital will be allocated to the investment goods sector, resulting in a lower price level of consumption relative to investment in autarky. If an economy integrates with a financially more developed rest of the world, capital will flow into the investment goods sector and out of the consumption goods sector. As a result, the domestic capital allocation will become more efficient and the relative price will increase. The opposite should be observed in a country that integrates with a less developed rest of the world. Overall, financial integration implies converging relative prices across regions. I formalize this mechanism in an analytical model of a small open economy with two sectors and provide empirical support through a panel data analysis that covers 113 countries from 1996 to 2010.

CHAPTER 2: "External Asset Positions, Demography and Life-cycle Portfolio Choice" establishes international differences in demographic aging as a new explanation for external imbalances in safe and risky assets. This chapter is joint work with Margaret Davenport.
We document that vis-à-vis a group of developed European economies, the United States hold, on net, risky assets alongside safe liabilities. While existing explanations focus on characteristics of the financial system, we explore demographic differences as a potential driver for this risk asymmetry. The population in the European countries is already older than the US population, and is projected to age at a faster speed over the next decades.
In our structural model, the key element is the life-cycle savings behavior of households. Individual preferences for portfolio risk decline with age. We build this stylized fact into a model with overlapping generations, where individuals can adjust their portfolio at each age. We identify various channels through which the age structure of an economy affects the aggregate demand for safe and risky assets. Jointly, these channels imply that the younger region will export safe assets and import risky assets from the older region. While our simulated external asset positions match the data for 1990 to 2015 on average, we also predict the risk asymmetry to persist over the next decades. Additionally, in our simulation demographic aging generates a decline in both the safe and risky rate of return which is similar to the data.

CHAPTER 3: The advances in automation technology and their potential effects on employment are much discussed. There are worries that automation may lead to job losses and that the effects are unevenly distributed across occupations or sectors. While there exists a large literature that addresses these questions, the findings are partly contradictory. In "Benign Effects of Automation: New Evidence from Patent Texts", Lukas Püttmann and I argue that this may be due to difficulties in measuring automation. We provide a novel indicator of automation, which measures automation technology as the outcome of an innovative process, and which is both granular and comprehensive.
Our proposed measure is based on patent grant texts, which we classify into automation and non-automation innovations. We use a machine-learning algorithm to analyze the texts of all patents that were granted in the United States between 1976 and 2014. According to our classification, both the absolute and the relative number of automation patents have increased strongly over time. We link the automation patents to the industries where they are likely to be applied, and – through the local industry structure – to US commuting zones. In a panel data analysis, we find that automation has a positive net effect on employment in the United States. Distinguishing between manufacturing and non-manufacturing industries, we show that the gains from automation are unevenly distributed: the service sector experiences job gains, whereas manufacturing jobs are destroyed. Automation thus fuels structural change.

Chapter 1

Does Foreign Capital Go Where the Returns Are? Financial Integration and Capital Allocation Efficiency

This paper asks whether financial integration leads to a more efficient allocation of capital within economies. I build a model of a small country with an investment goods sector and a consumption goods sector. Frictions in the domestic financial system affect the investment goods sector relatively more, so that the cross-sectoral allocation of capital is distorted in autarky. When the economy integrates with a financially more developed rest of the world, capital flows into the investment goods sector and out of the consumption goods sector. In consequence, the capital allocation within the economy improves. The opposite holds if the rest of the world is less financially developed. Overall, capital allocations become more similar across countries. I test the model implications empirically using the price level of consumption relative to investment as a measure of capital allocation. A panel data analyses for 113 countries from 1996 to 2010 lends support to the model predictions.

1. Introduction

The benefits of financial integration[1] are heavily debated. The literature studying international capital flows in relation with economic growth has provided mixed results. While some authors have found positive growth effects (Bussiere and Fratzscher, 2008; Mody and Murshid, 2005; Quinn and Toyoda, 2008), others found no or negative effects (Alesina and Milesi-Ferretti, 1994; Rodrik, 1998; Edison, Levine, Ricci, and Sløk, 2002). Still others present mixed results.[2] Despite this disagreement, considerably less research has been done on the specific channels through which financial integration affects growth and the real economy more generally. In particular, international capital flows may influence the allocation of capital across economic sectors. Depending on where international investments are directed,

[1]Throughout the paper, this term refers to *de facto* financial integration, the actual amount of external assets and liabilities that a country holds. A contrasting concept is *de jure* financial integration, the degree to which capital account openness is allowed for or restricted by law. Section 3 provides a more detailed discussion.

[2]For a survey of the literature, see Kose, Prasad, Rogoff, and Wei (2010).

financial integration could thus make the economy more or less productive. The current paper tries to shed light on this issue by asking whether and under which conditions financial integration improves the allocation of capital within economies.
To understand the link between international finance and the domestic capital allocation, I argue that domestic financial institutions are of central importance. It has been shown by Buera, Kaboski, and Shin (2011) that the development of the financial system affects the cross-sectoral allocation of capital in a closed economy. In economies with a poorly developed financial system, external financing is difficult to obtain. This affects some sectors more adversely than others. Rajan and Zingales (1998) and subsequent literature provide empirical evidence for this. The current paper builds upon those findings to study an open economy setting. I model a small economy that consists of an investment goods sector and a consumption goods sector. First, I consider how domestic financial institutions affect the cross-sectoral capital allocation in a closed economy, before looking at what happens when the economy integrates financially with the rest of the world. The effect of integration depends on the country's financial development relative to the rest of the world: If the economy is little financially developed, it benefits from cross-border capital flows. Foreign capital exploits the high returns in the under-financed sector, whereas unproductive domestic capital leaves the country. A negative effect of integration arises in a country with relatively high financial development. Overall, capital allocations become more similar across countries. I provide empirical evidence for the model predictions through various types of panel data analyses.

Sectors that produce a lower share of their output for consumption are more dependent on external finance, as Table 1.1 shows. So in a stylized two-sector world, the investment goods sector should suffer relatively more from domestic financial frictions like borrowing constraints. In consequence, we should observe a higher marginal product of capital in the investment goods sector relative to the consumption goods sector. The capital mis-allocation should also be reflected in the price level of consumption relative to investment, which is a measure easier to compute than marginal products.[3] The left panel of Figure 1.1 shows that relative prices are indeed positively correlated with a country's level of financial development. Financial integration is also positively correlated with relative prices (right panel), but the association could possibly be non-linear. At the same time, there is a strong negative correlation between the international dispersion of relative prices and the worldwide level of financial integration, as Figure 1.2 shows. So integration is associated with an international convergence in the allocation of capital between investment and consumption goods sector.

My model builds upon these stylized facts. The two-sector economy resembles that of Galor (1992), to which I add a credit market and a borrowing constraint, closely following Von Hagen and Zhang (2014). Borrowing plays an important role because only certain types of agents (called entrepreneurs) can invest directly in the investment goods sector, whereas additional capital needs to be acquired externally from the other agents (households). Solving the model for a closed economy, I show that the stronger the borrowing constraint, the less efficiently capital gets allocated: Too little capital is invested in the investment goods sector relative to the consumption goods sector. The mis-allocation is reflected by the price of consumption goods relative to investment goods, which is a negative function of the frictions.

[3]On the aggregate level, Lucas (1990) and Caselli and Feyrer (2007) have shown that wrongful calculations of the marginal product are the reason why many researchers have concluded that the volume of international capital flows observed in reality is too low compared to standard model predictions.

Table 1.1: External dependence and final consumption share

NAICS	Consumption Share 1997-2006	External Dependence	
		Rajan and Zingales (1998): 1980s	Cetorelli and Strahan (2006): 1990-2007
311	0.6215	-0.55	0.01
313	0.2603	-0.14	0.06
315	0.8849	-0.24	-0.14
321	0.0206	0.24	0.43
322	0.1103	-0.04	0.19
323	0.0369	-0.37	0.07
324	0.3308	-0.01	0.25
325	0.2943	-0.41	0.14
326	0.1112	-0.11	0.01
327	0.0701	-0.19	0.11
331	0.0073	0.09	0.27
332	0.0282	-0.43	0.07
333	0.0259	-0.33	-0.11
334	0.1157	-0.04	0.09
335	0.2203	-0.24	0.06
336	0.2861	-0.35	-0.10
337	0.3974	-0.32	0.19
339	0.4645	-0.18	-0.03
Corr. with Cons. Share		**-0.3481**	**-0.4347**

Note: The consumption share is private consumption expenditure over total commodity output minus net exports (U.S. Bureau of Economic Analysis). External dependence is defined as in Rajan and Zingales (1998), reported for NAICS industries by Haltenhof, Lee, and Stebunovs (2014). In accordance with Rajan and Zingales (1998), the United States is treated as the benchmark case where, due to high financial development, capital should be allocated efficiently.

I next assume that the small economy opens up financially towards the rest of the world, allowing for various degrees of integration. International capital flows will affect relative prices as long as the level of financial development of the country differs from the rest of the world. If the country is relatively less developed, foreign capital will flow into the investment goods sector whereas domestic savings will flow out of the consumption goods sector. This happens because the return to investing in the investment goods sector is higher at home, whereas the return in the consumption goods sector is higher abroad. In turn, the allocation of capital across the two sectors of the domestic economy becomes more efficient, which is reflected by an increase in the consumption-investment price ratio. If the economy is financially more developed than the rest of the world, the opposite will be observed: the favorable domestic financing conditions allow investments abroad, but simultaneously drain the domestic investment goods sector of capital. Additionally, foreign capital flows into the consumption goods sector. In consequence, the domestic capital allocation becomes less efficient, as evidenced by a decreasing relative price. International capital flows always make

Figure 1.1: Financial development, financial integration and relative prices

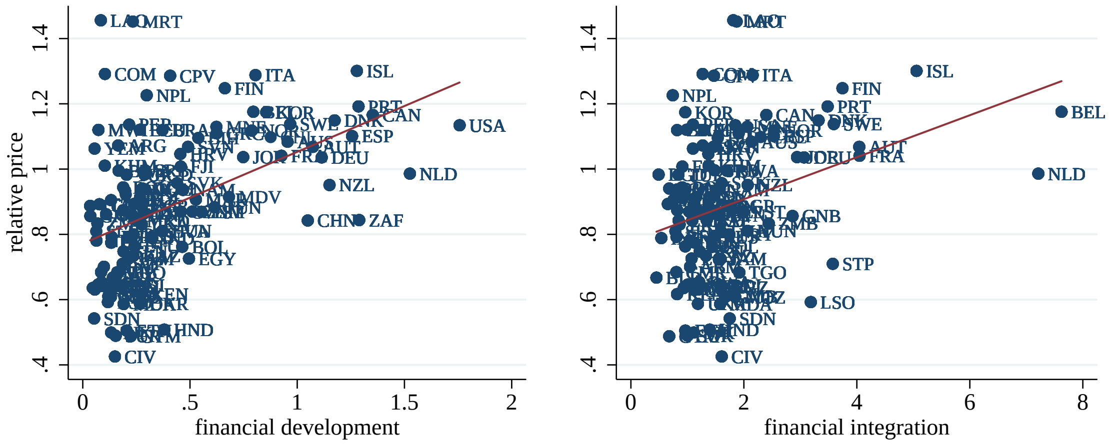

Note: Financial development is private credit by the domestic financial sector as fraction of GDP (Beck et al. (2000)). Financial integration is external assets and liabilities as fraction of GDP (Lane and Milesi-Ferretti (2007)). The relative price is the price level of household consumption divided by the price level of capital formation (*Penn World Table* (Feenstra et al. (2015))). All values are country means 1996-2010.

cross-sectoral capital allocations more similar across countries.

The model can be tested empirically, as price levels of consumption and investment are available for a large number of countries from the Penn World Tables. The sample contains annual data for 113 countries from 1996 to 2010. I apply two different types of panel data analyses: First, to test whether relative prices converge internationally upon integration, I regress deviations of a country's relative price from the respective cross-sectional mean on financial integration. The estimated coefficient on financial integration is negative and significant, which provides support for the convergence effect of integration predicted by the model.

Second, I test whether the effect of financial integration differs by the level of financial development of a country relative to the rest of the world. To this end, I sort country-time observations into a high development and a low development subsample. I find that in the high development group, financial integration has a significantly negative effect on the consumption-investment relative price, whereas the effect is mostly insignificant in the low development cluster. This means that the international convergence in capital allocations rather stems from a negative allocation effect in countries with a high level of financial development than from a positive effect in countries with low financial development. This has important policy implications. While highly financially developed countries may still benefit from financial integration in ways not captured in this paper, the findings may weaken the rationale for capital account opening in those countries.

The plan of the paper is as follows: I next place my work in the context of the existing literature. Section 3 features the model. I derive the steady state in a closed economy set-up before showing how the key economic variables react when the country opens its capital account. In section 3, I lay out the empirical strategy. I describe the estimation technique and data, before presenting the regression results and some robustness checks. Section 4 concludes.

Figure 1.2: Financial integration and relative price dispersion

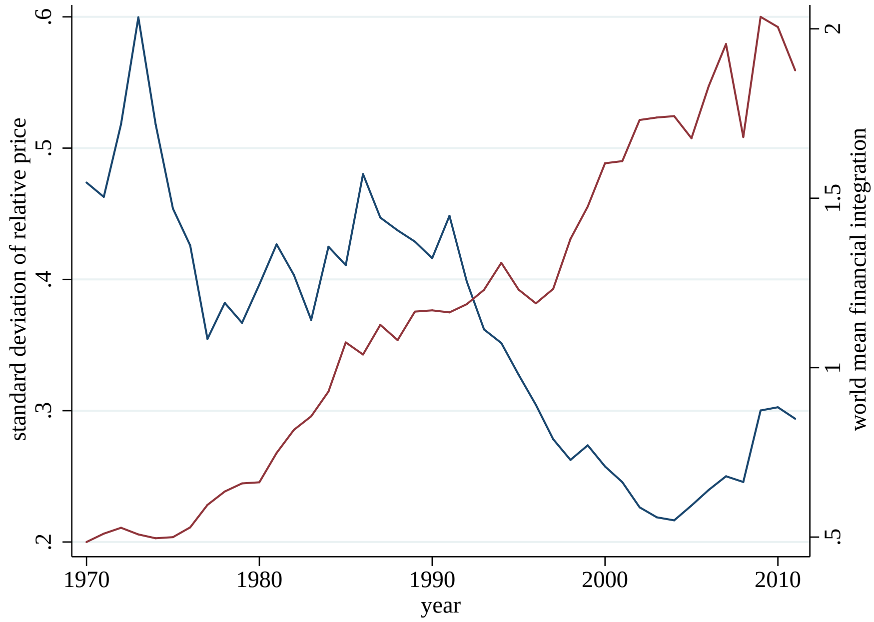

Note: The red line is the mean of all countries' financial integration level. The blue line is the standard deviation of the relative price across countries.

1.1 Related Literature

To my knowledge, I am the first to model the effect of financial integration on cross-sectoral capital allocations and on the relative price between consumption and investment goods. Papers in international finance have so far mostly focused on the aggregate capital stock of integrating economies. Ju and Wei (2010) and Von Hagen and Zhang (2014) show that foreign capital can increase an economy's overall investment efficiency. In both papers, financial integration creates the opportunity to circumvent inefficient domestic financial institutions. A different set-up, but with similar implications, is chosen by Caballero, Farhi, and Gourinchas (2008). They construct an equilibrium model of a world with two fully integrated regions and show how a deterioration in the quality of the financial system in one region will lead to capital outflows from that region. Mendoza, Quadrini, and Ríos-Rull (2009) explain capital flows and global imbalances by differences in country portfolios. In a world with risky and safe assets, if countries differ in the degree of market completeness, they will choose to hold a different mix of the two asset types. The resulting net foreign asset positions will, in general, not be zero. Some papers distinguish between sectors, but in a different way from my paper. Von Hagen and Zhang (2011) construct a model of a two-sector world where both sectors produce intermediate goods that are then combined to a final consumption good. Reis (2013) studies the capital allocation between productive units within the nontradables sector. Such models are less easily testable empirically across a large group of countries.

There exist several empirical studies that explore the link between financial integration and cross-sectoral capital allocations. Pang and Wu (2009), drawing on Wurgler (2000), aim to assess in an industry-level study whether financial integration redirects investment from sunset to sunrise industries. Bonfiglioli (2008) and Levchenko, Rancière, and Thoenig (2009) analyze the effects of financial integration on capital accumulation and productivity, but in

Figure 1.3: Relative prices, financial and economic development

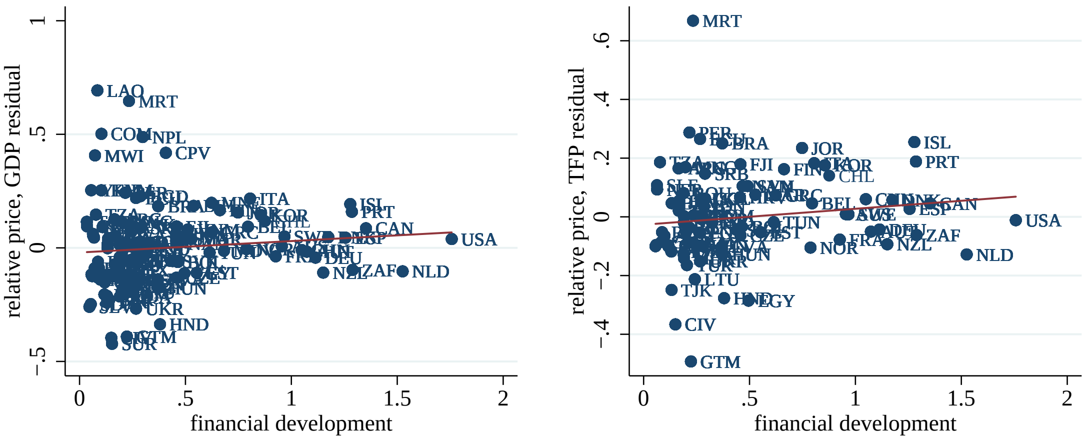

Note: The relative price residuals are residuals from a regression of the relative price on GDP per capita (left panel) or TFP (right panel). The coefficient in the left panel is 0.10, in the right panel 0.15, both significant at the 1% level.

contrast to my paper focus on the total amount of capital invested and not on its allocation. Galindo, Schiantarelli, and Weiss (2007) develop an index of investment efficiency at the firm level that compares hypothetical and real returns to investment, and use it to assess the role of financial integration. Benigno, Converse, and Fornaro (2015) analyze in an event study of the manufacturing sector whether episodes of large capital inflows are associated with changes in the relative investment in this sector. None of these papers links cross-sectoral capital allocations to international differences in relative prices.

Hsieh and Klenow (2007) analyze how dispersions in the price of investment relative to consumption can explain why investment rates differ across countries. They rely on the observation by Balassa (1964) and Samuelson (1964) that poor countries have a low productivity in tradable goods relative to nontradable goods. If investment goods are tradable and consumption goods (mostly) nontradable, poor countries will have a lower price of consumption whereas the price of investment will be roughly the same across countries. So while in my paper, the observed differences in relative prices are rooted in international differences in financial development, the explanation of Hsieh and Klenow (2007) builds on international differences in productivity. Both approaches should be seen as complementary. Figure 1.3 shows that the positive relation between relative prices and financial development persists (albeit weaker) when controlling for GDP or TFP.

2. Model

2.1 Capital Allocation in a Closed Economy

Firms

The economy consists of two production sectors, sector Y, which produces investment goods, and sector X, which produces non-durable consumption goods. Both sectors use capital and labor as inputs. Throughout the paper, I make the assumption that good Y is tradable and good X is non-tradable. This simplifying duality is meant to make the model more tractable. It can be motivated e.g. by assuming that good X has a high service content, whereas good Y is a manufactured good.[4]

Production functions exhibit constant returns to scale:

$$Y_t = z^y \left(K_t^y\right)^{\alpha} \left(L_t^y\right)^{1-\alpha} \quad X_t = z^x \left(K_t^x\right)^{\varepsilon} \left(L_t^x\right)^{1-\varepsilon} \tag{1.1}$$

where α, $\varepsilon \in (0,1)$. $K_t^x, K_t^y, L_t^x, L_t^y$ are sector capital and labor inputs with $K_t^x + K_t^y = K_t$ and $L_t^x + L_t^y = L_t$ the economy-wide capital and labor demand. z^y, z^x is technology, which is set to be larger than one because capital in this economy fully depreciates after one period. Investment takes one period to mature, so the investment good produced equals the new capital stock one period later, $Y_t = K_{t+1}$.

Input factors are rewarded by their respective marginal products, where r_t is the return to capital and w_t the return to labor. Let p_t be the price of one unit of good X, whereas the price of good Y is normalized to 1 (so p_t is at the same time the relative price of consumption to investment goods). Then,

$$r_t^y = \frac{\partial Y_t}{\partial K_t^y} = \alpha z^y \left(\frac{K_t^y}{L_t^y}\right)^{\alpha-1} \quad r_t^x = p_t \frac{\partial X_t}{\partial K_t^x} = p_t \varepsilon z^x \left(\frac{K_t^x}{L_t^x}\right)^{\varepsilon-1} \tag{1.2}$$

$$w_t^y = \frac{\partial Y_t}{\partial L_t^y} = (1-\alpha) z^y \left(\frac{K_t^y}{L_t^y}\right)^{\alpha} \quad w_t^x = p_t \frac{\partial X_t}{\partial L_t^x} = p_t (1-\varepsilon) z^x \left(\frac{K_t^x}{L_t^x}\right)^{\varepsilon} \tag{1.3}$$

Labour is assumed to be perfectly mobile between sectors, so $w_t^x = w_t^y = w_t$. In an economy without financial frictions, it also holds that $r_t^x = r_t^y$.

Agents

Agents live for two periods (young, old) and there is no population growth. Each generation consists of a mass of one. When young, the agents inelastically supply one unit of labor and earn a wage w_t (denoted in units of the investment good), of which they spend $p_t c_t$ on non-durable consumption goods and s_t on investment goods. They can use the investment goods in two different ways: either lend them to the credit market, where they earn an interest in the period after, or use them to operate a firm in the X- or Y-sector. Since that investment in both sectors takes one period to mature, firms will become operational only when the agents are old. In their second period of life, agents thus earn the returns on their investment and/or interest. They consume all of their income and exit the economy at the end of the period.

Each generation consists of two types of agents, entrepreneurs (of mass $\eta \in (0,1)$) and households (of mass $1-\eta$), which differ only in terms of their production opportunities: Entrepreneurs can produce both investment and consumption goods, whereas households can produce only consumption goods.[5]

Agents born in t maximize lifetime utility

$$U_t^j = ln\left(c_{t,young}^j\right) + \beta ln\left(c_{t+1,old}^j\right) \tag{1.4}$$

[4] A paper that chooses the same duality in a model set-up is Hsieh and Klenow (2007), who seek to explain cross-country differences in investment rates.

[5] These shares can be endogenized, as shown by Zhang (2014).

where $j = h, e$ denotes households and entrepreneurs. The first part denotes consumption when young and the second part consumption when old. β is a discount factor $\in (0, 1)$. In each of the two periods, the agents are subject to a budget constraint:

$$w_t = s_t^j + p_t c_{t,young}^j s_t^j r_{t+1}^j = p_{t+1} c_{t+1,old}^j \tag{1.5}$$

The maximization problem yields

$$s_t^j = \frac{\beta}{1+\beta} w_t c_{t,young}^j = \frac{1}{1+\beta} \frac{w_t}{p_t} c_{t+1,old}^j = \frac{r_{t+1}^j}{p_{t+1}} \frac{\beta}{1+\beta} w_t \tag{1.6}$$

External Capital Requirements and Financial Frictions

If all agents were able to produce in each of the two sectors, capital would be allocated optimally across sectors, implying $r_t^x = r_t^y = r_t$. This outcome would also be achieved if the fraction of entrepreneurs, η, were larger than the fraction of capital optimally allocated to the Y-sector. However, as shown in Appendix A, if $\eta < \frac{(1-\varepsilon)\alpha\beta}{\alpha\beta+\varepsilon}$, the entrepreneurs will want to borrow from the households until rates of return in both sectors equalize. In the following, I will assume this to be the case. In this way, the Y-sector is more financially dependent than the X-sector.

Assume now that there are financial frictions in the economy, which determine how much capital an entrepreneur can borrow. The entrepreneur, when young, will invest an amount of i_t^e in the Y-production, using the own savings into the investment good, s_t^e, and external capital n_t. The household will invest an amount i_t^h in the X-production and lend to the entrepreneur an amount of savings d_t:

$$i_t^e = s_t^e + n_t s_t^h = i_t^h + d_t \tag{1.7}$$

Denote by r_t^e the return to entrepreneurial capital and by r_t^h the return to lenders. Then, the total project return is

$$i_{t-1}^e r_t^y = s_{t-1}^e r_t^e + n_{t-1} r_t^h \tag{1.8}$$

The household earns on its savings $i_{t-1}^h r_t^x + d_{t-1} r_t^h$. No arbitrage implies $r_t^x = r_t^h$: the household gets the same returns whether it invests in an own project or lends to the entrepreneur. So the total return to the household can equally be written as $s_{t-1}^h r_t^x$.

Financial frictions impact the economy by imposing a limit on external capital. The entrepreneur can borrow only up to a certain fraction of total project returns,

$$r_t^h n_{t-1} = r_t^h (i_{t-1}^e - s_{t-1}^e) \leq \theta r_t^y i_{t-1}^e \tag{1.9}$$

where $\theta \in [0, 1]$ is a measure of financial development, which may refer to the development of the domestic financial sector as well as to other country-specific institutions like property rights protection, quality of corporate governance and corruption protection.[6] When the borrowing constraint is binding, θ limits the demand for credit (the supply of credit being fixed). In this way, it creates a wedge between the returns to entrepreneurs and lenders.

[6] In this formulation, I follow Von Hagen and Zhang (2011, 2014) and Matsuyama (2004, 2007). The borrowing constraint is kept as general as possible in order to not add unnecessary complexity to the model. However, Eq. (1.9) does not have to be a black box, but can be motivated by different kinds of moral hazard stories, for example following Bernanke and Gertler (1989), Kiyotaki and Moore (1997) or Holmstrom and Tirole (1997).

The return to an entrepreneur's capital is

$$r_t^e = \frac{r_t^y i_{t-1}^e - r_t^h n_{t-1}}{s_{t-1}^e} = r_t^y + \left(r_t^y - r_t^h\right)\left(\frac{i_{t-1}^e}{s_{t-1}^e} - 1\right) \tag{1.10}$$

where $\left(\frac{i_{t-1}^e}{s_{t-1}^e} - 1\right)$ is the amount borrowed per unit of investment. As Eq. (1.10) makes clear, the return on entrepreneurial capital consists of two parts: the marginal return on the project and an extra return earned on the external capital employed. The entrepreneur prefers to carry out the project instead of lending the savings to the credit market if $r_t^e \geq r_t^h$. This is the entrepreneur's participation constraint. If $r_t^e = r_t^h$, the borrowing constraint is slack and capital is allocated efficiently in the economy. Then it follows that $r_t^x = r_t^y$. If $r_t^e > r_t^h$, the borrowing constraint is binding and the entrepreneur will borrow up to the limit to exploit the leverage effect. In that case, the constraint distorts the allocation of capital by creating a wedge between the returns to external capital and entrepreneurial capital. This can be seen more clearly by combining Eqs. (1.9) and (1.10) to $r_t^e = \frac{r_t^y i_{t-1}^e (1-\theta)}{s_{t-1}^e}$. If the borrowing constraint is binding, r_t^e depends negatively on θ: returns to the entrepreneur are higher when frictions are larger.

Market Clearing Conditions

The labor market clears:

$$L_t = 1 \tag{1.11}$$

The investment good market clears:

$$(1-\eta)\, s_t^h + \eta s_t^e \equiv S_t = Y_t \tag{1.12}$$

The consumption good market clears:

$$(1-\eta)\left(c_{t,young}^h + c_{t,old}^h\right) + \eta\left(c_{t,young}^e + c_{t,old}^e\right) \equiv C_t = X_t \tag{1.13}$$

The credit market clears:

$$\begin{aligned}(1-\eta)\, d_t = \eta n_t \quad &\iff \quad (1-\eta)\left(s_t^h - i_t^h\right) = \eta\left(i_t^e - s_t^e\right) \\ &\iff \quad \underbrace{(1-\eta)\, i_t^h}_{K_{t+1}^x} + \underbrace{\eta i_t^e}_{K_{t+1}^y} \equiv I_t = S_t \end{aligned} \tag{1.14}$$

Autarky equilibrium

Given the outcomes to firms' profit maximization of Eqs. (1.2) and (1.3), and agents' utility maximization Eq. (1.6), the market clearing conditions Eq. (1.11) through Eq. (1.14), wage rate equalization, the definition of returns on entrepreneurial capital Eq. (1.10), the no-arbitrage condition and the borrowing constraint Eq. (1.9) (in combination with the participation constraint), a unique closed economy equilibrium can be derived for a given θ.

As shown in Appendix A, there exists a threshold value below which the borrowing constraint is binding (and consequently distorts capital allocation): $\overline{\theta} = 1 - \frac{\eta}{\psi} - \eta$ where

$\psi = \frac{1-\varepsilon}{\varepsilon}\frac{\alpha\beta}{1+\alpha\beta}$. For $\theta \in [0, \overline{\theta})$, the borrowing constraint is binding, for $\theta \in \left[\overline{\theta}, 1\right]$ it is slack. The borrowing constraint is stricter ($\overline{\theta}$ larger) if η is small (there are less entrepreneurs in the economy), if β is large (agents are patient and want to accumulate a large amount of savings) and if the Y-sector is more capital-intense relative to the X-sector.

In the case when the constraint is binding, the model solution is

$$r_t^{y-aut} = \alpha \frac{w_t^{aut}}{w_{t-1}^{aut}} \frac{1}{\frac{\psi\theta+\eta}{1+\psi\theta}} \qquad r_t^{x-aut} = \frac{\alpha}{\psi} \frac{w_t^{aut}}{w_{t-1}^{aut}} \frac{1}{1-\frac{\psi\theta+\eta}{1+\psi\theta}}$$

$$r_t^{e-aut} = \frac{w_t^{aut}}{w_{t-1}^{aut}} \left(\frac{\alpha\beta+\varepsilon}{(1-\varepsilon)\beta\eta} - \frac{\alpha(1-\eta)}{\psi\eta} \frac{1}{1-\frac{\psi\theta+\eta}{1+\psi\theta}} \right) \tag{1.15}$$

$$K_t^{y-aut} = \frac{\beta}{1+\beta} w_{t-1}^{aut} \frac{\psi\theta+\eta}{1+\psi\theta} K_t^{x-aut} = \frac{\beta}{1+\beta} w_{t-1}^{aut} \left(1 - \frac{\psi\theta+\eta}{1+\psi\theta}\right) \tag{1.16}$$

$$p_t^{aut} = \frac{1}{1-\varepsilon} \left(\frac{1+\alpha\beta}{\beta}\right)^{\varepsilon} \frac{1}{z^x} \frac{w_t^{aut}}{\left(w_{t-1}^{aut}\right)^{\varepsilon}} \left(1 - \frac{\psi\theta+\eta}{1+\psi\theta}\right)^{-\varepsilon} \tag{1.17}$$

$$w_t^{aut} = (w_{t-1}^{aut})^{\alpha} \left(\frac{\psi\theta+\eta}{1+\psi\theta}\right)^{\alpha} (1-\alpha)^{1-\alpha} z^y \tag{1.18}$$

The superscript *aut* stands for autarky equilibrium and is meant to distinguish this equilibrium from the integrated equilibrium described below. The model dynamics are characterized by Eq. (1.18), the dynamic equation of wages. Since $\alpha \in (0, 1)$, a unique and stable steady state exists with $w^{aut} = (1-\alpha)(z^y)^{\frac{1}{1-\alpha}} \left(\frac{\psi\theta+\eta}{1+\psi\theta}\right)^{\frac{\alpha}{1-\alpha}}$.

K_t^{y-aut} is a positive function of θ whereas K_t^{x-aut} depends negatively on θ: Frictions divert investment from the Y-sector to the X-sector. Consequently, r_t^{x-aut} is depressed and r_t^{y-aut} and r_t^{e-aut} are inflated relative to the frictionless case where $r_t^x = r_t^y = r_t^e$. Also, when η, the fraction of entrepreneurs in the economy, rises, r_t^{x-aut} rises and r_t^{y-aut} falls, i.e. the allocation becomes more efficient. The relative price p_t^{aut} is driven upwards by financial development θ: the consumption good is more expensive relative to the capital good, the higher the level of financial development of the country. Wages depend positively on θ as well. (The relationship becomes even stronger when wages are expressed in units of the consumption good, as $p_t w_t$).

Proposition 1: The model has a unique and stable steady state. In the case when the borrowing constraint is binding, $\theta \in [0, \overline{\theta})$, the wage rate is $w^{aut} = (1-\alpha)(z^y)^{\frac{1}{1-\alpha}} \left(\frac{\psi\theta+\eta}{1+\psi\theta}\right)^{\frac{\alpha}{1-\alpha}}$, which is smaller than in the non-binding case. A binding borrowing constraint creates a wedge between rates of return, $r_t^{e-aut} > r_t^{y-aut} > r_t^{x-aut}$. The wedge is larger, the less developed the country is. Relative prices are an increasing function of θ.

2.2 Financial Integration

Consider now a scenario where the country portrayed above integrates - partially or fully - with the rest of the world (RoW). The country is a small economy whereas RoW is large.

Other than in size, the two regions differ only in their level of financial development, so either $\theta > \theta^{RoW}$ or $\theta^{RoW} > \theta$. Assume that the borrowing constraint is binding everywhere.

In a financially integrated world, individuals are allowed to invest financial capital and entrepreneurial capital in both regions. This means that entrepreneurs can operate a firm either in their home country or abroad. Households, on the other hand, can lend their savings to financial intermediaries of both regions, or run an X-sector firm in the region of their choice. They continue to be excluded from operating a firm in the Y-sector.

Let Φ_t denote net outflows of financial capital, i.e. savings invested by domestic agents in a foreign financial market minus savings invested by foreign agents in the domestic economy. Conversely, Ω_t denotes net outflows of entrepreneurial capital, i.e. investments by domestic entrepreneurs in operating a firm abroad minus investments by foreign entrepreneurs in a domestic firm. Hence, positive values denote capital outflows, negative values capital inflows. Note that in terms of the international investment position, Φ_t are net debt assets and Ω_t are net assets in portfolio equity or foreign direct investments. The decision in which region to invest is taken by young agents. Compared to the closed economy model described above, the following equations change:
Credit market clearing now incorporates international capital flows, so Eq. (1.14) is replaced by

$$K_t^y + K_t^x = \eta s_{t-1}^e + (1-\eta)\, s_{t-1}^h - \Omega_{t-1} - \Phi_{t-1} \tag{1.19}$$

The left-hand side denotes investments made in the economy (i.e. the demand for savings, whether by residents or foreigners), whereas the right-hand side denotes the amount of savings which is available in the country (i.e. the supply of savings).

Assume that entrepreneurs borrow in the country where they produce.[7] Then, the borrowing constraint takes the form

$$r_t^x \left[(1-\eta)\, s_{t-1}^h - K_t^x - \Phi_{t-1}\right] \leq \theta r_t^y \left[\eta s_{t-1}^e - \Omega_{t-1} + (1-\eta)\, s_{t-1}^h - K_t^x - \Phi_{t-1}\right] \tag{1.20}$$

On the right-hand side, the term in square brackets denotes the total amount of capital that is employed in the country's Y-sector: entrepreneurial capital of local and foreign entrepreneurs ($\eta s_{t-1}^e - \Omega_{t-1}$) as well as financial capital of local and foreign households ($(1-\eta)\, s_{t-1}^h - K_t^x - \Phi_{t-1}$). The term in square brackets on the left-hand side denotes the external capital in the Y-production.

All other equilibrium conditions carry through. Eq. (1.12) remains unchanged because on the one hand, the demand for good Y is reduced by net capital outflows, but at the same time, the supply of Y is reduced by net exports of the same amount. Trade flows need to be exact mirror images of financial flows because in this model, the balance of payments must always be balanced. (The unit of account is the same for both K and Y.) Labor is assumed to be immobile across borders, but continues to be mobile nationally. Since good Y is tradable, prices equalize in that sector. Being normalized to 1 in both regions under autarky, there will be no visible effect. Note that in this way, the model abstracts from any effect of trade on factor prices. I control for trade integration and its possible effects on relative prices in the empirical analysis.

[7] Instead of the residence principle, the source principle could be imposed (entrepreneurs borrow in their country of origin). In that case, the effect on prices would be even stronger. See Appendix A for an argumentative proof.

Integrated Equilibrium

If the country fully integrates with RoW, unrestricted cross-border capital flows will equalize returns on entrepreneurial capital and financial capital across the two regions. Given that the country is small whereas RoW is large, the country's rates of return will converge to those of RoW, $r_t^x = r_t^{x-RoW}$ and $r_t^e = r_t^{e-RoW}$. If the country does not fully integrate, the interest rates will only partially converge to those of RoW. The degree of return rate equalization can thus be used as a measure of financial integration. It ultimately translates into equilibrium levels of capital flows.

The new equilibrium rates of return can be expressed as

$$r_t^{x-int} = \lambda r_t^{x-RoW} + (1-\lambda)\, r_t^{x-aut} r_t^{e-int} = \lambda r_t^{e-RoW} + (1-\lambda)\, r_t^{e-aut} \tag{1.21}$$

where $\lambda \in [0,1]$ is the degree of financial integration, or financial openness, of the country. Higher values signify more integration. The superscript *int* denotes (partial) financial integration values. If $\lambda = 1$, full integration is achieved. $\lambda = 0$ corresponds to complete financial autarky, in which case the equilibrium described below will correspond to that of the previous section.

In dependence of equilibrium rates of return, the equilibrium values of the other key variables can be expressed as

$$r_t^{y-int} = \left[\frac{\theta}{r_t^{x-int}} + \frac{(1-\theta)}{r_t^{e-int}}\right]^{-1} \tag{1.22}$$

$$K_t^{y-int} = \frac{\beta}{1+\beta}\alpha w_t^{int}\left[\frac{\theta}{r_t^{x-int}} + \frac{(1-\theta)}{r_t^{e-int}}\right] K_t^{x-int} = \frac{\beta}{1+\beta} w_t^{int}\frac{1}{r_t^{x-int}}\frac{\alpha}{\psi} \tag{1.23}$$

$$p_t^{int} = \frac{1}{\varepsilon z^x}\left(\frac{w_t^{int}}{1-\varepsilon}\right)^{1-\varepsilon}\left(r_t^{x-int}\right)^{\varepsilon} \tag{1.24}$$

$$w_t^{int} = (1-\alpha)\,(z^y)^{\frac{1}{1-\alpha}}\left[\alpha\left(\frac{\theta}{r_t^{x-int}} + \frac{(1-\theta)}{r_t^{e-int}}\right)\right]^{\frac{\alpha}{1-\alpha}} \tag{1.25}$$

$$\Phi_{t-1} = \frac{\beta}{1+\beta} w_{t-1}^{int}\left[(1-\eta) - \frac{w_t^{int}}{w_{t-1}^{int}}\frac{\alpha}{\psi}\frac{(1+\theta\psi)}{r_t^{x-int}}\right]$$

$$\Omega_{t-1} = \frac{\beta}{1+\beta} w_{t-1}^{int}\left[\eta - \frac{w_t^{int}}{w_{t-1}^{int}}\alpha\frac{(1-\theta)}{r_t^{e-int}}\right] \tag{1.26}$$

Lemma 2: In the model with international capital flows, there is a unique and stable steady state for given values of $\theta^{RoW}, \theta, \eta, \alpha, \beta, \varepsilon, z^x$ and z^y.

The steady state interest rates, prices and wages for a given integration level λ are

$$r^{x-int} = \frac{\alpha}{1-\eta}\frac{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}{\psi} r^{e-int} = \frac{\alpha}{\eta}\left(1-\theta-\lambda\left(\theta^{RoW}-\theta\right)\right) \tag{1.27}$$

$$p^{int} = \frac{1}{\varepsilon z^x}\left(\frac{w^{int}}{1-\varepsilon}\right)^{1-\varepsilon}\left(\frac{\alpha}{1-\eta}\frac{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}{\psi}\right)^{\varepsilon} \tag{1.28}$$

$$w^{int} = (1-\alpha)\left(z^y\right)^{\frac{1}{1-\alpha}}\left[\frac{(1-\eta)\,\psi\theta}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}+\frac{\eta(1-\theta)}{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}\right]^{\frac{\alpha}{1-\alpha}} \tag{1.29}$$

The cross-border steady state capital flows are

$$\Phi = \frac{\beta}{1+\beta}w^{int}(1-\eta)\left[1-\frac{1+\psi\theta}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}\right]$$

$$\Omega = \frac{\beta}{1+\beta}w^{int}\eta\left[1-\frac{1-\theta}{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}\right] \tag{1.30}$$

Eq. (1.30) creates a direct link between the abstract notion of financial integration used in the model and denoted by λ, and a real-world equivalent, the total volume of *de facto* capital flows $|\Phi| + |\Omega|$. As both $|\Phi|$ and $|\Omega|$ are strictly monotonically increasing functions of λ, $|\Phi| + |\Omega|$ must also be strictly monotonically increasing in λ. Conversely, λ is a positive function of $|\Phi| + |\Omega|$. This will be important in the empirical part of the paper.

Proposition 2: If the country is less developed than the rest of the world, $\theta < \theta^{RoW}$, the wedge between rates of return decreases in the level of financial integration λ. Then, the more the country is financially integrated, the higher are wages and the relative price. The relative price of the country approaches that of RoW from below: $p^{RoW} > p^{int} \geq p^{aut}$ (strict inequality if $\lambda > 0$). Whenever $\lambda > 0$, the country exports financial capital (savings) and imports entrepreneurial capital: $\Phi > 0, \Omega < 0$.
If the country is financially more developed than the rest of the world, $\theta > \theta^{RoW}$, the opposite holds: the return wedge increases in λ, wages and the relative price decrease in λ. The relative price approaches that of RoW from above: $p^{RoW} < p^{int} \leq p^{aut}$. Whenever $\lambda > 0$, the country exports entrepreneurial capital and imports financial capital: $\Phi < 0, \Omega > 0$.

Proposition 2 states the most important result of the model. It implies that financial integration will lead to a convergence in relative prices by increasing the price of consumption goods relative to investment goods in financially underdeveloped countries and depressing it in countries with high financial development.[8] Thus, the within-country capital allocation will become more efficient in countries with low financial development relative to the rest of the world, but less efficient in countries with high financial development. The predictions on the direction of capital flows are in line with what is often observed in the real world: Many emerging economies have positive net equity positions ($\Omega < 0$) and negative net debt positions ($\Phi < 0$) vis-à-vis the rest of the world, whereas the opposite holds for advanced economies (Lane and Milesi-Ferretti, 2007).[9]

[8]Of course, if all world regions are fully financially developed, financial integration will have no effect. In reality, we observe also capital flows between highly financially developed regions, but these should be governed by motives outside of this model, for example hedging purposes.

[9]The model makes no clear prediction on the direction of net capital flows. Depending on the parameter

The next part presents empirical evidence on relative prices and their link to financial integration.

3. Empirical Analysis

Proposition 2 can be translated into two testable hypotheses about the effect of financial integration on relative prices. The first hypothesis is that integration leads to a global convergence in relative prices. It is thus concerned with the *direction* of changes in relative prices. Countries with a relative price above that of the rest of the world will see the relative price decrease, whereas in countries with a lower relative price than the rest of the world, the relative price will increase. The second hypothesis describes the *channel* through which cross-border capital flows affect relative prices: the development of the domestic financial sector. According to the model, cross-border capital flows lower the relative price in countries with a high level of financial development and increase it in countries with a low level of financial development.

Testing the model hypotheses requires a global sample of countries. Data on the price-level of investment and consumption for a large number of countries are available on the aggregate level, as described in detail below, but not on the industry level[10]. This is why the empirical analysis is carried out exclusively on the country-level.

3.1 Estimation Technique

I test each hypothesis through a different type of regression set-up. Both build on fixed-effects panel data estimation. Fixed effects estimation is the most common way of estimating a panel model of the type described below, building on the assumption that country-specific effects are correlated with the explanatory variables of interest, and therefore need to be controlled for in the regression. In the present sample, the choice of fixed effects over random effects is motivated by a Hausman test.

A natural way of testing the first hypothesis is to consider directly the cross-country dispersion of relative prices and to check whether a higher level of financial integration decreases the distance of a country's relative price from the price prevailing in the rest of the world. As a proxy for the RoW relative price I use the GDP-weighted cross-sectional mean. The cross-sectional mean includes the country itself, but due to the large sample size, the error introduced in this way is negligible. Indeed, excluding China and the U.S. as countries with the highest GDP (where the bias would be largest) does not change the estimation results. The estimation equation is described by

$$|relprice_{it} - \overline{relprice_t}| = fi_{it}\beta_1 + fd_{it}\beta_2 + X'_{it}\beta_4 + u_{it} \tag{3.1}$$

with $$u_{it} = \alpha_i + \mu_t + \epsilon_{it}.$$

$i = 1, ..., N$ labels the cross-sectional and $t = 1, ..., T$ the time dimension of the panel. $relprice_{it}$ is the relative price, fi_{it} is financial integration, fd_{it} stands for the level of financial

values of θ^{RoW}, η and ψ, capital flows can be either uphill or downhill. The empirical evidence on this is also not clear. While there has been a large discussion about uphill net capital flows for developing countries (see e.g. Prasad et al., 2006, 2007), this might be due to public rather than private flows (Gourinchas and Jeanne, 2013, Alfaro, Kalemli-Ozcan, and Volosovych, 2014).

[10]For example, the EU KLEMS database is limited to a small sample of mostly highly developed countries.

development, and X_{it} is a vector that comprises other control variables. The errors u_{it} are composed of an unobserved country-specific effect α_i, an unobserved time-specific effect μ_t and an idiosyncratic component ϵ_{it}, which is assumed to follow a normal distribution with mean zero and variance σ_ϵ^2. Tests indicate that in the sample, there is serial correlation as well as cross-sectional dependence of the errors. I therefore use the Driscoll and Kraay (1998) estimator throughout all specifications, which corrects standard errors for both types of correlations. For the regressions to lend support to the model, β_1 should be negative. The model does not make any statement on the sign of the coefficients on the other variables.

I test the second hypothesis by splitting the sample in two subsets depending on the level of financial development. An easy way to do this would be to separate countries into those that have a level of financial development higher than the cross-sectional mean and those that have a lower level. However, this might not properly reflect the distribution of the observations. A more sophisticated technique is by *kmeans clustering*. This method aims at minimizing the Euclidean distance of (cross-sectional) observations from a previously specified number k of centers. Starting point are randomly or purposely chosen initial positions of the centers. The clustering algorithm assigns each observation to the closest center, and then recalculates the position of the center so that its distance from all observations within the cluster is minimized. This procedure is carried out iteratively until the assignment of observations to clusters does not change anymore. In a panel data set-up, kmeans clustering can be carried out separately for each year, which potentially results in a different country sample for each year.

In this paper, given the results from the theory part, it is natural to set the number of clusters to two. As initial sorting criterion, I use the cross-sectional mean. The algorithm sorts more countries into the low development group than into the high development group. Note that the cluster assignment of some countries changes throughout the sample period, with the result that the sum of the number of countries in each cluster is larger than the total number of countries in the sample.

For each cluster, I estimate the same regression equation, which is specified as

$$relprice_{it} = fi_{it}\beta_1 + fd_{it}\beta_2 + X_{it}'\beta_3 + u_{it} \tag{3.2}$$

where u_{it} is composed as before, and the same assumptions apply. According to the model, β_1 should be positive for the low development-cluster whereas it should be negative for the high development-cluster. β_2 should be positive always. The coefficient on GDP is expected to be positive as well.

Throughout most parts of the analysis, t refers to annual data. Prices should react immediately to changes in the supply and demand for capital, which makes a lagged response unlikely. Furthermore, I control for short-run fluctuations at the global level by including time fixed effects and for country-specific shocks to the domestic financial system through a banking crisis dummy described below. Nevertheless, it is interesting to additionally take a medium-run perspective on the relation between cross-border capital flows and relative prices. This can be seen as a robustness check for the annual regression results. In an extension, I therefore consider non-overlapping five-year averages.

3.2 Data, Measurement and Sources

The sample consists of annual data for 113 countries from 1996 to 2010. The choice of observations is determined first and foremost by the availability of data. In addition, several countries are excluded from the dataset: First, countries which are offshore financial centers, since financial flows into these countries might be governed by motives other than return maximization. As an indicator, I use a classification by the International Monetary Fund (IMF, 2000). Second, I drop countries where revenues from oil exports make up more than 50% of GDP. Investments in these countries are likely to be directed mostly at the oil-producing industry, which differs from other industries in that it produces mainly for export. I also experimented with dropping countries from the dataset that are no functioning market economies. One could argue that in these countries, prices for consumption and investment goods might not (or not completely) be determined by market forces, and therefore not reflect the purchasing power of the country accurately. Limiting the sample to WTO members does however not change the results, so these concerns do not seem to matter.

Appendix B provides an overview of the countries in the dataset. The dataset is summarized in Table 1.2.

Relative Prices

A standard source for data on purchasing power parities and price levels is the Penn World Table (PWT, version 8.1 used in this paper). It relies on the United Nations' International Comparison Program, which at regular intervals conducts surveys on local prices in a wide range of countries and product/service categories. These are aggregated into broad expenditure categories. For the non-benchmark years, PWT uses estimates, which are derived by interpolation. Price levels by expenditure category are calculated as the purchasing power parity divided by the market exchange rate. They are expressed relative to the price level of GDP of the United States in 2005.[11]

I use PWT price levels of household consumption and investment, and calculate the relative price, the dependent variable in the regressions, as the ratio of the two. The basket of investment goods includes machinery and equipment, construction and related products. The majority of these products should be tradable. In the category of household consumption fall not only consumption goods and services, but also housing, energy expenses, health and education expenditures. The assumption of non-tradability thus seems to hold fairly well for this category, with the exception of some consumption goods and energy. Including energy expenditures is also problematic because energy prices are largely driven by supply and demand fluctuations in international energy markets. However, the share of energy in consumption expenditures is not very high.

[11] Details on the methodology as well as a discussion of the problems connected to it can be found in Feenstra, Inklaar, and Timmer (2015).

Table 1.2: Summary statistics

Summary statistics for the dataset comprising of 113 countries, annual data 1996-2010

Variable	Description	Source	Unit of Measurement	Mean	Overall Std. Dev.	Between Std. Dev.	Within Std. Dev.	Min	Max
relprice	relative price	PWT 8.1	price level of household consumption over price level of capital formation	0.8881	0.2582	0.2194	0.1389	0.1357	2.0850
fi	*de facto* financial integration	Lane and Milesi-Ferretti (2007)	(Absolute value of) external assets plus external liabilites as fraction of GDP	1.6623	1.3254	1.1287	0.6948	0.2429	12.4728
fi_ci	*de jure* financial integration	Chinn and Ito (2008)	Capital account openness index (0 = completely closed, 1= completely open)	0.5131	0.3644	0.3412	0.1282	0	1
fd_credit	financial development (credit)	Beck et al. (2000)	fraction of GDP	0.4026	0.4148	0.3799	0.1577	0.0001	2.7281
fd_joint[1]	financial development (credit and bonds)	Beck et al. (2000)	fraction of GDP	0.4114	0.2398	0.2159	0.0998	0.0713	1.3560
gdp	log of GDP per capita	WDI	constant 2005 US$	7.7537	1.6117	1.5962	1.5867	4.8530	11.1244
open	trade openness	WDI	fraction of GDP	0.7566	0.3368	0.3224	0.1157	0.1493	2.2306
rule_law[2]	rule of law	WGI	composite index reflecting quality of contract enforcement, property rights etc., in units of std. normal distribution	-0.1192	0.9507	0.9354	0.1636	-2.07	2

[1]: 34 countries only: Argentina, Australia, Austria, Belgium, Brazil, Canada, Chile, China, Colombia, Czech Republic, Germany, Denmark, Spain, Finland, France, Greece, Hungary, Indonesia, India, Iceland, Italy, South Korea, Mexico, Netherlands, Norway, Peru, Poland, Portugal, Slovak Republic, Slovenia, Sweden, Turkey, United States and South Africa.

[2]: From 1996 to 2002, *rule_law* is available biannually only. Intermediate observations are obtained by linear interpolation.

Financial Integration

In quantifying financial integration, often the distinction is made between *de facto* and *de jure* integration (see e.g. Kose et al., 2010). *De jure* integration describes the degree to which capital account openness is allowed for or restricted by law. *De facto* integration refers to the actual links of a country to international capital markets. These two measures are not equivalent. A country can have a completely open capital account, but still experience no cross-border capital flows (like many African countries). On the other hand, if legal restrictions are not complied with or not enforced, capital flows may take place even in the presence of limited *de jure* integration. In general, for developed countries, both integration measures are highly correlated, but the correlation is less strong for developing countries (Kose et al., 2010: 4291). Depending on the type of research question, either one or the other measure may be more appropriate.

In the model, the index of financial integration λ refers to the extent to which rates of return equalize across countries. By Eq. (1.30), there is a direct link from λ to cross-border entrepreneurial capital and financial capital flows - which, in the simplified model world with full depreciation, are equal to the stock of international capital. A *de facto* measure of integration is therefore an appropriate representation. The dataset compiled by Lane and Milesi-Ferretti (2007) provides a detailed overview of external assets and liabilities for different asset classes. It has the advantage of an extensive coverage of countries at various stages of development. I use their measure of gross external assets and liabilities relative to GDP as my primary measure of financial integration.[12] Unfortunately, the dataset does not distinguish between private and public assets and liabilities. In particular, the external debt component of the indicator includes sovereign debt. The model, not containing a government sector, abstracts from this differentiation, but in practice, public capital flows are likely to be governed by motives other than return maximization. Ideally, I would therefore like to use a measure of private external positions only. Bonfiglioli (2008) suggests introducing an interaction term between *de facto* and *de jure* financial integration indices. Capital flows in and out of countries with a closed capital account, she argues, should reflect public transactions. But the distinction is not as clear, since private capital frequently finds loopholes for evading capital controls. I nevertheless use this interaction in an extension of my baseline regression model. As measure of *de jure* integration, I use the Chinn and Ito (2008) indicator. Based on the IMF's Annual Report on Exchange Arrangements and Exchange Restrictions (AREAER), this index is the first principal component of a principal component analysis on four binary indicators that codify restrictions on cross-border financial transactions.

De facto measures of financial integration have been criticized for not being exogenous in growth regressions, making it difficult to pin down causal effects (see Collins, 2007). Even though reverse causality might be present in the regressions presented below, the bias would go in the opposite direction from the effect that I will try to identify. If a country reacts to capital misallocation by opening its capital account or by enacting measures to promote international investments, then misallocation will lead to financial integration. In contrast, my aim is to show that financial integration will make allocations more efficient, not less efficient. Therefore, reverse causality introduces a downward bias to the coefficient on financial integration. The regression results therefore represent a lower bound.

[12] All regression results are robust to using only portfolio debt, portfolio equity, FDI and the other investment category of the International Investment Position.

Financial Development

Another key variable in my model is the level of development of the domestic financial system. It is a summary measure of the size and structure of the financial sector, and the quality of its institutions. A proxy often used in the literature is the ratio of private-sector credit to GDP.[13] But credit is only one type of financing means available to economic actors. In general, countries differ in the extent to which firms rely on banks and bank-like institutions to finance investment. Among developed countries, continental European countries have historically been characterized by a bank-based system whereas Anglo-Saxon countries have market-based systems (see Allen and Gale, 2000 and the references therein). Beck, Demirgüç-Kunt, and Levine (2000) have assembled more refined data on financial sector development, which, next to banks, cover other financial institutions like corporate bond markets, insurance companies, or pension funds. I use two series of this dataset as proxies for financial development in order to capture the duality of bank- vs. market-based systems. The bank-based measure is private credit by deposit money banks and other financial institutions as fraction of GDP. The market-based measure is private bond market capitalization as fraction of GDP.[14] Bond market capitalization refers to the outstanding domestic debt securities issued by private domestic entities. Unfortunately, bond market data have limited scope and are available for only 34 (developed) countries of my sample. This is why, in the benchmark specification, I rely only on credit market development. I take bond markets development into account in an extension presented in Appendix B, where I combine credit and bond market development into a joint measure. It is calculated as the natural logarithm of one plus the average sum of both measures. Note that credit and bond market development are actually highly correlated with a correlation coefficient of 0.70. This means that, at least for advanced economies, financing investment through banks or corporate bonds markets seem to be complements rather than substitutes.

Additional Controls

While in the model, countries differ only in terms of financial development, in reality there are other distinguishing features that need to be controlled for.

It is a well documented observation that the price of consumption goods relative to investment goods is lower in poor countries than in rich countries (see Parente and Prescott, 2000, Hsieh and Klenow, 2007 and the references therein). This is due to differences in TFP or human capital, which might not only manifest themselves in relative price differences as explained above, but could also affect relative prices in other ways. At the same time, rich countries are financially more integrated. I believe it possible that the level of economic development of a country has an independent influence on relative prices, which could erroneously be captured by the financial integration variable. Also, GDP is the denominator of the financial development as well as the *de facto* integration variable, which automatically creates a correlation. Therefore, I include per capita GDP.

[13] For its use in the finance and growth literature, see Levine (2005).

[14] In the literature, sometimes the size of the stock market is taken into account instead. However, financing firm activities through the emission of shares is risky, and small and medium-sized enterprises usually do not rely on the stock market for financing their investments. Therefore, stock market development does not cover the financing schemes preferred by a large share of the economy.

The model makes the simplifying assumption that only investment goods are tradable, whereas consumption goods are not. Trade does not serve any other purpose than to balance the current account. In reality, trade integration might, similarly to financial integration, reduce relative price differences. Classical trade theory - going back to the Heckscher-Ohlin model - argues that trade and financial integration are substitutes: Trade in goods automatically implies factor price equalization as long as the number of factors is not larger than the number of goods. This is why I include trade as a fraction of GDP (World Development Indicators) in the regression.

Ju and Wei (2010) argue that overall institutional quality has a different effect on capital allocation than financial development: Low institutional quality might distort investments independently of the sector concerned, depressing the overall level of investment in a country. I therefore include institutional quality as an additional control variable. As a proxy, I use the rule of law-index of the World Bank's Worldwide Governance Indicators. It is a composite index reflecting the quality of contract enforcement, property rights, police and the court system. The Worldwide Governance Indicators are based on survey data of enterprises, citizens and experts. The rule of law-index is reported in units of a standard normal distribution, with mean zero and standard deviation of one. It runs from approximately -2.5 to 2.5, with higher values corresponding to better governance. Data is available from 1996, on a yearly basis only from 2002. This variable thus puts the strongest limitations on the time dimension of my dataset.

Unfavorable macroeconomic conditions, weak institutions and unsustainable policies may trigger banking crises. During a banking crisis, the domestic financial sector experiences defaults and solvency problems of systemically important banks, which leads to a sharp increase in non-performing loans, depressed asset prices and hikes in interest rates. Banking crises will distort the borrowing and lending conditions in a country, and sectors dependent on external finance presumingly will be more strongly affected. The resulting capital misallocation should attract foreign capital to the capital-intensive sector, leading to relative price increases. At the same time, banking crises are periods of high uncertainty and are often accompanied by other types of crises like sovereign debt crises or currency crises (“twin crises"), which will deter foreign investors or change their investment behavior. The effect of a banking crisis on the overall level of financial integration is therefore not clear. In order to account for it, I include a dummy for year-country observations with banking crises in the regressions. Additionally, I introduce an interaction term of the crisis dummy with the measure of financial integration, which is meant to capture the differential effect that integration has during crisis episodes. Laeven and Valencia (2008, 2012) have compiled a global dataset in which they report the occurrence of systemic banking crises and their effects starting from 1970. Between 1996 and 2010, they identify 45 crisis episodes. Almost half of them occurred during the global financial crisis that started in 2007. While these affected mostly developed countries, the majority of banking crises prior to 2007 happened in developing countries.

Currency crises might affect cross-border capital flows in a more direct way. Sharp nominal depreciations of the currency, as observed particularly in developing countries, often lead to capital flow slowdowns and sudden stops. While they certainly have an effect on the size and composition of capital flows, it is not clear whether they also impact on the cross-sectoral allocation of investments. Prices in both sectors should react almost instantaneously to changes in the supply of capital, and therefore accurately reflect the capital allocation in a country. I experimented with including a currency crises dummy as a control variable in the regressions (reported also by Laeven and Valencia, 2008), but it never had any significant effect.

3.3 Regression Results

Convergence Regressions

Table 1.3 presents the baseline regressions for the analysis of cross-country relative price convergence. The measure of financial development used is credit market development. For each regression specification, I report estimations with only country fixed effects, and with country and time fixed effects. Columns (1) and (2) present a simplified regression-setup without any controls beyond financial integration and financial development. In Columns (3) and (4), the set-up is extended to control for other explanatory variables. Columns (5) and (6) additionally contain the banking crises dummy and an interaction term of financial integration with banking crises. The last columns (7) and (8) instead include *de jure* financial integration and an interaction with *de facto* integration.[15]

The coefficient on financial integration is negative and statistically significant throughout all specifications. There is thus strong evidence for a convergence effect. It is smaller when time fixed effects are used, which indicates that in the specifications with only country fixed effects, *fi* might partly pick up other time-specific influences. The coefficient in column (4) can be interpreted such that a one-unit increase in financial integration moves the relative price of a country closer to the cross-sectional mean by -0.0113. A one-unit increase in financial integration corresponds for example to the difference between the average 1996 to 2010 level of financial integration of Peru to that of New Zealand, or to the increase in financial integration experienced by Austria between 2000 and 2004. The deviation from the cross-sectional mean is on average 0.23, so that around 5 percent of the gap is closed when *fi* rises by 1.

When testing for nonlinearity of the effect, a quadratic financial integration term is not significant. Integration does not seem to have stronger effect for countries with particularly high or low levels of integration. An interaction term between financial integration and credit market development equally does not produce a coefficient significantly different from zero.

Credit market development has a positive, mostly significant effect on relative price dispersions. The coefficient of GDP is positive and significant in the specifications with both types of fixed effects, so in richer countries, the relative price diverges more from the weighted cross-sectional mean. As mentioned above, the model does not offer any interpretation of these results. The coefficient on trade openness is negative whenever significant: Trade integration, similarly to financial integration, seems to lead to convergence of relative prices. But since the effect becomes insignificant in specifications with time fixed effects, it might simply reflect some time trend of relative prices.

The coefficient on the banking crises dummy in column (5) is positive and significant. Banking crises seem to exacerbate the distortive effect of financial frictions on domestic capital allocation at the mean level of financial integration. When including time fixed effects in column (6), the coefficient becomes insignificant, however the interaction term between

[15]The results are similar if both the banking crisis with interaction, and *de jure* integration with interaction are included.

Table 1.3: Convergence regressions with credit market development

dependent variable: relprice deviations from cross-country mean

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES								
fi	-0.0175***	-0.0131**	-0.0142***	-0.0113**	-0.0153***	-0.0093*	-0.0147*	-0.0165*
	(0.0034)	(0.0046)	(0.0027)	(0.0044)	(0.0028)	(0.0050)	(0.0076)	(0.0078)
fd_credit	0.0508	0.0442**	0.0487**	0.0421**	0.0386*	0.0381**	0.0512***	0.0388***
	(0.0328)	(0.0167)	(0.0188)	(0.0166)	(0.0182)	(0.0160)	(0.0137)	(0.0118)
open			-0.0758***	-0.0135	-0.0741***	-0.0131	-0.0758***	-0.0127
			(0.0150)	(0.0131)	(0.0142)	(0.0131)	(0.0166)	(0.0134)
gdp			-0.0016	0.0438**	0.0030	0.0425**	0.0029	0.0455*
			(0.0297)	(0.0188)	(0.0295)	(0.0184)	(0.0361)	(0.0242)
rule_law			-0.0207	-0.0191	-0.0165	-0.0183	-0.0189	-0.0162
			(0.0156)	(0.0155)	(0.0172)	(0.0155)	(0.0162)	(0.0154)
bcrisis					0.0247**	0.0123		
					(0.0087)	(0.0096)		
bcrisis*fi					-0.0004	-0.0044**		
					(0.0025)	(0.0016)		
fi_ci*fi							0.0000	0.0097
							(0.0101)	(0.0089)
fi_ci							-0.0217	-0.0024
							(0.0346)	(0.0318)
Constant	0.2406***	0.2681***	0.2997	0.0000	0.2669	0.0000	0.2771	-0.0890
	(0.0169)	(0.0062)	(0.2340)	(0.0000)	(0.2316)	(0.0000)	(0.2652)	(0.1734)
Observations	1,618	1,618	1,496	1,496	1,496	1,496	1,476	1,476
No. of groups	113	113	113	113	113	113	111	111
R-squared	0.0099	0.0637	0.0148	0.0627	0.0173	0.0631	0.0655	0.0155
Country FE	YES	YES	YES	YES	YES	YES	YES	Yes
Year FE	NO	YES	NO	YES	NO	YES	NO	YES

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

banking crises and financial integration is negative and significant. This provides evidence that when a country is going through a banking crisis, financial integration becomes more effective in promoting the convergence of the relative price. In columns (7) and (8), I aim to determine what type of capital flows lead to convergence in relative prices. While the coefficient on *de facto* integration stays negative and significant, the interaction term with *de jure* integration is insignificant as well as the coefficient on *de jure* integration. The effect of *fi* does not seem to depend on capital account openness, which could mean that private capital flows rather than public capital flows drive the convergence in relative prices. As these might be governed to a larger extent by profit maximization, this further supports the model.

As shown in Appendix B, the results for the coefficient on financial integration are robust to using the joint measure of bank and bond market development instead of the simple credit measure as a proxy for financial development. The coefficient becomes even more significant and slightly larger, which might however be due to the different sample size.

Cluster Regressions

Next, I investigate the second hypothesis, according to which the level of financial development determines the direction of price changes initiated by international capital flows. The regressions above considered squared relative price deviations, so nothing could be said on this account. Table 1.4 shows separate regressions for countries with high and low levels of credit market development. Columns (1) through (4) display the results for the low development cluster, columns (5) through (8) for the high development cluster. In order to check the validity of the model, compare the coefficients of (1) with (5), (2) with (6), (3) with (7), and (4) with (8).

Table 1.4: Cluster regressions

dependent variable: relprice

	Low Development				High Development			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
VARIABLES								
fi	-0.0296	-0.0097	-0.0249**	-0.0892***	-0.0224***	-0.0239***	-0.0305***	-0.0131
	(0.0170)	(0.0097)	(0.0094)	(0.0177)	(0.0050)	(0.0051)	(0.0087)	(0.0077)
fd_credit	0.1455***	0.2537***	0.2504***	0.2202**	0.0342	0.0565	0.0580*	0.0418*
	(0.0419)	(0.0791)	(0.0829)	(0.0746)	(0.0327)	(0.0345)	(0.0302)	(0.0213)
open		0.0844**	0.0858**	0.0886*		0.2419***	0.2227***	0.2435***
		(0.0337)	(0.0333)	(0.0429)		(0.0668)	(0.0707)	(0.0661)
gdp		-0.3155***	-0.3271***	-0.3038***		0.1732**	0.2254**	0.1668
		(0.0305)	(0.0349)	(0.0393)		(0.0765)	(0.0909)	(0.1104)
rule_law		0.1764***	0.1883***	0.1614***		0.1136**	0.1234**	0.0744
		(0.0145)	(0.0188)	(0.0144)		(0.0527)	(0.0535)	(0.0616)
bcrisis			-0.0419*				0.0167	
			(0.0199)				(0.0388)	
bcrisis*fi			0.0521***				0.0100	
			(0.0136)				(0.0095)	
fi_ci				-0.0495				0.3248*
				(0.0408)				(0.1573)
fi_ci*fi				0.1417***				0.0014
				(0.0229)				(0.0235)
Constant	0.8217***	3.0964***	3.2032***	0.0000	1.1658***	-0.8980	-1.4214	-1.0868
	(0.0208)	(0.1848)	(0.2168)	(0.0000)	(0.0542)	(0.7893)	(0.9329)	(0.9709)
Observations	1,261	1,160	1,160	1,142	357	336	336	334
No. of groups	94	94	94	92	36	36	36	35
R-squared	0.0202	0.1025	0.1095	0.1290	0.1890	0.2420	0.2615	0.2871
Country FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	YES	YES	YES	YES	YES	YES	YES	YES

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

First, note that in the regressions that only estimate a level effect of financial integration, the coefficient on *fi* is insignificant for the low development cluster (columns (1) and (2)) and significantly negative for the high financial development cluster (columns (5) and (6)). This can also be seen in Figure (1.4), which plots the residuals of a fixed effects regression of the

Figure 1.4: Relative price residuals and financial integration

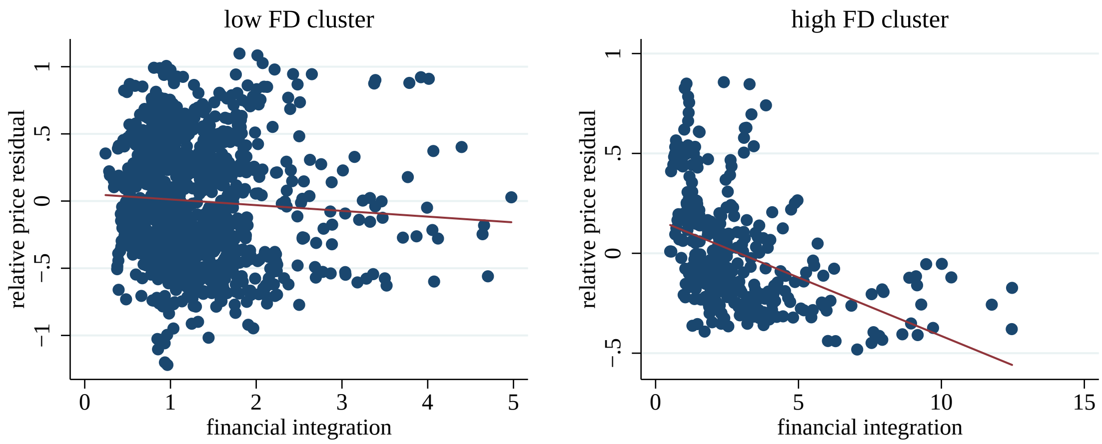

relative price on fd, *open*, *gdp* and *rule_law* against financial integration (corresponding to columns (2) and (6) in Table 1.4). Whereas in the high financial development cluster, the correlation between relative price residuals and financial integration is clearly negative, the dots in the low financial development cluster are very dispersed. Low development countries seem to be a very heterogeneous group, with some countries behaving as predicted by the model, but others not.[16] But overall, integration seems to drive relative prices of high-price countries downwards rather than driving those of low-price countries upwards. This is a potentially problematic result: even though it was shown above that capital allocations become more similar across countries, this apparently works through worsening the allocations in high development countries rather than by improving them in low development countries. This finding is in line with Reis (2013), who finds negative effects of financial integration on productivity in the nontradable goods sector of Portugal, a country that is part of the highly developed cluster. Benigno et al. (2015) find that in periods of large capital inflows, international capital are associated with negative effects on within-country capital allocation in a sample of advanced and some emerging economies.

In columns (3) and (7), differential effect during crisis and non-crisis times is taken into account. In normal times, financial integration affects relative prices negatively in both clusters, but more so in the high development cluster. During banking crises, however, the effect turns positive in the low development cluster, to a total of 0.0272. International capital then seems to be particularly directed to the externally dependent capital goods producing sector, and/or domestic capital is flowing out of the less dependent consumption goods producing sector. This provides additional evidence that financial integration is beneficial when domestic lending channels are interrupted - however only for countries with a low level of financial development.

In columns (4) and (8), the additional interaction term with *de jure* integration is highly significant and positive for the low development cluster, whereas it is close to zero and insignificant for the high development cluster. This may be interpreted in the sense that in countries with low financial development, public capital flows have a larger effect on in-

[16]For example, the correlation turns positive if we would exclude Argentina, Brazil, Sao Tomé and Príncipe and Lesotho.

creasing the efficiency of the capital allocation than private capital flows. However, in low development countries, it might be easier for private capital to circumvent official capital controls than in high development countries, which is why this result is to be interpreted with caution.

The coefficient on financial development is positive and significant throughout all specifications. This is in line with the model predictions. The coefficient is smaller in the high development cluster, which means that the effect is nonlinear. Trade openness has a positive effect on relative prices, and so does rule of law, although the coefficient is insignificant in the high development cluster. The effect of GDP is negative in the low development cluster, but positive in the high development cluster. Whereas GDP was shown above to induce a convergence of relative prices, this is not true across groups with different levels of financial development.

Medium-Run Effects

Table 1.5: Convergence regressions for five-year averages

dependent variable: relprice deviations from cross-country mean

	(1)	(2)	(3)	(4)
VARIABLES				
fi	-0.0181	-0.0175*	-0.0272	-0.0245
	(0.0062)	(0.0056)	(0.0123)	(0.0177)
fd_credit	0.0289	0.0469	0.0384	0.0413**
	(0.0416)	(0.0257)	(0.0140)	(0.0087)
open		-0.0650**	-0.0655**	-0.0596
		(0.0095)	(0.0126)	(0.0215)
gdp		-0.0256	-0.0177	-0.0328
		(0.0367)	(0.0427)	(0.0520)
rule_law		0.0277	0.0348	0.0278
		(0.0250)	(0.0197)	(0.0308)
banking_crisis			0.0009	
			(0.0043)	
bcrisis_fi			0.0085	
			(0.0054)	
fi_ci				-0.0112
				(0.0757)
fi_ci_fi				0.0153
				(0.0267)
Constant	0.2455***	0.4844	0.4404	0.5397
	(0.0123)	(0.2828)	(0.3119)	(0.3579)
Observations	287	281	281	275
No. of groups	108	106	106	104
R-squared	0.0218	0.0205	0.0259	0.0218
Country FE	YES	YES	YES	YES
Year FE	NO	NO	NO	NO

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The cross-sectional convergence and development cluster regressions are replicated for non-overlapping five-year averages. The panel is thereby reduced to three time periods: 1996-2000, 2001-2005 and 2006-2010. Table 1.5 replicates Table 1.3 for five-year averages. The results mostly confirm those of the regressions with annual data. Lower significance levels can be explained by a much smaller sample size, whereas the point estimates are of similar magnitude. Thus, the effects of financial integration on the allocation of capital hold also in the medium-run.

Table 1.6 conveys largely the same message as Table 1.4: Focusing on the level effect, integration leads to a lower relative price in countries with high financial development (column 2), but has no significant effect for countries with low financial development (column 5). The interaction of *de facto* with *de jure* financial integration again is shown to be significant only for the low development cluster, where public capital flows may have a more positive effect on the cross-sectoral efficiency of capital allocation than private capital flows.

Table 1.6: Cluster regressions for five-year averages

dependent variable: relprice

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
fi	-0.0422	-0.0501	-0.2381**	-0.0142	-0.0260***	-0.0112
	(0.0193)	(0.0248)	(0.0484)	(0.0053)	(0.0019)	(0.0093)
fd_credit	0.2153	0.3367*	0.3204**	0.0975	0.0552	0.0054
	(0.0792)	(0.0840)	(0.0504)	(0.0378)	(0.0275)	(0.0108)
open		0.2191**	0.2205**		0.3225**	0.0387
		(0.0500)	(0.0370)		(0.0611)	(0.1725)
gdp		-0.1695**	-0.2038**		0.2181**	0.2131*
		(0.0330)	(0.0275)		(0.0252)	(0.0667)
rule_law		0.1633***	0.1293***		0.0441	0.0105
		(0.0057)	(0.0063)		(0.1369)	(0.1471)
fi_ci			-0.1937			0.4875*
			(0.0880)			(0.1202)
fi_ci*fi			0.3570**			0.0223
			(0.0746)			(0.0315)
Constant	0.8374***	1.9533***	2.3054***	1.0149***	-1.3349**	-1.4676*
	(0.0368)	(0.1542)	(0.1347)	(0.0171)	(0.3034)	(0.4157)
Observations	226	220	214	61	61	61
No. of groups	88	86	84	29	29	29
R-squared	0.0231	0.1123	0.2429	0.0400	0.2784	0.3915
Country FE	YES	YES	YES	YES	YES	YES
Year FE	NO	NO	NO	NO	NO	NO

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

4. Conclusion

This paper presents evidence on how financial integration influences the cross-sectoral allocation of capital in integrating economies. In the intorduced two-sector model of a small economy at different stages of financial integration, the level of financial development is a key variable. Cross-country capital flows improve the allocation of capital in financially underdeveloped countries, whereas they have the opposite effect in developed countries. Financial integration contributes to making the cross-sectoral capital allocations of integrating economies more similar.

The second part of the paper tests the model predictions empirically. The price of consumption relative to capital goods serves as a proxy for the cross-sectoral efficiency of capital allocation in an economy. A panel data analysis of 113 countries over 15 years finds strong evidence that financial integration leads to a more similar allocation of capital across countries. It also supports the model prediction that the level of financial development is important for determining whether the cross-sectoral allocation of capital improves upon integration. However, this works through a negative allocation effect for countries with high financial development rather than a positive effect for countries with low financial development. In the future, I plan to look deeper into this result. Firm-level studies can provide additional insights on the relation between financial frictions and capital allocation (for a recent study, see Te Kaat, 2016).

In focusing on capital allocation efficiency, this paper has looked into one channel through which financial integration affects the real economy. The motives for and effects of financial integration are of course more widespread. For example, this analysis has abstracted from income risks. In reality, portfolio diversification can serve as an insurance against idiosyncratic income shocks. Risk hedging is an important motivation particularly for North-North capital flows and should bring additional benefits for the integrating economies. Thus, while the paper has called into question efficiency gains in cross-sectoral capital allocations for countries with high financial development, these countries may still enjoy gains from integration for other reasons.

Appendices

A. Proofs and Further Analytical Results

A.1 Non-Binding Solution and Threshold for Borrowing Constraint

The equilibrium values for the case when the borrowing constraint is slack are

$$K_t^{y-\overline{aut}} = \frac{\beta}{1+\beta} w_{t-1}^{\overline{aut}} \frac{\psi}{1+\psi} K_t^{x-\overline{aut}} = \frac{\beta}{1+\beta} w_{t-1}^{\overline{aut}} \left(1 - \frac{\psi}{1+\psi}\right) \tag{A.1}$$

$$r_t^{y-\overline{aut}} = \alpha \frac{w_t^{\overline{aut}}}{w_{t-1}^{\overline{aut}}} \frac{1}{\frac{\psi}{1+\psi}} \qquad r_t^{x-\overline{aut}} = \left(r_t^{y-\overline{aut}} =\right) \frac{\alpha}{\psi} \frac{w_t^{\overline{aut}}}{w_{t-1}^{\overline{aut}}} \frac{1}{1-\frac{\psi}{1+\psi}}$$
$$r_t^{e-\overline{aut}} = \frac{w_t^{\overline{aut}}}{w_{t-1}^{\overline{aut}}} \left(\frac{\alpha\beta + \varepsilon}{(1-\varepsilon)\beta\eta} - \frac{\alpha(1-\eta)}{\psi\eta} \frac{1}{1-\frac{\psi}{1+\psi}} \right) \tag{A.2}$$

$$p_t^{\overline{aut}} = \frac{1}{1-\varepsilon} \left(\frac{1+\alpha\beta}{\beta}\right)^{\varepsilon} \frac{1}{z^x} \frac{w_t^{\overline{aut}}}{(w_{t-1}^{\overline{aut}})^{\varepsilon}} \left(1 - \frac{\psi}{1+\psi}\right)^{-\varepsilon} \tag{A.3}$$

$$w_t^{\overline{aut}} = \left(w_{t-1}^{\overline{aut}}\right)^{\alpha} \left(\frac{\psi}{1+\psi}\right)^{\alpha} (1-\alpha)^{1-\alpha} z^y \tag{A.4}$$

When the borrowing constraint is weakly binding, this solution is equivalent to the solution for the binding constraint as laid out in the main text. Thus,

$$\frac{\psi}{1+\psi} = \frac{\psi\theta + \eta}{1 + \psi\theta} \Longleftrightarrow \theta = 1 - \frac{\eta}{\psi} - \eta = 1 - \frac{1+\psi}{\psi}\eta \equiv \overline{\theta} \tag{A.5}$$

If $\theta < \overline{\theta}$, the borrowing constraint is binding. In that case, return rate equalization does not take place and p_t will be too low compared to the frictionless situation. In order to ensure that $\overline{\theta} < 1$, it needs to hold that $\eta < \frac{\psi}{1+\psi} = \frac{(1-\varepsilon)\alpha\beta}{\alpha\beta+\varepsilon}$. This is also the requirement stated in the main text for that entrepreneurs want to use external capital in the production.

A.2 Proof of Proposition 1

I prove that (1.15) - (1.18) are the model solutions in the case where the borrowing constraint is binding.

Use (1.6) to replace utility maximising consumption and saving in all equations. Derive optimal labor allocation across sectors: (1.12) and (1.3) yield $L_t^y = \frac{\beta(1-\alpha)}{1+\beta}$ and with (1.11), $L_t^x = \frac{1+\beta\alpha}{1+\beta}$. (1.14) tells how capital is distributed between X- and Y-producing firms, $K_t^x = \frac{\beta}{1+\beta} w_{t-1} - K_t^y$, and is used to replace K_t^x. Then, prices and interest rates can be expressed as function only of K_t^y and w_t. Let $\varphi_t \equiv \frac{1+\beta}{\beta} K_t^y$. Then,

$$r_t^y = \alpha w_t \frac{1}{\varphi_t} r_t^x = \frac{\varepsilon(1+\beta\alpha)}{(1-\varepsilon)\beta} w_t (w_{t-1} - \varphi_t)^{-1} r_t^e = \frac{\alpha\beta + \varepsilon}{(1-\varepsilon)\beta\eta} \frac{w_t}{w_{t-1}} - \frac{1-\eta}{\eta} r_t^x \tag{A.6}$$

$$p_t = \frac{1-\alpha}{1-\varepsilon}\frac{z^y}{z^x}\left(\frac{\varphi_t}{1-\alpha}\right)^{\alpha}\left(\frac{\beta}{1+\alpha\beta}(w_{t-1}-\varphi_t)\right)^{-\varepsilon} \tag{A.7}$$

When the economy is constrained by financial frictions, the borrowing constraint (1.9), reformulated as $\left(1-\theta\frac{r_t^y}{r_t^x}\right)\varphi_t = \eta w_{t-1}$, is used to derive (1.15) - (1.17). In the case when the borrowing constraint is slack, return rate equalization between X- and Y-sector instead closes the model.

In both cases, the only time-varying element in the solutions for interest rates and prices are wage rates. As explained in the main text, the dynamic equation for wages, (1.18), guarantees that the model has a unique steady state since $\alpha \in (0,1)$. (1.18) is derived as follows:

Combining (1.2) and (1.3) to $K_t^y = \alpha\frac{w_t}{r_t^y}\frac{\beta}{1+\beta}$, equilibrium output of sector Y is

$$Y_t = z^y\left(\frac{w_t}{r_t^y}\right)^{\alpha}\left(\frac{\alpha}{1-\alpha}\right)^{\alpha}\frac{\beta(1-\alpha)}{1+\beta} \tag{A.8}$$

Using this again in (1.2) in its formulation $r_t^y = \alpha\frac{Y_t}{K_t^y}$, the unique relationship between the two factor rewards is

$$\left(r_t^y\right)^{\alpha}(w_t)^{1-\alpha} = z^y(1-\alpha)\left(\frac{\alpha}{1-\alpha}\right)^{\alpha} \tag{A.9}$$

Thus, for given parameter values $\alpha, \beta, \varepsilon, z^x, z^y, \theta$ and η, there exists a unique and time-invariant solution to the interest rates given a wage level w_{t-1}.

The equivalent to (1.18) for the non-binding case is

$$w^{\overline{aut}} = (1-\alpha)\left(z^y\right)^{\frac{1}{1-\alpha}}\left(\frac{\psi}{1+\psi}\right)^{\frac{\alpha}{1-\alpha}} > w^{aut} \tag{A.10}$$

A.3 Proof of Lemma 2

In the model set-up, it was assumed that both regions are in their respective steady states before starting to integrate. Thus, r_t^{RoW} and r_t^{aut} are time-invariant, and so by (1.21) is r_t^{int}. Hence, by (1.29), there exists a unique and stable wage level. Since wages are the only dynamic element in the financial integration equilibrium (1.22) - (1.30), the system has a unique and stable steady state, and all variables in the equilibrium can be expressed as functions uniquely of $\alpha, \beta, \varepsilon, z^x, z^y, \theta^{RoW}$, θ and η.

A.4 Proof of Proposition 2

Rates of Return

By Proposition 1, r^x increases in θ whereas r^e decreases in θ. If $\theta^{RoW} > \theta$, it must therefore be that $r^{x-aut} < r^{x-RoW}$ and $r^{e-aut} > r^{e-RoW}$. Then, by (1.21), the wedge increases in λ. The opposite holds if $\theta^{RoW} < \theta$.

Wages

I prove in the following that w^{int} monotonically increases in λ when $\theta < \theta^{RoW}$. The proof works in the opposite direction for $\theta^{RoW} > \theta$.

Consider two different degrees of financial integration, λ and $\widetilde{\lambda}$, with $\lambda > \widetilde{\lambda}$. Then,

$$\left(w^{int}\right)^{\frac{1-\alpha}{\alpha}}\left[(1-\alpha)(z^{y})^{\frac{1}{1-\alpha}}\right]^{-1} = (1-\eta)\frac{\psi\theta}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)} + \eta\frac{1-\theta}{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)} \tag{A.11}$$

and

$$\left(\widetilde{w}^{int}\right)^{\frac{1-\alpha}{\alpha}}\left[(1-\alpha)(z^{y})^{\frac{1}{1-\alpha}}\right]^{-1} = (1-\eta)\frac{\psi\theta}{1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)} + \eta\frac{1-\theta}{1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)} \tag{A.12}$$

Now assume that $w^{int} > \widetilde{w}^{int}$. Then,

$$\begin{aligned}
&\frac{(1-\eta)\psi\theta}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}+\frac{\eta(1-\theta)}{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)} > \frac{(1-\eta)\psi\theta}{1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)}+\frac{\eta(1-\theta)}{1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)} \\
&\iff \frac{\psi\theta}{1-\theta}\left(\frac{\psi\left(\widetilde{\lambda}-\lambda\right)\left(\theta^{RoW}-\theta\right)}{\left[1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)\right]\left[1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)\right]}\right) \\
&\qquad < \frac{\eta}{1-\eta}\left(\frac{\left(\widetilde{\lambda}-\lambda\right)\left(\theta^{RoW}-\theta\right)}{\left[1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right]\left[1-\theta-\lambda\left(\theta^{RoW}-\theta\right)\right]}\right)
\end{aligned} \tag{A.13}$$

Since $\widetilde{\lambda} < \lambda$,

$$\iff \frac{\psi^2\theta}{1-\theta}\frac{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)} < \frac{\eta}{1-\eta}\frac{1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)}{1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)} \tag{A.14}$$

We know that $\frac{\eta}{1-\eta} < \psi$ (from the assumption that in both regions, the borrowing constraint holds). Then,

$$\begin{aligned}
&\frac{\psi^2\theta}{1-\theta}\frac{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)} < \psi\frac{1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)}{1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)} \\
&\iff \underbrace{\frac{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}{1-\theta}}_{<1}\underbrace{\frac{\psi\theta}{1+\psi\left(\theta+\lambda\left(\theta^{RoW}-\theta\right)\right)}}_{<1} < \underbrace{\frac{1+\psi\left(\theta+\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)\right)}{1-\theta-\widetilde{\lambda}\left(\theta^{RoW}-\theta\right)}}_{>1}
\end{aligned} \tag{A.15}$$

which is a true statement. So from $\lambda > \widetilde{\lambda}$ follows $w^{int} > \widetilde{w}^{int}$. More generally, w^{int} is a monotonically increasing function of λ.

Note that even when $\lambda = 1$, the country's wage will not fully converge to that of RoW, since it continues to be a negative function of θ.

Prices

The steady state price in both autarky and integration is described by

$$p = \frac{1}{\varepsilon z^{x}}\left(\frac{w}{1-\varepsilon}\right)^{1-\varepsilon}(r^{x})^{\varepsilon} \tag{A.16}$$

The direction of price change that results from opening the economy thus depends on the direction of change in r^x and w, which have been derived above: If $\theta < \theta^{RoW}$, both r^x and w are monotonically increasing in λ, and so is p. If $\theta > \theta^{RoW}$, the opposite holds. Note that since wages never fully equalize between the two regions unless $\theta = \theta^{RoW}$, prices will not equalize, either.

Steady State Capital Flows

By (1.30),

$$\Phi = \frac{\beta}{1+\beta} w^{int}(1-\eta)\left[1 - \frac{1+\psi\theta}{1+\psi\theta+\psi\lambda\left(\theta^{RoW}-\theta\right)}\right] \quad \begin{cases} \geq 0 & if\, \theta^{RoW} > \theta \\ \leq 0 & if\, \theta^{RoW} < \theta \end{cases}$$

$$\Omega = \frac{\beta}{1+\beta} w^{int}\eta\left[1 - \frac{1-\theta}{1-\theta-\lambda\left(\theta^{RoW}-\theta\right)}\right] \quad \begin{cases} \leq 0 & if\, \theta^{RoW} > \theta \\ \geq 0 & if\, \theta^{RoW} < \theta \end{cases}$$

Both expressions equal zero only when $\lambda = 0$, and are strictly monotone functions of λ. Without making additional assumptions on the parameter values, it is not possible to determine the direction of net capital flows. However, for the extreme case where $\lambda = 1$, the direction of net capital flows can be determined by the following proof by contradiction: Assume that $\theta^{RoW} > \theta$ and at the same time $\Phi + \Omega > 0$ (and consequently, $\Phi^{RoW} + \Omega^{RoW} < 0$). Using the more general definition $\Phi + \Omega = \frac{\beta}{1+\beta} w\left[1 - \alpha\frac{1}{r^x}\left(\frac{1}{\psi}+\theta\right) - \alpha\left(1-\theta\right)\frac{1}{r^e}\right]$, we get

$$\begin{aligned} &\frac{\beta}{1+\beta} w^{int}\left[1 - \alpha\frac{1}{r^x}\left(\frac{1}{\psi}+\theta\right) - \alpha\left(1-\theta\right)\frac{1}{r^e}\right] > 0 \\ &\qquad > \frac{\beta}{1+\beta} w^{RoW}\left[1 - \alpha\frac{1}{r^x}\left(\frac{1}{\psi}+\theta^{RoW}\right) - \alpha\left(1-\theta^{RoW}\right)\frac{1}{r^e}\right] \end{aligned} \tag{A.17}$$

Since wages are positive, for the inequality to hold it must be that

$$\begin{aligned} 1 - \alpha\frac{1}{r^x\psi} - \alpha\frac{1}{r^x}\theta^{RoW} - \alpha\left(1-\theta^{RoW}\right)\frac{1}{r^e} > 0 > 1 - \alpha\frac{1}{r^x\psi} - \alpha\frac{1}{r^x}\theta - \alpha\left(1-\theta\right)\frac{1}{r^e} \\ \iff \frac{1}{r^x}\theta^{RoW} + \left(1-\theta^{RoW}\right)\frac{1}{r^e} < \frac{1}{r^x}\theta + \left(1-\theta\right)\frac{1}{r^e} \\ \iff \left(\theta^{RoW}-\theta\right)r^e < \left(\theta^{RoW}-\theta\right)r^x \end{aligned} \tag{A.18}$$

We know that $r^e > r^x$, so with $\theta^{RoW} > \theta$, this statement is wrong. If the rest of the world is more developed, it must hence be that $\Phi + \Omega < 0$. The proof works in the opposite direction when the rest of the world is less developed.

A.5 Price Dynamics Under the Source Principle

Assume the same set-up as in the baseline model, but now entrepreneurs borrow in and under the conditions of their country of origin. Entrepreneurs in the small open economy thus face $\theta \equiv \theta^{SOE}$, and entrepreneurs in RoW θ^{RoW}. In a financially integrated world, both entrepreneurial capital and financial capital can flow freely across borders. This will, for each type of agents in each region, lead to cross-border return rate equalization: Households of SOE get the same returns r^{h-SOE} whether investing/lending money in SOE or RoW, households of RoW get r^{h-RoW} everywhere, entrepreneurs of SOE r^{e-SOE} and those of RoW r^{e-RoW}. Households are not restricted in their investment decisions, and therefore $r^{h-RoW} = r^{h-SOE}$ must hold in the financial integration equilibrium. (The same does not need to be true for entrepreneurs.)

Combine (1.7), (1.8) and (1.9) to

$$\frac{1}{r_t^{y-f}} = \frac{\theta^f}{r_t^{h-f}} + \frac{(1-\theta^f)}{r_t^{e-f}}$$

where $f = SOE, RoW$ denotes the region of origin. A region f-entrepreneur, whether operating in SOE or RoW, faces the same θ^f and r_t^{h-f}, and gets the same rate of return on entrepreneurial capital r_t^{e-f}, so that the project return r_t^{y-f} must also be the same independently of the region in which the firm operates.

At the same time, labor mobility within countries and competitive markets imply that all Y-sector firms operating in one region must earn the same return on capital r^y. But then, we can invoke (A.9) to derive that wages equalize across countries. It follows from (A.16) that prices equalize across countries.

B. Data

B.1 Countries in the Dataset

The countries included in the regressions are: Albania, Argentina, Armenia, Australia, Austria, Burundi, Belgium, Benin, Burkina Faso, Bangladesh, Bulgaria, Bosnia Herzegovina, Belarus, Bolivia, Brazil, Bhutan, Botswana, Central African Republic, Canada, Chile, China, Ivory Coast, Cameroon, Columbia, Comoros, Cabo Verde, Czech Republic, Germany, Denmark, Dominican Republic, Ecuador, Egypt, Spain, Estonia, Ethiopia, Finland, Fiji, France, Georgia, Ghana, Gambia, Guinea-Bissau, Greece, Guatemala, Honduras, Croatia, Hungary, Indonesia, India, Icleand, Italy, Jamaica, Jordan, Kazakhstan, Kenia, Kyrgyz Republic, Cambodia, South Korea, Laos, Sri Lanka, Lesotho, Lithuania, Latvia, Morocco, Moldova, Madagascar, Maldives, Mexico, Macedonia, Mali, Mongolia, Mozambique, Mauritania, Malawi, Namibia, Niger, Netherlands, Norway, Nepal, New Zealand, Pakistan, Peru, Poland, Portugal, Paraguay, Romania, Russia, Rwanda, Sudan, Senegal, Sierra Leone, El Salvador, Serbia, Sao Tome and Principe, Suriname, Slovakia, Slovenia, Sweden, Chad, Togo, Tajikistan, Tunesia, Turkey, Tanzania, Uganda, Ukraine, United States, Venezuela, Vietnam, Yemen, South Africa, Zambia.

B.2 Additional Regression Results

Table 1.7 reproduces Table 1.3, but I replace credit market development with a joint measure of credit and bonds market development. Due to the fact that data on bonds market development is available only for some, mostly economically developed, countries, the sample is

Table 1.7: Convergence regressions with credit and bond market development

dependent variable: relprice deviations from cross-country mean

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
fi	-0.0211***	-0.0168***	-0.0166***	-0.0192***	-0.0188***	-0.0198***	-0.0238***	-0.0264***
	(0.0030)	(0.0038)	(0.0028)	(0.0036)	(0.0034)	(0.0040)	(0.0044)	(0.0045)
fd_joint	0.1797***	0.2133***	0.2174***	0.2007***	0.2085***	0.1993***	0.1872***	0.1713***
	(0.0507)	(0.0394)	(0.0347)	(0.0318)	(0.0290)	(0.0323)	(0.0290)	(0.0302)
open			-0.0190	0.0079	-0.0263	0.0058	-0.0313	-0.0007
			(0.0494)	(0.0449)	(0.0474)	(0.0447)	(0.0523)	(0.0465)
gdp			-0.1456***	-0.1943***	-0.1356***	-0.1894***	-0.1513***	-0.1772***
			(0.0391)	(0.0282)	(0.0367)	(0.0256)	(0.0466)	(0.0343)
rule_law			0.1064**	0.1092**	0.1103**	0.1108**	0.0973**	0.0963**
			(0.0377)	(0.0402)	(0.0401)	(0.0441)	(0.0387)	(0.0409)
bcrisis					0.0062	0.0021		
					(0.0166)	(0.0180)		
bcrisis*fi					0.0027	0.0007		
					(0.0038)	(0.0033)		
fi_ci							0.0081	0.0067
							(0.0325)	(0.0343)
fi_ci*fi							0.0151***	0.0175***
							(0.0047)	(0.0035)
Constant	0.1237***	0.1076***	1.4063***	1.8861***	1.3197***	1.8514***	1.4695***	1.7241***
	(0.0320)	(0.0171)	(0.3509)	(0.2331)	(0.3321)	(0.2138)	(0.4055)	(0.2840)
Observations	467	467	436	436	436	436	436	436
No. of groups	34	34	34	34	34	34	34	34
R-squared	0.0807	0.1816	0.1457	0.2317	0.1504	0.2319	0.1578	0.2457
Country FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

not representative of the universe of countries anymore. Note that the relative price cross-sectional mean still refers to the full sample of 113 countries. In these regressions, again, the convergence hypothesis is strongly supported. The coefficient is slightly higher. *fd_joint* is more significant than *fd_credit* before. Some of the other coefficients are changed in sign or significance level of the estimates. This is however due to the different sample composition rather than the different specification of financial development: They are reproduced when using the same sample, but the previous definition of financial development as credit market development. It seems that in richer countries, economic development has a converging effect, whereas rule of law has a diverging effect. Countries with open capital accounts exhibit less of a convergence effect. The underlying reasons for these findings, though certainly interesting, are beyond the scope of this paper.

Chapter 2

External Asset Positions, Demography and Life-cycle Portfolio Choice

With Margaret Davenport

How do demographic differences between regions affect external positions in safe and risky assets? We answer this question focusing on the US vis-à-vis 15 European Union member states. The US bilateral position is characterized by risky assets alongside safe liabilities. At the same time, the US population is relatively younger. We present a structural model of two fully integrated regions, which differ by the age structure of their populations. There are multiple overlapping generations of agents, who choose a portfolio of safe and risky assets over the life-cycle. We show that the younger region has a higher relative demand for risky assets, which induces international asset trades. In a simulation starting in 1990, we replicate the observed positions between the US and the European countries both in sign and in magnitude. We predict the risk asymmetry to persist until the end of the century, whereas both safe and risky returns decline persistently.

1. Introduction

Over the last few decades, the US has accumulated net risky assets alongside net safe liabilities vis-à-vis the rest of the world. The literature has offered purely financial explanations for this asymmetry: First, the US as the center of the international monetary system acts as a global liquidity provider by exporting safe, short-term assets (Gourinchas and Rey, 2007). Second, a lack of financial development in many of its counterparts not only attracts foreigners to US debt, but makes US investors turn towards foreign destinations in their search for yield (Bernanke, 2005; Mendoza, Quadrini, and Ríos-Rull, 2009; Von Hagen and Zhang, 2014). This paper provides evidence of such portfolio imbalances between the US and a group of developed European countries. Most of them are reserve issuers and have highly developed financial systems. This suggests another driving force for the US risk profile. The current paper establishes diverging demographic trends as a suitable candidate.

The external position of the US vis-à-vis 15 European Union member states (henceforth "EU")[1] is characterized by net equity assets of around 4 percent of GDP, and net debt liabilities of 5 percent. At the same time, the US and EU populations differ in their demographic composition: The US population is at present younger than the population in the EU, and over the next decades is projected to age at a slower speed. This paper shows that demographic differences can explain bilateral imbalances. We build a structural model with overlapping generations. There are two regions, whose populations differ in fertility and longevity. Agents earn a high labor income during working life, and a low pension income after a fixed retirement age. They smooth consumption by saving into two types of assets: a safe asset with a certain return, and a risky asset with a stochastic return. Asset supply is proportional to population size. Savings patterns are age-dependent, with portfolio risk decreasing over the life-cycle. As a result, aggregate demand for safe and risky assets reflects the age structure of the population. In a financially integrated world, the relatively younger region imports risky assets and exports safe assets. Calibrating the model to the US and EU economies for 1990 onwards, we obtain bilateral equity and debt positions of magnitudes similar to the data. Since demographic trends are persistent, so are bilateral imbalances.

How do differences in demographics affect regional asset demand, and ultimately external positions? Using data from the US Survey of Consumer Finances, we estimate life-cycle asset holdings and portfolio choice. Savings are hump-shaped over the life-cycle, whereas the portfolio share of risky assets declines with age. Aggregate asset demand in a region should therefore reflect the age composition of the population. When there is a larger share of older cohorts in the economy, the aggregate country portfolio should be shifted towards safe assets. The total amount of savings should be higher, the larger the amount of individuals near or at retirement age. At the same time, with fixed retirement age and constant retirement income, changes in longevity will shift the share of life that individuals expect to spend in retirement with low non-financial income. In order to smooth consumption over the life-cycle, they will make adjustments to their optimal savings and portfolio choice.

While per capita asset demand should thus vary with demographics, per capita asset supply is kept constant in the model. Each individual is endowed with an identical Lucas (1978) tree that pays a safe and a risky dividend. Therefore, in order for markets to clear, there will be an effect on rates of return. When there are two regions whose populations age at different speeds, they will have different market-clearing returns in autarky. Under financial integration, they will trade assets with each other until returns equalize. The resulting market-clearing returns should reflect the age structure of both regions together.

We simulate the demographic transitions of the US and the EU between two steady states: 1950 and 2095. We assume that the two regions are financially integrated from 1990 onwards. This corresponds to the point in time when EU countries undertook major steps towards opening their capital accounts. The resulting US net external position is characterized by risky assets and safe liabilities throughout the whole simulation period. Magnitudes vary over time, reflecting how the share of retirees evolves in the US relative to the EU: Whenever the US has a relatively lower share of retirees (dis-savers), its net external position is more positive. On average, the model predicts safe US liabilities of 5 percent of GDP and risky

[1]This includes all members of the European Union prior to the 2004 enlargement: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden and the United Kingdom.

assets of 2 percent of GDP, which comes very close to the data. Furthermore, the model can explain almost all the decline in safe returns between 1990 and 2015, and a large share of the movements in the risk premium. The rates of return continue to fall until the end of the simulation period. An exception is a short period between 2020 and 2040, in which the baby boomer cohorts run down their assets during retirement.

We decompose the effect of demographics on aggregate asset demand and on the portfolio share of risky assets into three main channels: A *distribution effect*, acting through the shift in the population age structure from younger cohorts towards older cohorts. A *life-cycle effect* capturing the effect that increasing individual longevity has on individual decision-making. Lastly, a *valuation effect*, which reflects changes in the equilibrium demand for safe and risky assets in response to movements in market-clearing rates of return. In a counterfactual exercise, we find that the distribution effect accounts for more than 50 percent of aggregate asset demands in the US by the end of the demographic transition, and increases both over time and with cohort age. For the risky share, the valuation effect is the most important one.

Our paper contributes to understanding the observed asymmetries in the international asset position of the US. To our knowledge, we are the first to systematically analyze demographics as an explanatory factor. The results of this paper show that differences in the age structure of the population can be powerful determinants of debt and equity positions between two developed regions, which the existing literature has focused on to a lesser extent. However, the insights of this paper are more general, and extend to the US external position vis-à-vis other countries. Our results also allow for a long-run outlook on external positions. Future demographics are to a large extent known well in advance. Therefore, our predictions involve only a limited degree of uncertainty. The bilateral equity and debt positions that we predict arise from optimal individual decision-making in two financially integrated regions with different demographic structure. They are thus neither destabilizing nor unsustainable.

In what follows, we place our contribution in the context of the related literature. Section 2 presents the empirical evidence that motivates our research. We introduce the model in Section 3, before explaining how it is calibrated in Section 4. Section 5 presents the simulation results for the effects of demographic change on external positions and rates of return. We further assess the three channels through which demographics operate. Section 6 concludes.

1.1 Related Literature

A growing literature in international macroeconomics has found demographics to be an important determinant of the behavior of capital flows and returns, not distinguishing between asset types (Backus, Cooley, and Henriksen, 2014; Krueger and Ludwig, 2007; Börsch-Supan, Ludwig, and Winter, 2006; Attanasio, Kitao, and Violante, 2007; Barany, Coeurdacier, and Guibaud, 2015; Ferrero, 2010). Krueger and Ludwig (2007) present a complex structural model with several world regions and a large number of cohorts. They emphasize the need for an open economy framework to study the impact of demographic change as the effect on rates of return is more extreme if less favorable demographics are imported from other world regions. Backus et al. (2014) build a relatively simple model that can match persistent capital flows in the US, Japan, China and Germany and project declining rates of return until 2030. Barany et al. (2015) analyze the interaction of demographics with financial development and pension availability in a structural model of a developed and a developing economy that can replicate global imbalances. We are not aware of structural models that study the link between regional demographic differences and trade in *different* types of assets. In an empirical analysis De Santis and Lührmann (2009) consider - among other things - the link between demographics and various types of capital flows. They find that a higher old-age dependency ratio induces net equity inflows and net outflows of debt instruments. This is in line with our results.

Our paper also relates to closed economy studies of the effects of demographic change on prices and returns of assets with different degrees of riskiness. This literature builds on life-cycle models of asset accumulation and portfolio choice (Modigliani and Brumberg, 1954, 1990; Merton, 1971). A consensus emerges that the entry of a large cohort into the labor market will result in downward pressure on asset returns, as asset demand increases during prime saving years (Abel, 2001, 2003; Poterba, 2001; Brooks, 2002, 2004; Geanakoplos, Magill, and Quinzii, 2004; Gagnon, Johannsen, and Lopez-Salido, 2016; Carvalho, Ferrero, and Nechio, 2016). However, there is disagreement about the effect on the equity premium. In a general equilibrium model with overlapping generations, Brooks (2002) finds that the equity risk premium should fall when a large cohort retires, but later shows that this is dependent on unconstrained borrowing of young agents and a zero net supply assumption (Brooks, 2004). Kuhle, Ludwig, and Boersch-Supan (2007) present a richer model with bonds in positive net supply and find that the risk premium first falls with increased saving prior to retirement and then increases as the large cohort retires. In a recent paper, Kuhle (2017) analytically derives a positive relationship between fertility and the equity premium that holds for concave utility functions and bonds in positive net supply, due to a stronger reaction of the risky rate of return to asset demand. In discussing our simulation results for the rates of return, we will refer to his insights.

Finally, our paper motivates the asymmetric composition of the US external position with its relative demographic developments, and thus relates to papers that document this asymmetry and provide potential explanations. Lane and Milesi-Ferretti (2007) provide a detailed analysis of US external debt and equity positions, whereas Gourinchas and Rey (2007) explore the US "exorbitant privilege" – the fact that despite being a net debtor vis-à-vis the rest of the world, the US earns a positive net investment income. They explain this by the role of the US as the center country of the international monetary system. By issuing safe, liquid assets, the US provides insurance to the rest of the world, and in return enjoys an insurance premium from risk and maturity transformation. This role has been further investigated by Gourinchas, Rey, and Govillot (2010) and Curcuru, Dvorak, and Warnock (2008), among others. A second strand of the literature has focused on cross-country differences in the development of the financial sector Bernanke, 2005; Mendoza et al., 2009; Von Hagen and Zhang, 2014; Ju and Wei, 2014; see also Chapter 1). Due to a limited capacity to generate safe, liquid financial assets, demand from financially less developed countries for this type of assets needs to be satisfied abroad. Conversely, foreign investors may bypass domestic institutions and exploit profitable risky investment opportunities. Both liquidity provision and lacking institutional capacity may be less relevant for bilateral US-EU asset positions. Our paper therefore contributes to understanding imbalances between two developed regions.

2. Stylized Facts

Our research is motivated by two observations: pronounced bilateral external asset positions between the US and the EU, and differences in the age structure between the two regions. In the following, we present empirical evidence for both. Then, we show estimates of household life-cycle asset holdings and portfolio choice based on data from the US Survey of Consumer Finances. To motivate that demographic trends can potentially have large effects on aggregate savings, we carry out a statistical exercise: We combine our life-cycle estimates with data and predictions for relative cohort sizes at various points in time.

2.1 Bilateral Positions, US vis-à-vis EU

External asset positions between the US and the EU have been characterized by an asymmetry by asset type since the 1990s. Figure 2.1 depicts bilateral positions in safe and risky assets over GDP for the US vis-à-vis the EU. Safe assets are comprised of short-term and long-term debt instruments and claims from the banking sector. This corresponds to the net debt position of the portfolio investment category and the net position from the other investment category of the International Investment Position. Risky assets correspond to equity. We use data on the net equity position of the portfolio investment category. In our definition of safe and risky assets, we follow Gourinchas and Rey (2007) and Gourinchas et al. (2010), but exclude foreign direct investment, as we focus on household portfolio choice. To construct the time series, we combine data from the IMF's Coordinated Portfolio Investment Survey (CPIS) on equity and long-term debt with data on net claims from the banking sector for short term instruments from the US Treasury International Capital (TIC). A detailed description of the construction of these positions is presented in Appendix A.1, along with a definition of asset types in Appendix A.2. The different time coverage of safe and risky positions is due to the use of different data sources.

The net bilateral asset position remains small but exhibits a slight downward trend: Since the 2000s, the US has been a net debtor vis-à-vis the EU. The net position conceals considerable divergence in external positions by asset type. The US has a net external asset position in risky assets over GDP between 2 and 5 percent between 1994 and 2015. This is offset by a net external liability position in safe assets between 2 and 7 percent of GDP. Thus, the US is a net exporter of safe assets and a net importer of risky assets.[2]

On a global scale, bilateral asset holdings between the US and the EU are important: The US and the EU are by far the largest international investors, jointly holding between 60 and 70 percent of the world's external debt assets and more than 70 percent of equity assets between 1990 and 2014 (see Figure 2.15 in the Appendix). Also, the US is the most important investor

[2]EU debt instruments include government bonds issued by countries that were perceived as having a high default probability in some of the recent years, and in particular by Greece, which negotiated a debt haircut in 2011. This could potentially make the safe external position to some extent risky. However, Greek debt makes for less than 0.3 percent of US bilateral debt assets in CPIS data and the joint share of Greece, Spain, Portugal and Ireland – the countries that received financial assistance by the European Stability Mechanism – is 7.7 percent. Excluding this type of debt does not change the overall bilateral pattern. The safe asset position also includes low quality corporate and asset-backed debt of both regions. In the TIC data, corporate bonds make for 60% of EU holdings of US long-term securities between 2011 to 2016, of which however only about 6% should be asset-backed debt securities, this last number being based on less detailed data on US holdings vis-à-vis the rest of the world.

Figure 2.1: Bilateral debt and equity positions US vs. EU

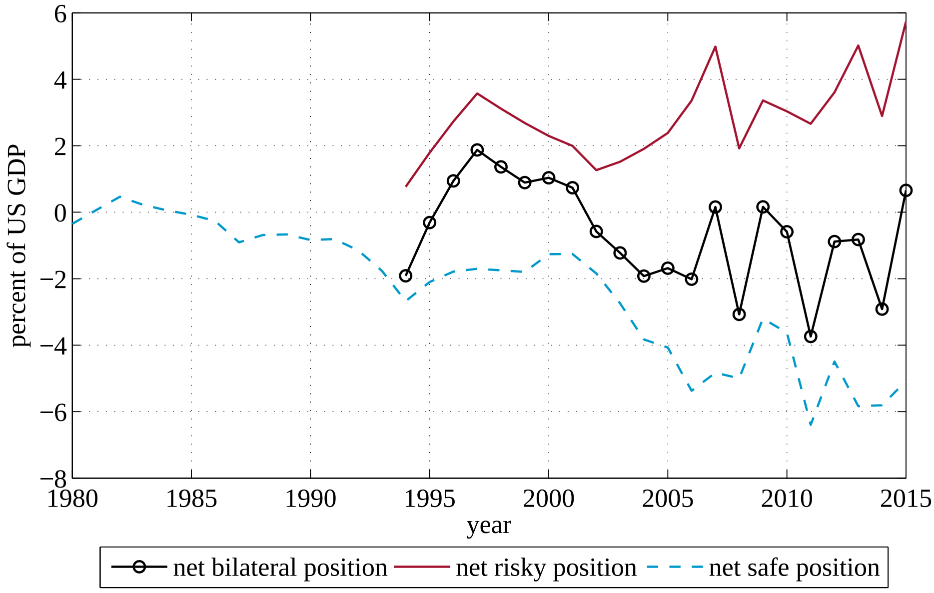

Net safe position: net banking claims plus net portfolio debt assets of the US vis-à-vis the EU. Net risky position: net portfolio equity assets. Net bilateral position: Net safe position + net risky position. *Source:* Coordinated Portfolio Investment Survey (IMF), US Treasury International Capital.

in the EU and vice versa. In Appendix Table 2.6 we consider the share of total US assets that originates from the EU, and vice versa. For both debt and equity, the bilateral share lies between 0.4 and 0.5 in the US as well as in the EU. The shares are larger when including in the bilateral position assets from offshore financial centers, which act as intermediators of US-EU investments (Milesi-Ferretti and Lane, 2010). The numbers necessarily exclude any investment that is unreported, a large portion of which should also be bilateral (Zucman, 2013). This motivates our two-region approach, supported by the fact that the two regions exhibit a high degree of de-jure financial openness since the 1990's (Chinn and Ito, 2008).

However, the association between demographics and the risk-content of external positions can also be found in a larger set of countries (see Appendix A.5 for a detailed analysis).

2.2 Demographic Trends in the US and the EU

The US and the EU exhibit large variations in demographic trends. The left panel of Figure 2.2 shows data and projections for the two regions' old-age dependency ratios (the number of people aged 65 or older relative to the number of people 20 to 64 years old) between 1950 and 2095 (United Nations, 2015). While both regions are aging quite dramatically over this period, the level and trends differ: Dependency ratios were similar until the 1980s, but have been diverging since, with the EU aging faster. The difference between the dependency ratios is projected to reach 0.2 by 2050. There are also important differential trends throughout the transition. In particular, large increases in the dependency ratios in both regions are predicted from now on until around 2040, due to the retirement and death of the large baby boomer cohorts on both sides of the Atlantic. This makes demographic change a particularly relevant

Figure 2.2: Demographic trends US vs. EU, old-age dependency ratio

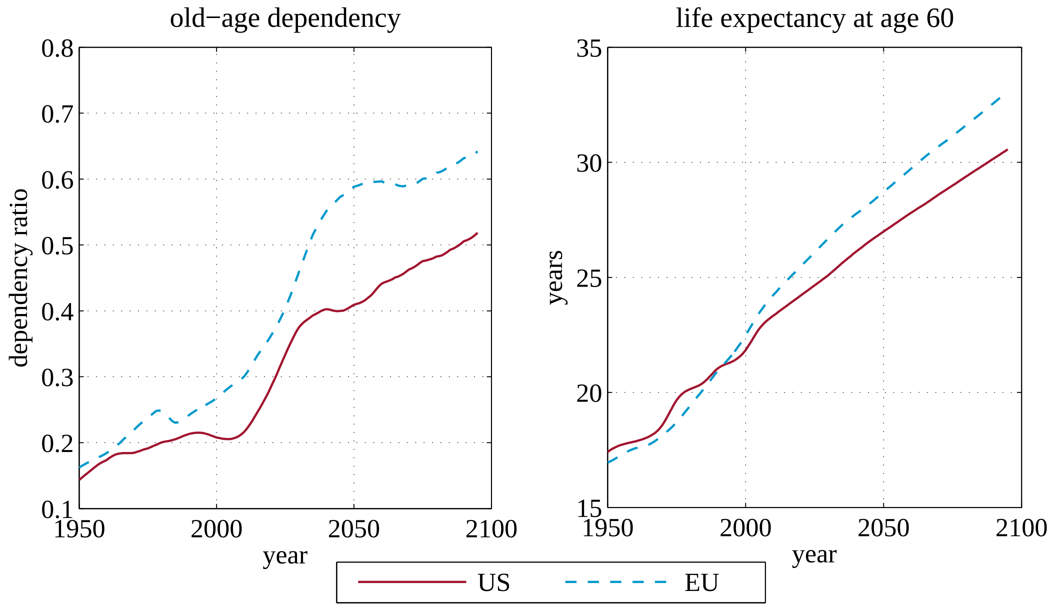

Left panel: Old age dependency ratios (population aged 65 and older divided by population aged 20 to 64) of the US and the EU. The graph displays data for 1950 to 2015 and projections for a medium variant for fertility and normal migration afterwards. Right panel: Life expectancy at age 60 for the medium variant for fertility and normal migration is shown for the US and the EU. *Source:* United Nations Population Division.

phenomenon within the next decades.

The old age dependency ratio reflects the age composition of the population due to both declining birth rates and increasing longevity. The right panel of Figure 2.2 shows data and projections for the life expectancy at age 60 in the US and EU. While individuals in the US were on average living longer until the mid-1990s, this has since reversed and by 2095, individuals in the EU can expect to live 2 years longer at age 60 than their US counterparts.

2.3 Savings, Portfolio Choice, and Demographics

How might the level and trend of demographic change affect aggregate savings and portfolios in the two regions? Figure 2.3 shows estimates of total financial assets and the share of risky assets of the household head by age, conditional on participation in the stock market.[3] We use data from 9 waves of the US Survey of Consumer Finances (SCF) between 1989 and 2013.

In order to extract age-specific effects from repeated cross-sectional data, we apply a technique developed by Deaton and Paxson (1994) and Deaton (1997), which is based on the assumption that only cohort- and age-specific effects exhibit a trend, whereas time-specific effects are due to business-cycle fluctuations and should therefore net to zero.[4] For the risky

[3]The definitions of safe and risky assets are given in Appendix A.2 and correspond to how assets are classified in the Survey of Consumer Finances. We focus on the risky share conditional on participation in the stock market because we do not explicitly model the participation decision of households.

[4]The underlying problem is that we have to account for time effects (e.g., due to aggregate economic condi-

Figure 2.3: Total financial assets and risky portfolio share by age, United States

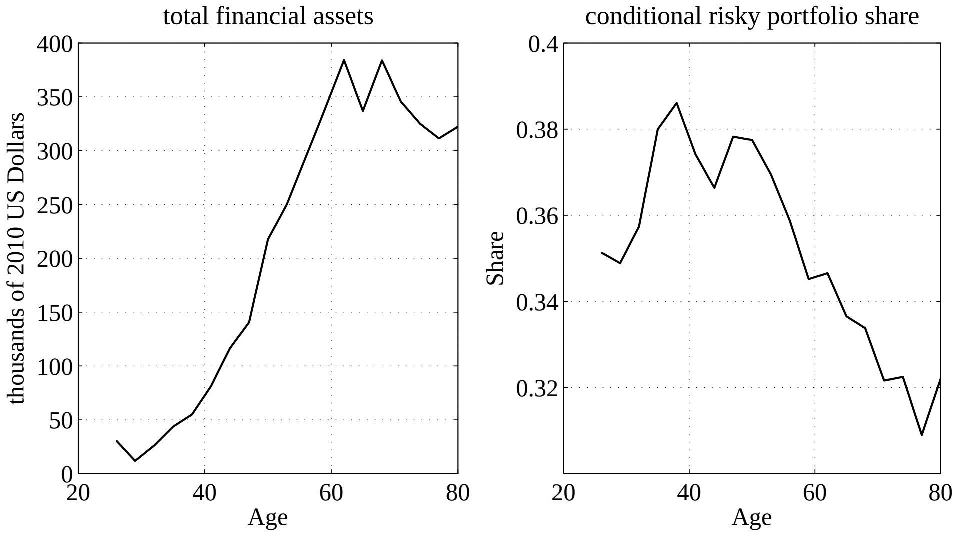

Left panel: Estimated total financial assets in thousands of 2010 US dollars. Financial assets are defined as the sum of safe and risky assets, as given in Table 2.3 in the Appendix. Right panel: Estimated conditional risky share. *Source:* Estimated from 1989-2013 waves of the Survey of Consumer Finances.

share, we apply a two-stage Heckman estimation that takes into account a correlation between the decision to participate in the stock market and the risky share of assets, which is well documented in the household finance literature (Guiso, Haliassos, and Jappelli, 2002; Guiso and Sodini, 2012). Details on this estimation procedure are provided in Appendix A.3.

Savings are hump-shaped over the life-cycle, increasing smoothly until retirement age, and gradually declining thereafter. The conditional portfolio share of risky assets peaks in the 30s and declines over the rest of the life-cycle. A risky share that is decreasing with age is in line with recent evidence presented in Fagereng, Gottlieb, and Guiso (2017) and further discussed in Guiso and Sodini (2012).[5] It is also consistent with findings that the willingness to take risk declines linearly from early adulthood until retirement age (Dohmen, Falk, Golsteyn, Huffman, and Sunde, 2017).

Ideally, we would apply an identical estimation approach to analyze household behavior in the EU, however, survey data is not available over a long enough period. Recently, the Eurosystem has conducted the Household Finance and Consumption Survey (HFCS) that covers 12 out of the 15 countries that comprise our EU region. We compare the data from SCF and HFCS in Figure 2.19 in the Appendix for pooled waves of each survey. The asset holdings, while slightly lower for the EU, exhibit a similar age-dependent structure and the unconditional risky share across age groups is very similar, which is in line with observations

tions) and cohort effects (e.g., due to economic conditions encountered at the beginning of working life, shown to affect economic decisions by Cogley and Sargent (2008) and Malmendier and Nagel (2011)). Because of a linear relationship between calendar year, birth year and age, it is impossible to estimate each of these three effects simultaneously from repeated cross-sections of survey data.

[5]Earlier literature, for example Heaton and Lucas (2000), Poterba and Samwick (2001) and Ameriks and Zeldes (2004), has reported a more pronounced hump-shape of the conditional risky share over the life-cycle. The differences are due to the recent use of more sophisticated estimation techniques and more comprehensive data. See Guiso and Sodini (2012) for details.

Figure 2.4: Hypothetical asset holdings and risky share

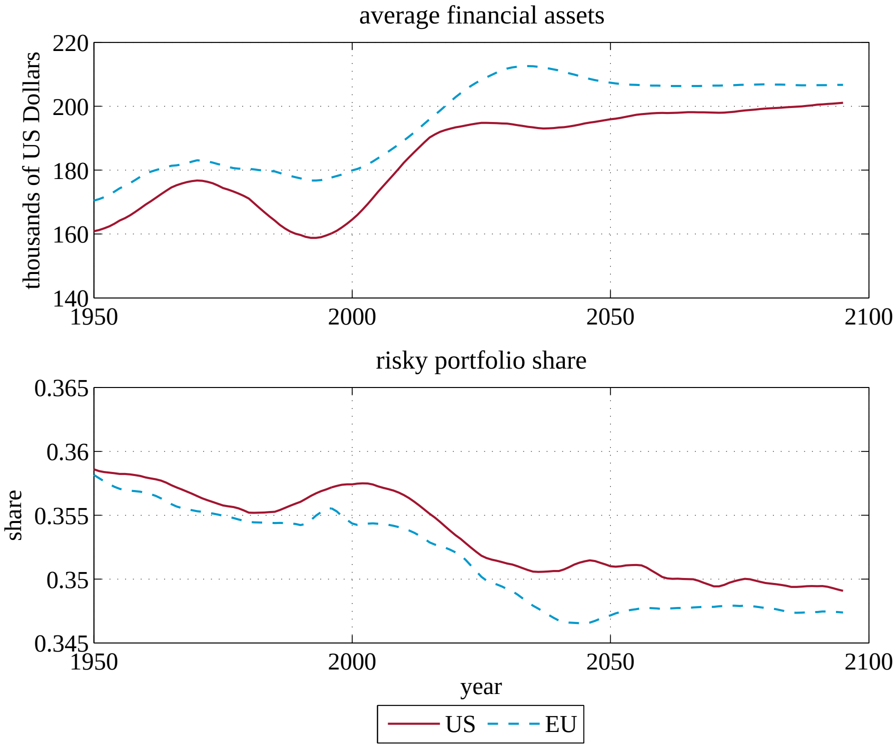

Upper panel: per capita financial asset holdings in thousands of 2010 US Dollars. Lower panel: per capita risky portfolio share. *Source:* Authors' calculations using data from SCF and UN Population Prospects.

of cross-country differences in asset holdings and portfolios (Guiso et al., 2002; Guiso and Sodini, 2012).

The life-cycle pattern of financial assets and portfolio allocation is important in the context of the demographic trends described above. We want to explore the implications of a changing population age structure on aggregate asset holdings and the risky portfolio share. To this end, we carry out the following statistical exercise: We keep the level of financial assets and the risky share by age group constant at (moving averages of) the levels shown in Figure 2.3. Then we multiply by the actual or projected number of people at each age group and divide by the total population of 26- to 79-year olds in each year. This assumes that life-cycle financial decisions will stay unchanged over the next decades, but demographics will change, in line with the UN predictions. The resulting hypothetical average per capita financial assets and the risky share by year are shown in Figure 2.4. Figure 2.16 in the Appendix shows hypothetical financial asset holdings by age group for both regions.

The aggregate implications of a changing age structure of the population are large. If life-cycle savings behavior and portfolio choice did not change (and supply reacted to satisfy any demand changes at current interest rates), per capita asset holdings in the EU between 1950 and 2095 would increase by 25 percent, and between 1990 and 2095 by 26 percent. In the US, the increase would be 21 percent between 1950 and 2095, and 19 percent between 1990 and 2095. Particularly strong increases can be observed between the years 2000 and 2040, the time period during which the baby boomers retire and finally die. In levels, agents in the EU would always choose to hold a larger level of assets on average than agents in the US, due to the fact that the EU is older during the time period shown. There would also be considerable changes in the portfolio share held in risky assets: Both regions would experience a (small) decline in the per capita risky share, but the downwards trend is stronger in the EU, as the EU population ages faster. These calculations suggest that the relative age distribution of the two regions can be a powerful determinant of relative asset demands for safe and risky assets. In addition to changes in the per capita asset demand, the total population size is predicted to increase, which will contribute to further inflating the demand for assets in both regions.

The simple statistical exercise highlights just one channel through which demographics take effect, which we will later discuss as the *distribution effect*. There are two aspects that this exercise is missing: First, we would not realistically expect that individuals keep their life-cycle pattern of savings and portfolio choice unchanged when their life expectancy increases. Second, if asset supply is not fully elastic, there would need to be a reaction in market-clearing rates of return in response to movements in aggregate demand. These would in turn lead to a re-evaluation of the households' portfolios and affect their optimal decisions. A piece of evidence pointing at the importance of valuation changes is the decline in real bond returns observed since the 1980s both in the US and in the largest EU economies (see Figure 2.18 in the Appendix), which some of the literature has attributed to demographics. Thus, while a statistical exercise can partly identify the effect of demographics on aggregate asset demand and portfolio choice, the model and simulation introduced in the following sections will allow us to obtain a more complete picture.

3. Model

We augment the workhorse overlapping generations model of Auerbach and Kotlikoff (1987) along several dimensions. First, we allow for cohort specific survival probabilities and birth rates, which enable us to model the impact of a full demographic transition. Second, we consider a two-region world, where regions differ only by their demographics. The two regions are closed until period T and completely open afterwards.[6] Third, we incorporate endogenous portfolio choice of individuals by introducing a safe asset, or bond, and a risky asset, or stock, for which returns are endogenous and perfectly correlated across regions. Finally, individuals take decisions in a stochastic environment with uninsurable idiosyncratic risk in the form of shocks to labor and pension income, and aggregate risk related to stock returns. Idiosyncratic risk generates a precautionary motive for asset accumulation, whereas aggregate risk, and how it relates to agents' future marginal utility of consumption, will determine the portfolio allocation, as well as the equilibrium risk premium that results from our general equilibrium model.

In general, we distinguish between individual choice variables and aggregate variables by referring to aggregate variables in capital letters. To keep the exposition clean, we drop country superscripts until we discuss the open economy and market clearing conditions in Section 3.5.

[6] In reality, the opening of the EU and US financial markets that took place around the 1990s happened more gradually. We will consider this in an extension in Section 5.

3.1 Demographics

We model demographics as a combination of birth rates (fertility) and survival probabilities (longevity). Agents take into account a cohort-specific series of survival probabilities when making decisions. At the same time, birth rates and survival probabilities jointly impact the aggregate size of the population, and the distribution of the population across cohorts.

In each period, a continuum of individuals, indexed by i, is born at working age N^b, forming a cohort. Each successive cohort grows at a time-varying rate, γ_t. The size of the youngest cohort $L_{N^b,t}$ evolves according to

$$L_{N^b,t+1} = (1 + \gamma_{t+1})L_{N^b,t}. \tag{3.1}$$

Individuals live to a maximum age of N^d. Prior to age N^d, survival to age n is stochastic with probability $\delta_{n-1,t}$ conditional on being alive at age $n-1$ in period t. The size of a cohort of age n in period t is

$$L_{n,t} = \left(\prod_{l=0}^{n-1} \delta_{N^b+l,t-n+l}\right) L_{N^b,t-n}. \tag{3.2}$$

Total population size is the sum over all cohorts alive at time t: $L_t = \sum_{n=N^b}^{N^d} L_{n,t}$.

Agents retire at a fixed age N^r. This implies that the share of the life-cycle that is spent in retirement will vary with changes in $\delta_{n,t}$. While this is potentially a strong assumption, it allows us to study the effect of demographics on asset demand independently of potential reforms of the retirement system. A constant retirement age is furthermore in line with data on the average effective retirement age in OECD countries since the 1970s, and with projections until 2060 (OECD, 2015).

3.2 Market Structure

There are two financial assets in the economy. Safe assets b_t, or bonds, pay a certain dividend of d_t in each period and trade at price Q_t, whereas risky assets s_t, or stocks, pay a stochastic dividend, $\tilde{d}_t$ and trade at price $\tilde{Q}_t$. Gross returns of each asset are

$$R_{t+1} = \frac{Q_{t+1} + d_{t+1}}{Q_t} \qquad \tilde{R}_{t+1} = \frac{\tilde{Q}_{t+1} + \tilde{d}_{t+1}}{\tilde{Q}_t}. \tag{3.3}$$

Returns consist of a capital gain due to changes in asset prices, $\frac{Q_{t+1}}{Q_t}$, and a dividend yield $\frac{d_{t+1}}{Q_t}$. The risky dividend is defined as $\tilde{d}_t = d_t + rp_t + \epsilon_t$, where rp_t is a risk premium and ϵ_t an *i.i.d.* shock with distribution $\ln \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

Asset supply is generated by Lucas (1978) trees. We assume that when an individual is born, two Lucas trees become available to provide financial assets to the economy: a risky tree and a safe tree. The trees live as long as the individual. Each period, the safe tree pays a safe dividend and the risky tree pays a risky dividend; claims to the dividend streams paid by the trees can be traded separately in the economy. When individuals die so do their Lucas trees, but claims traded in the previous period are still paid out. Thus, all trees of the same type are identical to potential savers, and claims will be traded at the same price, regardless of the tree owner's age. In this formulation we keep the supply side of assets simple but responsive to demographic change.

Shocks to the risky rate of return are assumed to be perfectly correlated across regions. In addition, we abstract from modeling (real) exchange rates and impose that the law of one price always holds. In the model under full financial integration, domestic and foreign assets are therefore perfect substitutes. Moreover, the international positions resulting from asset trades are net positions, whereas the gross positions are indeterminate. While this simplifies the portfolio choice problem[7], it gets down to the core of how we think about the effect of demographic change on asset positions: Demographics-induced external positions represent excess saving into a certain type of asset at world prices, motivated by regional differences in portfolio preferences. They should therefore be largely unrelated to hedging behavior, which clearly motivates international asset purchases in reality, but should affect gross positions rather than net positions. One limitation is that we cannot motivate home bias in equity and bond holdings, which the literature on the international diversification puzzle tries to explain (see Lewis, 1999; Gourinchas and Rey, 2013, for surveys). We will refer to this point in Section 5.1 when discussing the simulation results.

The market for goods consists of one single consumption good $c_{i,n,t}$, which is identical across regions. Its price is normalized to 1. With our assumption on real exchange rates, this implies that we can abstract from explicitly modeling international trade in goods. Thus, we isolate the impact of demographic change on international asset trade.

3.3 Labor and Pension Income

Labor income for individual i of age n at time t is defined as

$$y_{i,n,t} = P_{i,n,t}\theta_{i,n,t}\zeta_{i,n,t}, \tag{3.4}$$

where $\theta_{i,n,t}$ is an idiosyncratic, transitory income shock that is distributed lognormally, $\ln\theta_{i,n,t} \sim \mathcal{N}(0, \sigma^2_\theta)$ and $\zeta_{i,n,t}$ is an *i.i.d.* zero income shock that occurs with probability p. We take a broad interpretation of this shock as an unemployment spell or an employment break due to medical reasons. $P_{i,n,t}$ is a permanent income component, which evolves according to

$$P_{i,n,t} = G_n P_{i,n,t-1}\eta_{i,n,t}, \tag{3.5}$$

where G_n is an age-specific component of income. It generates a hump shape in the labor income over the life-cycle, reflecting a premium on work experience or age-dependent productivity level. The permanent income shock, $\eta_{i,n,t}$, is assumed to be log-normally distributed, $\ln\eta_{i,n,t} \sim \mathcal{N}(0, \sigma^2_\eta)$, and uncorrelated with the transitory income shock. A similar specification of the labor income process is frequently used in the life-cycle literature, going back to Zeldes (1989) and Carroll (1992), and generates both saving motivated by life-cycle consumption smoothing and a precautionary motive.[8] All shocks to labor income are assumed

[7]On the added complexity when solving for gross positions in a setting where risky returns are independent across regions, see Tille and Van Wincoop (2010) and Devereux and Sutherland (2011).

[8]Standard specifications usually do not contain zero income shocks, however Carroll (1992) shows that they are empirically important. We include them to address an issue well known in the literature on portfolio choice over the life-cycle: a model with standard values of risk aversion generates a counterfactually high risky share of assets over the life-cycle, particularly for the very young and the very old (Merton, 1971; Heaton and Lucas, 1997; Constantinides et al., 2002; Guiso et al., 2002; Cocco et al., 2005; Fagereng et al., 2017).

to be independent of the shocks to stock returns across individuals and over time. While this assumption keeps the mechanism of our model clear, it is also empirically plausible. The literature has sometimes found a small positive correlation between labor income and stock returns (Campbell, Cocco, Gomes, and Maenhout, 2001) and sometimes a small negative correlation (Heaton and Lucas, 2000). Furthermore, recent work finds that the stock component of human capital, or how much it co-moves with aggregate business cycle risk, is much lower than that which rationalizes observed portfolio allocation over the life-cycle (Huggett and Kaplan, 2016).

Retirees receive pension income as a fixed share $\phi \in [0, 1]$ of the deterministic labor income earned in the period directly preceding retirement, subject to a zero income shock $\zeta_{i,n,t}$ with the same statistical properties as during working age:

$$\bar{y}_{i,n,t} = \phi P_{i,N^r,t}\zeta_{i,n,t} \quad \forall n = N^r + 1, ..., N^d. \tag{3.6}$$

In general, zero income periods should not be encountered by retirees that receive a public pension, but we consider these shocks rather to reflect periods of high expenditure, e.g. due to health shocks.[9]

The replacement rate ϕ summarizes the provision of public pensions (for example, social security payments in the US). We choose a simple constant replacement rate pension scheme rather than explicitly modeling the pension system to highlight the direct effect of demographic change on individual behavior through the effect that it has on the effective subjective discount rate.[10] A constant replacement rate can be rationalized by interpreting pensions as home production or by assuming that gross aggregate labor income grows over time so that a government, by increasing taxes such that net income is kept fixed, could finance pension payments to an increasing number of pensioners.

3.4 Individual's Optimization Problem

Individuals choose consumption and portfolio allocation by maximizing expected lifetime utility:

$$\mathbb{E}\sum_{n=N^b}^{N^d}\left(\prod_{l=N^b}^{n-1}\delta_{l,t+l-N^b}\right)\beta^{n-N^b}u(c_{i,n,t+n-N^b}), \tag{3.7}$$

where $c_{i,n,t}$ is consumption of individual i who is n years old at time t. $\beta\delta_{n,t}$ is the subjective discount factor that takes the region-specific stochastic survival probability into account. Period utility is of the CRRA type, $u(c) = \frac{c^{1-\vartheta}-1}{1-\vartheta}$.

Each working individual faces the flow budget constraint,

$$c_{i,n,t} + Q_t b_{i,n+1,t+1} + \tilde{Q}_t s_{i,n+1,t+1} = \overbrace{(Q_t + d_t)b_{i,n,t} + (\tilde{Q}_t + \tilde{d}_t)s_{i,n,t} + y_{i,n,t}}^{x_{i,n,t}} \quad \forall\ n < N^r. \tag{3.8}$$

[9]On the role of medical expenses of the elderly, see De Nardi, French, and Jones (2010), De Nardi, French, Jones, and McCauley (2016), and Ameriks, Briggs, Caplin, Shapiro, and Tonetti (2015).

[10]There are several different pension systems around the world, among them fully funded and pay-as-you-go (PAYGo) systems. Imposing a PAYGo system would require us to specify how the system remains solvent in the presence of demographic change: Through an increase in the contribution rate, a decrease in the replacement rate, or through an increase in the retirement age. All three of these influence the amount of life-cycle saving and would add additional dimensions through which demographics influence aggregate asset demand.

In the budget constraint for retired individuals, public pension income replaces labor income,

$$c_{i,n,t} + Q_t b_{i,n+1,t+1} + \tilde{Q}_t s_{i,n+1,t+1} = \overbrace{(Q_t + d_t) b_{i,n,t} + (\tilde{Q}_t + \tilde{d}_t) s_{i,n,t} + \tilde{y}_{i,n,t}}^{x_{i,n,t}} \quad \forall \; n \geq N^r. \tag{3.9}$$

$x_{i,n,t}$ is the state variable in our model and characterizes an agent's total available resources (cash on hand) in each period. $b_{i,n+1,t+1}$ and $s_{i,n+1,t+1}$ are bond and stock holdings carried over into period $t+1$ by a household of age n in period t.

Due to the zero income shock, a natural borrowing constraint arises, because agents want to ensure positive consumption at all times and states of the world:

$$Q_t b_{i,n,t+1} + \tilde{Q}_t s_{i,n,t+1} \geq 0 \qquad \forall \, i, n, t \tag{3.10}$$

Additionally, we impose the stronger constraint that agents are not allowed to short stocks or bonds,

$$b_{i,n,t+1} \geq 0 \qquad s_{i,n,t+1} \geq 0 \qquad \forall \, i, n, t \tag{3.11}$$

This has been suggested by Constantinides, Donaldson, and Mehra (2002) as a way to generate equity risk premia closer to the data.

The timing of the model is as follows: At the beginning of each period, agents receive the returns on their assets bought in the previous period. Afterwards, dying agents exit the economy and a new cohort enters.[11] Newborn agents do not own any financial wealth. The number of trees adjusts according to the population. Next, all living agents receive labor or pension income. Given their total resources, agents decide on consumption and savings and on their portfolio such that they maximize Eq. (3.7) with respect to Eqs. (3.8), (3.9) and (3.11), which will generate their financial wealth in the following period. They make this decision forming expectations over future return realizations and income shocks and take into account their age- and time-specific survival probability.

Savings and Portfolio Decisions

For simplicity, we drop individual and cohort subscripts and define end of period assets as $a_{t+1} \equiv Q_t b_{t+1} + \tilde{Q}_t s_{t+1}$ and the risky asset share as $\omega_{t+1} = \frac{\tilde{Q}_t s_{t+1}}{a_{t+1}}$. The first order condition with respect to consumption yields the following Euler equation

$$c_t^{-\vartheta} = \beta \delta_{t+1} \mathbb{E}\left[\left(\omega_{t+1} \tilde{R}_{t+1} + (1 - \omega_{t+1}) R_{t+1} \right) c_{t+1}^{-\vartheta} \right] \tag{3.12}$$

Individuals will consume optimally such that marginal utility gain of an increase in current consumption is offset by the expected loss in future consumption times the portfolio return.

The portfolio share optimality condition is

$$\beta \delta_{t+1} \mathbb{E}\left[c_{t+1}^{-\vartheta} a_{t+1} \left(\tilde{R}_{t+1} - R_{t+1} \right) \right] = \mu_b - \mu_s \tag{3.13}$$

where μ_b and μ_s are the Lagrangian multipliers on the no-leverage constraints spelled out in Eq. (3.11), which satisfy $(1 - \omega_{t+1}) \mu_b = 0$ and $\omega_{t+1} \mu_s = 0$.

[11] For simplicity, we abstract from accidental bequests and assume that the financial wealth of dying agents simply expires. Alternatively, it could be relocated to the individuals alive.

For an interior solution, the individual optimally allocates the share ω_{t+1} to risky assets such that the expected value of future marginal utility of consumption times the expected excess return equals zero. At an optimum, the individual should not be better off by rebalancing her portfolio towards either asset.

There is no closed form solution to the optimization problem. With individuals taking interest rates as given, we solve for age-dependent policy functions for the endogenous state variable cash on hand x_t, and the future returns R_{t+1} and $\tilde{R}_{t+1}$:

$$\begin{aligned}\hat{c}_{t,n,i} &= c(x_{t,n,i}, R_{t+1}, \tilde{R}_{t+1})\\ \hat{\omega}_{t+1,n+1,i} &= \omega(x_{t,n,i}, R_{t+1}, \tilde{R}_{t+1}).\end{aligned}$$

3.5 Aggregates and the Open Economy

So far, we have presented the individual's problem and market structure separately for the two regions we study. We now consider aggregation and market-clearing conditions for the time for which the economies are closed and after full financial integration.

Aggregate demand for goods, bonds and stocks in each region $j = e, u$ at time t is the sum over cohorts weighted by the cohort size:

$$C_t^j = \sum_{n=N^b}^{N^d} L_{n,t}^j c_{n,t}^j \qquad B_{t+1}^j = \sum_{n=N^b}^{N^d} L_{n,t}^j b_{n,t+1}^j \qquad S_{t+1}^j = \sum_{n=N^b}^{N^d} L_{n,t}^j s_{n,t+1}^j, \tag{3.14}$$

where in each region, individual consumption $c_{i,n,t}^j$, stock holdings, $s_{i,n,t+1}^j$ and bond holdings $b_{i,n,t+1}^j$ are averaged within cohorts, $c_{n,t}^j = \frac{1}{L_{n,t}^j}\int_0^{L_{n,t}^j} c_{i,n,t}^j \mathrm{di}$; $s_{n,t+1}^j = \frac{1}{L_{n,t}^j}\int_0^{L_{n,t}^j} s_{i,n,t+1}^j \mathrm{di}$; and $b_{n,t+1}^j = \frac{1}{L_{n,t}^j}\int_0^{L_{n,t}^j} b_{i,n,t+1}^j \mathrm{di}$.

Total financial assets demanded at time t to be carried over into the next period are $Q_t^j B_{t+1}^j + \tilde{Q}_t^j S_{t+1}^j = A_{t+1}^j$ and the aggregate share of risky assets $\Omega_{t+1}^j = \frac{\tilde{Q}_t^j S_{t+1}^j}{A_{t+1}^j}$.

Given the assumption of an endowment economy where the number of safe and risky Lucas trees is identical to the population size, the aggregate supply of stocks $\mathbb{S}_{t+1}^j$ and bonds $\mathbb{B}_{t+1}^j$ is proportional to the sum of all living cohorts:

$$\mathbb{B}_{t+1}^j = \sum_{n=N^b}^{N^d} \lambda L_{n,t}^j \qquad \mathbb{S}_{t+1}^j = \sum_{n=N^b}^{N^d} \tilde{\lambda} L_{n,t}^j, \tag{3.15}$$

where λ, $\tilde{\lambda}$ are scaling parameters for safe and risky assets, respectively. They are identical for both regions and can be interpreted as per capita asset supply of safe and risky assets. This assumption for asset supply is consistent with recent evidence that, while the total supply of financial assets in the economy is growing over time, the share that can be considered safe is roughly constant (Gorton, Lewellen, and Metrick, 2012).

The two regions we consider remain closed until period T. The countries are in autarky and agents can buy only domestic stocks and bonds and consume domestic goods. From $T+1$ onwards, financial asset and goods markets of e and u are perfectly integrated, so that agents of both regions are free to buy assets and consumption goods anywhere in the world.

Autarky

Until period T, local demand for assets is cleared by local supply at equilibrium asset prices in each region,

$$S^j_{t+1} = \mathbb{S}^j_{t+1} \qquad\qquad B^j_{t+1} = \mathbb{B}^j_{t+1}, \tag{3.16}$$

where stock and bond demands in each region are given in Eqs. (3.14) and (3.15).

Aggregating over cohorts, respectively, using Eq. (3.14) yields the aggregate resource constraint under autarky:

$$C_t + Q_t(B_{t+1} - B_t) + \tilde{Q}_t(S_{t+1} - S_t) = \overbrace{Y_t + d_t B_t + \tilde{d}_t S_t}^{\text{Aggregate Endowment}}. \tag{3.17}$$

where Y_t is the aggregate labor and pension income in the economy.

If the market clearing conditions under autarky, Eqs. (3.16) hold, the market for goods clears by Walras's Law and Eq. (3.17) becomes,

$$C_t = \overbrace{Y_t + \lambda d_t L_{t-1} + \tilde{\lambda}\tilde{d}_t L_{t-1}}^{\text{Aggregate Endowment}} - Q_t\lambda(L_t - L_{t-1}) - \tilde{Q}_t\tilde{\lambda}(L_t - L_{t-1}). \tag{3.18}$$

Asset prices under autarky, Q_t and $\tilde{Q}_t$, will be such that aggregate consumption will be equal to the aggregate endowment minus the additional savings that offset the change in the supply of assets due to population change.

Two-Region World

In a world consisting of two perfectly integrated regions $j = e, u$, the asset prices will adjust such that the markets for assets and goods clear globally. The relevant market-clearing conditions for stocks and bonds are:

$$\begin{aligned} S_{t+1} &\equiv S^e_{t+1} + S^u_{t+1} = \mathbb{S}_{t+1} \equiv \mathbb{S}^e_{t+1} + \mathbb{S}^u_{t+1} \\ B_{t+1} &\equiv B^e_{t+1} + B^u_{t+1} = \mathbb{B}_{t+1} \equiv \mathbb{B}^e_{t+1} + \mathbb{B}^u_{t+1}. \end{aligned} \tag{3.19}$$

Due to our assumption of perfectly correlated returns across regions, agents of both regions are indifferent between holding domestic and foreign assets. Total asset demand in region u will consist of local assets and cross-border asset holdings from region e. A region's net foreign asset position in each instrument (NFB_t and NFS_t) is the difference between local asset demand and local asset supply,

$$\begin{aligned} NFB^u_t &= Q_t\left(\overbrace{B^{uu}_{t+1} + B^{ue}_{t+1}}^{\text{bond demand by u}} - \overbrace{B^{uu}_{t+1} - B^{eu}_{t+1}}^{\text{bond supply by u}}\right) = Q_t\,(B^{ue}_{t+1} - B^{eu}_{t+1}) \\ NFS^u_t &= \tilde{Q}_t\left(\overbrace{S^{uu}_{t+1} + S^{ue}_{t+1}}^{\text{bond demand by u}} - \overbrace{S^{uu}_{t+1} - S^{eu}_{t+1}}^{\text{bond supply by u}}\right) = \tilde{Q}_t\,(S^{ue}_{t+1} - S^{eu}_{t+1}) \end{aligned} \tag{3.20}$$

where, for example, B^{ue}_{t+1} are claims by agents in region u on region e's safe trees. Region e's net foreign asset positions are computed analogously, such that $NFB^u_t = -NFB^e_t$ and

$NFS_t^u = -NFS_t^e.$

3.6 Equilibrium

For each period t, an equilibrium consists of a set of asset returns, R_t and $\tilde{R}_t$; age-dependent policy functions for consumption in each region $\hat{c}_{n,t}^j$, portfolio allocation, $\hat{\omega}_{n,t}^j$ and cash on hand $\hat{x}_{n,t}^j$; aggregate demand for stocks S_t^j and bonds B_t^j, aggregate consumption C_t^j and for post financial opening, net external positions for each region, NFB_t^j and NFS_t^j such that:

1. The asset returns are consistent with Eq. (3.3).
2. For given interest rates, the household's problem solves with policy functions, $\hat{c}_{n,t}^j$ and $\hat{\omega}_{n,t}^j$ for all ages n and in both regions j.
3. All markets clear: Prior to opening, the aggregate stock and bond positions given in Eq. (3.14) satisfy Eq. (3.16) and the aggregate budget constraint Eq. (3.17); after opening, the aggregate world stock and bond positions satisfy Eq. (3.19).
4. The aggregate behavior in the economy is consistent with the necessary conditions for all ages in both regions as given by Eqs. (3.12) and (3.13).

The solution algorithm that we use to solve the individual's problem and the general equilibrium is described in detail in Appendix B.

4. Calibration

In order to test our model's ability to explain international external positions, we perform an annual simulation of the demographic transition from 1950 to 2095. Regions are fully financially integrated starting in 1990. We generate pre- and post-demographic transition distributions of the population and asset holdings by solving the model for demographic variables in 1950 and 2095, respectively. Starting in 1950 the demographic transition begins and agents born in 1950 take into account the full demographic transition directly through survival probabilities and indirectly through the aggregate population structure affecting returns.

The solution and simulation of our large-scale model is not trivial. For each region, we simulate 146 transition years, plus 80 cohorts each for the pre- and post-transition steady states. In each simulation year, we must compute the cohort-specific consumption and portfolio allocation for given values of the endogenous state and expectations of future returns. Since returns in the subsequent period are a function of the entire distribution of asset holdings, this requires the current and expected future aggregate distribution of assets to enter as a state variable. We currently proceed by making two simplifying assumptions. While facilitating the simulation, they still allow us to highlight the demographic mechanisms in our results.

First, agents are myopic with respect to future returns. In their optimization, they set their expectations of next period returns equal to current equilibrium returns,

$$\mathbb{E}[R_{t+1}] = R_t \qquad \mathbb{E}[\tilde{R}_{t+1}] = \tilde{R}_t.$$

A similar simplifying assumption is made by Backus et al. (2014), who assume complete myopia of agents in terms of demographic variables and factor prices determined from aggregate distributions. This assumption simplifies the solution significantly because we avoid calculating current consumption and portfolio allocation for a large number of paths for interest rates. However, it comes at the cost of not being able to fully quantify the effect of demographic change on individual behavior that works through changing expectations of returns. For example, if agents can perfectly anticipate the future path of returns, they might rebalance away from the asset with lower equilibrium returns and, on the other hand, may adjust their level of savings to offset diminished asset returns.

Second, we run a deterministic simulation, where we set all realizations of the shocks to their expected values. Abstracting from idiosyncratic shocks makes agents within each cohort homogeneous. The simplification is not very restrictive in this case, given that we assume a continuum of individuals within each cohort, so that the law of large numbers implies average shock realizations to equal expected values. For the aggregate return rate shock, a deterministic simulation has potentially strong implications. As the composition of portfolios is age-dependent, large shocks to the risky asset returns could shift the distribution of asset holdings across age groups, which, in turn, determines equilibrium returns. However, stochastic simulations will introduce sampling noise that makes the policy functions dependent on a specific realization of the aggregate shock and such realizations would be concentrated around the mean, so that accuracy for a large realization of the aggregate shock would be low (Algan, Allais, and Haan, 2008). Our approach generates a mean path for aggregate asset demands as they respond to demographic change, which is the central focus of this paper.

The full calibration of our model is shown in Table 2.1.

Table 2.1: Parameters in simulation

Age at birth	N^b	20	(Cocco et al., 2005)
Maximum age	N^d	100	(Cocco et al., 2005)
Retirement age	N^r	65	Legal retirement age, United States
Subjective discount factor	β	0.96	(Cocco et al., 2005)
CRRA parameter	ϑ	8	(Mehra and Prescott, 1985)
Variance of shock to risk premium	σ_ϵ^2	0.025	(Fama and French, 2002)
Variance of transitory income shock	σ_θ^2	0.07	(Cocco et al., 2005)
Variance of permanent income shock	σ_η^2	0.01	(Cocco et al., 2005)
Probability of zero income	p	0.005	(Carroll, 1992)
Age-dependent income growth	G_n, P_n		(Cocco et al., 2005)
Pension income replacement rate	ϕ	0.68	(Cocco et al., 2005)
Survival probabilities	$\delta_{n,t}$		UN Population Prospects, medium variant, 1950 to 2095
Asset prices	Q_t, $\tilde{Q}_t$	1	normalization
Safe asset supply parameter	λ	160.6	Calibrated as in Section 4
Risky asset supply parameter	$\tilde{\lambda}$	83.2	Calibrated as in Section 4

Demographics

In accordance with our model assumption that agents are born at working age, we use data starting from age 20. Over the life-cycle agents survive with age-dependent survival probabilities, for which we use data and projections from the United Nations World Population Prospects (United Nations, 2015). Section A.1 provides further information on the construction of the data series. While we use the age- and time-specific survival probabilities in the individual optimization problem, we simulate the aggregate population using data on actual cohort sizes instead of using a combination of birth rates and survival probabilities. This has the advantage that we can take into account migration, and that defining 1950 artificially as a steady state does not skew the population distribution in subsequent simulation years (see Section A.7 for a discussion).

We impose a fixed retirement age of 65, which roughly matches the average effective retirement age in the US and the EU since the 1970s (OECD, 2015). There have been only slight changes over the last decades: While the average effective retirement age was 68 in 1970 in both the US and the EU, it decreased to 64 in the US and 61 in the EU during the 1990s and has been slightly rising to 65 (US) and 63 (EU) since. The OECD in its projections until 2060 does not expect major increases in the retirement age, which is why the assumption of a (roughly) constant retirement age over the whole simulation period does not seem unreasonable.

Labor and Pension Income

For the deterministic component of labor income G_n in Eq. (3.5), we use estimates by Cocco, Gomes, and Maenhout (2005) based on the US Panel Study of Income Dynamics (PSID). We use their results for the group with a medium level of education, high school graduates without college degree,

$$ln(G_n) = 0.6004 + 0.1682 \times age - 0.0323 \times \frac{age^2}{10} + 0.0020 \times \frac{age^3}{100}.$$

The variance of the transitory and permanent income shocks θ and η are also taken from Cocco et al. (2005). The probability of a zero income shock, p, is set to 0.5 percent as estimated by Carroll (1992).

For the replacement rate, we also use the value estimated by Cocco et al. (2005), $\phi = 0.68$. This equals the public pension income that individuals receive on average. Additionally, private pension schemes play a role, particularly in the US, but these are incorporated in the savings and portfolio choice problem of the agents.

Asset Prices

We normalize asset prices to 1 so that $Q_t = \tilde{Q}_t \quad \forall \; t$. Given the definitions for asset returns in Eq. (3.3), this assumption implies that safe asset returns are $R_t = 1 + d_t$ and risky returns are $\tilde{R}_t = 1 + d_t + rp_t + \epsilon_t$. We solve for the time-varying d_t and rp_t that clear asset markets. Assuming constant asset prices abstracts from capital gains in asset returns. Thus, changes in the valuation of assets work entirely through changes in the dividends d_t and in the risk premium rp_t. In consequence, we are not able to separate the effect of demographic change on capital gains and dividend yields, but since we are interested in the impact of demographic change on real asset returns, this is not of first order importance.

Parameters on Asset Supply

Until the 1990s, cross-border financial flows between the US and the EU are rather small (especially by today's standards). Therefore, we use realized US returns during the 1980s to calibrate λ and $\tilde{\lambda}$ in a partial equilibrium set-up for the US as a closed economy. We take real returns to a US ten-year government bond and historical risk premia estimated for the S&P 500 from Damodaran (2016) as market-clearing returns in our model economy. We then solve for asset demands in the economy at these rates, which are the aggregate of optimal individual household decisions. In the closed economy, these aggregate demands have to equal aggregate supply, so we obtain each period's λ and $\tilde{\lambda}$ by dividing by total US model population. We use a simple average of these values to represent the time-invariant per capita asset supply from 1990 onward,

$$\lambda = \frac{1}{10}\sum_{t=1980}^{1989}\frac{B_t}{L_t} \qquad \tilde{\lambda} = \frac{1}{10}\sum_{t=1980}^{1989}\frac{S_t}{L_t}$$

where B_t and S_t are asset demands in the US at observed interest rates and L_t is the US population. Our model results and calibrated supply parameters are very similar when using alternative specifications, namely different scaling parameters for the US and the EU, or the last scaling parameters that clear the markets in 1989.

5. Results

The following section contains the simulation results. We first present the results for equilibrium rates of return and external positions. After the baseline simulation, we introduce a version where we account for more gradual financial integration and home bias in equity and debt. Then, we discuss the various channels through which demographic change affects the demand for safe and risky assets. We provide intuition for the effect on life-cycle financial decisions of individuals, before more systematically decomposing the aggregate results.

5.1 Effect of Demographics on Aggregate Outcomes

Rates of Return

The simulated equilibrium return to safe assets and the risk premium are shown in Figure 2.5. We compare them to the same data series that were used to calibrate the supply side: real US 10-year government bond returns and risk premia estimated by Damodaran (2016).[12] Due to our calibration for asset supply, the level of the returns in the data and in the model are similar in 1990 before demographic change becomes more pronounced. For the risk premium, due to an increase during the 1980s, the average used for the calibration is slightly lower than the 1990 data value. In the years following, we predict a downward trend in safe and risky

[12] As can be seen from Figure 2.18, real bond returns for the largest European economies are very similar to the US rates, so that we can take US rates to represent the integrated world returns.

Figure 2.5: Rates of return

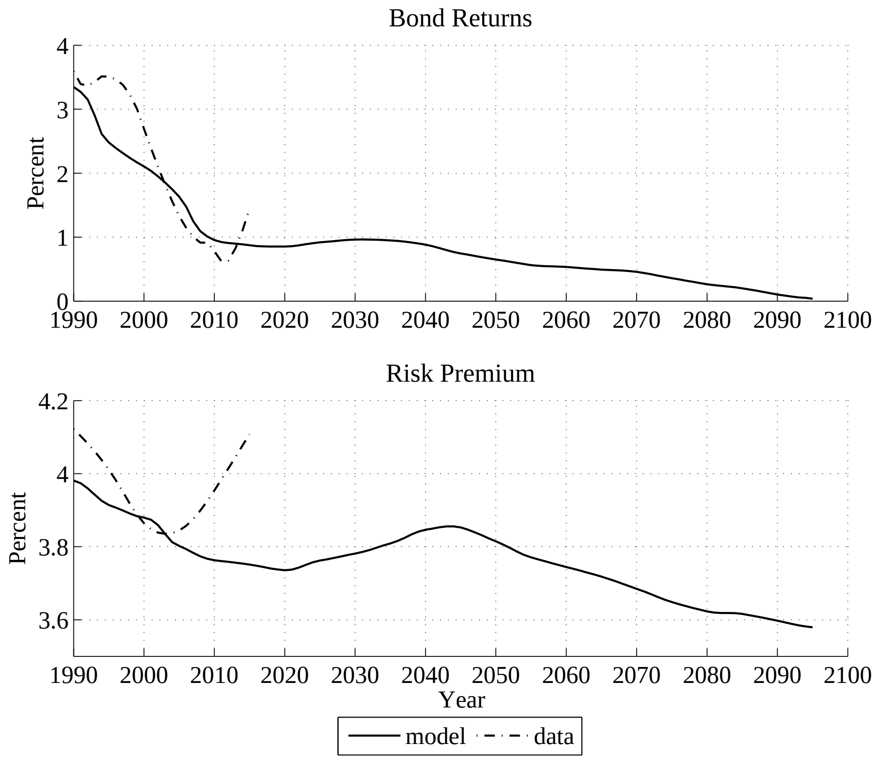

The upper panel shows model predictions for safe asset returns in the solid line and data for real returns to 10-year US treasury bills in the dashed line. The lower panel shows the model predictions for risk premia and data for historical risk premia. Data are from Damodaran (2016).

returns, which is consistent with the data. From 1990 to 2015, world demographics induce a decline in safe asset returns of about 200 basis points, which is almost exactly the decline in real US returns observed over the same period in the data. This number is slightly larger than what has been found by Carvalho et al. (2016) and Gagnon et al. (2016). Concerning the risk premium, we predict a continued downward trend during 1990 to 2015, which comes close to the data until around the year 2007. Afterwards, the risk premium increases in the data, which can likely be explained with the start of the global financial crisis. As this should be a transitory phenomenon, and one that is not accounted for in our model, we would expect the risk premium to eventually decrease again. In general, the model rationalizes the rather limited movements in the risk premium over the last years.

Demographic change implies a continued downward trend in safe asset returns until the end of the century. This reflects an ever increasing asset demand in combination with fixed per capita asset supply. The movement is however not monotonic. Between 2020 and 2040, the model predicts safe returns to increase, which is due to the baby boomer cohorts: As retirees, the baby boomers will run down their assets, leading to a temporary decline in aggregate asset demand. Once they have exited the model, the effect is gone.

The dynamics of the risk premium mirror those of the risk-free rate, so that the risky return always reacts more strongly to demographic change than the safe return. This is a general result from our type of model. As shown analytically by Kuhle (2017), in a portfolio choice

Figure 2.6: External positions

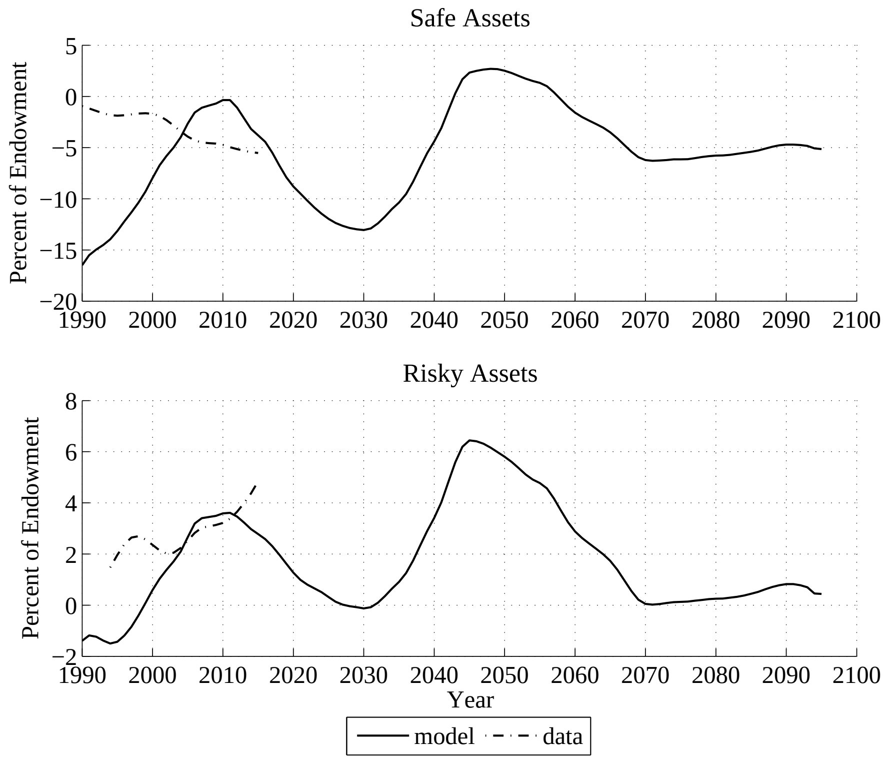

The upper panel shows the model predictions for US external safe asset position vis-à-vis the EU as percent of aggregate endowment in the solid line from 1990 to 2095. Data on safe bilateral positions as percent of GDP are constructed as in Figure 2.1 is shown in the dashed line. The lower panel shows model predictions for bilateral risky assets positions as percent of aggregate endowment and data on risky bilateral positions as percent of GDP.

problem with concave utility, agents are more sensitive to changes in the safe rate than in the risky rate when choosing their portfolio, provided that the risk premium is positive. Thus, an increase in both rates of return which leaves the equity premium unchanged will result in a lower portfolio share of the risky asset. In our model environment where the aggregate world risky share is required to stay constant due to the supply side specification, this implies that a demographics-induced decline in rates of return leads to a more pronounced decrease in risky rates in order for markets to clear, and the risk premium hence declines. This negative relation also results from the models by Brooks (2002) and Geanakoplos et al. (2004). Additionally, from the middle of the century onwards, the US will overtake the EU in terms of population size. Given that the US ages more slowly, this increases the relative size of younger cohorts, who prefer to hold a larger share of risky assets in their portfolio.

External Positions

Our model predicts sizable net asset trades between the two regions after financial opening in 1990. Figure 2.6 depicts the predicted net external positions in safe and risky assets as a ratio to the aggregate endowment in each simulation year. The results are compared to smoothed data on bilateral safe and risky assets shown in Section 2, Figure 2.1.

Figure 2.7: Relative old-age dependency ratios and net external positions, 1990-2095

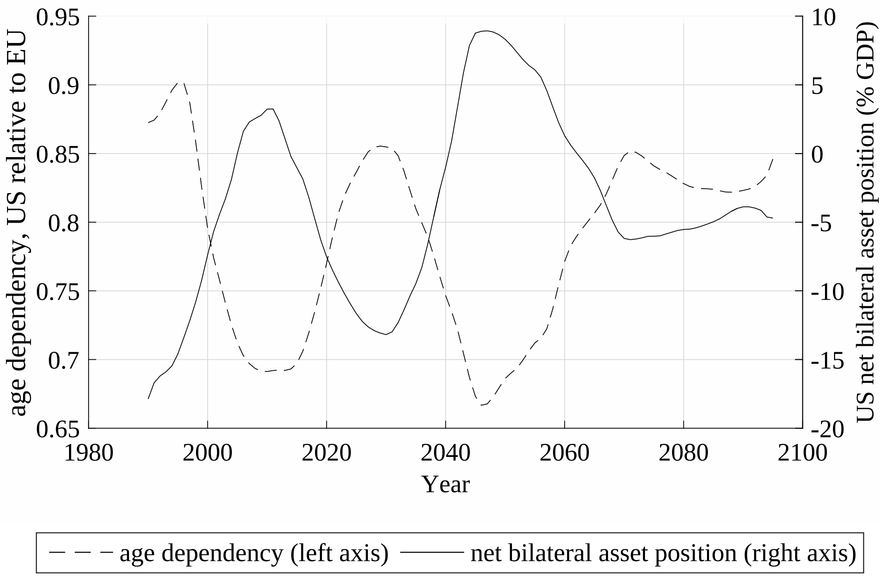

The figure shows the old-age dependency ratio of the US relative to the old-age dependency ratio of the EU in the dashed line and the model predictions for the US net bilateral asset position as percent of aggregate endowment in the solid line. There is a negative relationship between the relative DR and NFA position, with a correlation coefficient of -0.77.

The simple model cannot perfectly replicate all the nuances of the net US external asset position in the data. However, it does replicate the most relevant features. Consider first the direction of external positions. The model correctly predicts a bilateral liability position in safe assets and an asset position in risky assets for the US. Thus, the fact that the US population is younger than the EU population can explain the risk asymmetry in the bilateral positions. For the period 1994 to 2015, the model predicts an average of US risky assets of 1.7 percent of aggregate endowment, whereas in the data, the value is 2.8 percent of GDP. For safe assets, the model predicts a US liability position of 5.4 percent of aggregate endowment, compared to 3.5 percent in the data. Thus, on average, the model predictions come close to the data. The model fit is better from approximately the year 2000 onwards. For the period 1990 to 2000, the simulated positions in both types of assets are more extreme than those in the data, and in particular the large negative debt liabilities do not match the data. This may be linked to the fact that we assume full and immediate integration in 1990. While we explore the effect of a more gradual integration process further below, we also want to point out that until 1994, the safe position shown excludes long-term debt instruments, and in reality could thus be more negative. Overall, demographics can explain the observed external positions reasonably well.

For the period 2016 to 2095, the model predicts that the risk asymmetry in the bilateral asset positions will continue: During most of the projection years, the US is an importer of risky assets and an exporter of safe assets. The pronounced oscillations in the magnitude of positions are closely linked to relative demographics, as shown in Figure 2.7: Whenever the aging trend in the US slows relative to the EU, reflected by a lower relative dependency ratio, the US accumulates relatively more assets. This is because a lower old-age dependency ratio corresponds to a higher share of savers in the population. When a region gets younger, the domestic demand for assets rises faster than supply. Therefore, at rates that clear international markets, the country wants to save more externally. The difference in demographics is most extreme in 2040 to 2050, because the baby boom took place somewhat earlier in the US than in the EU: By 2040, many US baby boomers will have died and the US population be on average rejuvenated, whereas some EU baby boomers will still be alive. Dependency ratios become more similar towards the end of the century in the transition into the new steady state, at which point the US net bilateral position is only slightly negative. Still, the model predicts that the asymmetry in external positions between the US and the EU is not a transitory phenomenon, but will continue to be sizable over the next decades.

External Positions Under Gradual Integration

Although formal barriers to asset trades had to a large extent been removed by 1990, the assumption that capital can flow freely across borders since 1990 is a strong one: On the one hand, some formal restrictions on capital accounts in the EU remained in force until the 2000s. Indicative of this is the Chinn and Ito (2008) Index of de jure capital account openness, which for the EU has only converged to the full integration value of 1 by 2002 (Appendix Figure 2.17). On the other hand, informal barriers further limited asset trades. In particular information frictions, which are sometimes brought up as an explanation for observed home bias, may still have been pronounced in 1990, but decreased during the following years. Sørensen, Wu, Yosha, and Zhu (2007) estimate that between 1993 and 2003, the equity home bias in the EU decreased from 0.81 to 0.62 and the debt home bias decreased from 0.59 to 0.43.[13] In the US, the equity home bias decreased from 0.84 to 0.74 and the debt home bias from 0.78 to 0.73. In our model, agents always perceive home and foreign assets of the same type as identical, and therefore in the baseline simulation we are not able to capture the effect of these frictions.

In this section, we assume that financial opening took place gradually, from no integration in 1990 to full integration by 2000. We model this through restrictions on the degree of equalization in the rates of return. However, the assumptions are ad-hoc and therefore, the purpose of this simulation is not to obtain a fully realistic prediction. Rather, it is meant to serve as a benchmark against which we can assess how barriers to asset trades may affect the model outcome.

We assume that until 1990, autarky returns prevail, such that markets clear separately in the US and the EU. Over the next 10 years, rates gradually converge to the open economy market clearing rates so that by 2000, full integration is attained. Let ρ_t be a parameter linearly increasing from 0 in 1990 to 1 in 2000:

$$\rho_t = \frac{\text{simulation year} - 1990}{10} \qquad \text{for} \quad \text{simulation year} \in [1990, 2000]$$

Safe rates are defined as $(1-\rho_t)R_t^{aut} + \rho_t R_t^{int}$ and risky rates are $(1-\rho_t)\tilde{R}_t^{aut} + \rho_t \tilde{R}_t^{int}$ where the superscript *aut* denotes autarky rates and *int* full integration rates. We reproduce the aggregate results for these new rates.

[13]The numbers exclude Luxembourg (debt and euqity) and Ireland (equity).

Figure 2.8: External positions under gradual integration

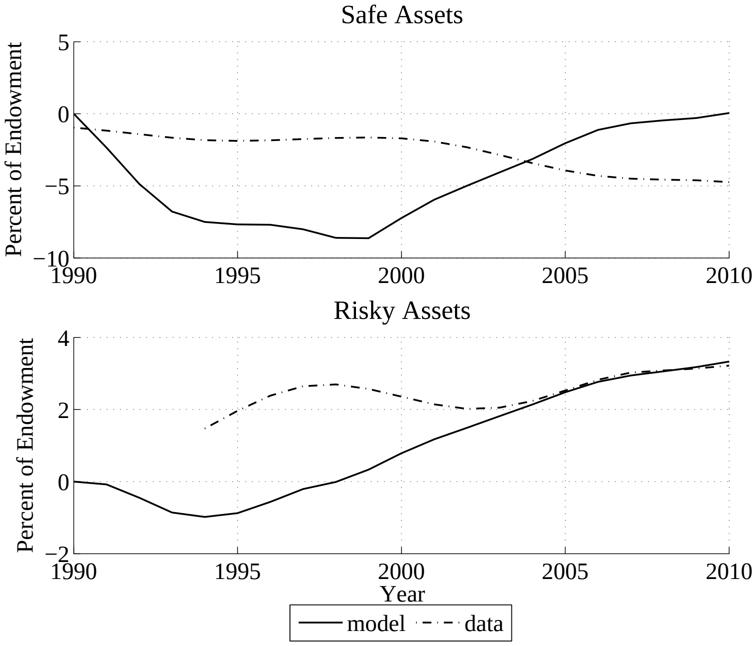

The upper panel shows the model predictions for the US external safe asset position vis-à-vis the EU as percent of aggregate endowment under gradual integration (solid line). The dashed line is the same data series as in Figure 2.6. The lower panel shows model predictions for the US external risky asset positions vis-à-vis the EU under gradual integration as percent of aggregate endowment (solid line) and the corresponding data series (dashed line).

Figure 2.8 shows the results. Gradual integration brings the model predictions closer to the data. In particular, the large negative safe position that our baseline model predicts for the US is now smaller. So it seems that the large initial positions in our baseline version are partly driven by our assumption on openness.

5.2 Effect of Demographics on Savings and Portfolio Choice

We have shown that demography has strong effects on external positions and equilibrium rates of return. Now we want to take a closer look at how demographic change impacts savings decisions and portfolio choice, both at the individual level and on aggregate. Figure 2.9 shows model predictions for a cross-section of agents in the US in 2010 at market-clearing rates of return. The left panel shows age-specific financial asset holdings, non-financial income consisting of labor and pension income, and consumption. The right panel shows the risky asset share. Total savings are hump-shaped over the cross-section: they increase with age until shortly before retirement and decrease afterwards. This savings pattern reflects a desire of agents to smooth consumption over the high labor income and the low pension income period of life. The risky share increases slightly at the very beginning of economic life, and then decreases continually. This is in line with the estimates presented in Figure 2.3 of Section 2.[14]

[14]The model-predicted risky share is initially too high, but generating realistic portfolio shares is generally difficult in this type of models (Merton, 1971; Heaton and Lucas, 1997; Guiso et al., 2002; Cocco et al., 2005;

Figure 2.9: Life-cycle results

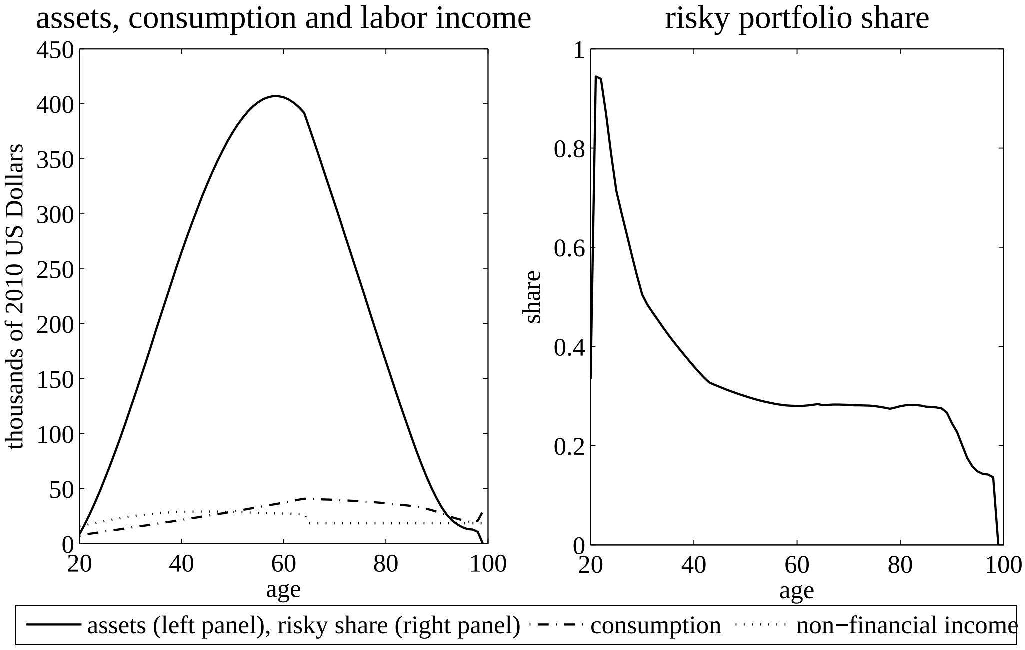

The left panel shows model-predicted financial assets, consumption and non-financial income for a cross-section of US individuals in 2010. The right panel shows the model-predicted risky share of financial assets.

The life-cycle profiles in asset holdings and portfolio choice create a strong link between aggregate asset demand and the age structure of the population. The higher the relative share of older individuals, the lower is the desired aggregate risky share. A larger population mass near or at retirement age will increase the overall demand for assets, whereas a shift towards more dissavers (retirees) will overturn this effect. As shown in the statistical exercise in Section 2, the population aging predicted for the US and the EU will work to increase desired aggregate savings between 1950 and 2095, and decrease the desired risky share of assets held.

This effect working through the relative population shares of age groups is one out of three channels through which demography impacts the aggregate demand for safe and risky assets. We call it the *distribution effect*. In the following, we describe the two other channels at work: The *life-cycle effect* and the *valuation effect*. To conclude, we show how important each channel is, both in the cross-section and over time.

Life-cycle Effect

The *life-cycle effect* comprises the direct effect of longevity on the optimal level of individual savings and the risky portfolio share. An increase in individual survival probabilities affects how much agents optimally want to save at different points of their lives, and what type of assets they prefer to hold. Figure 2.10 compares the total financial asset holdings and the portfolio share of risky assets of two individuals living in the US: one in the pre-demographic

Fagereng et al., 2017). However, our high CRRA parameter in combination with zero income shocks allows us to get closer to the data than models with more standard parameters.

Figure 2.10: Assets and risky share over the life-cycle: 1950 vs. 2095

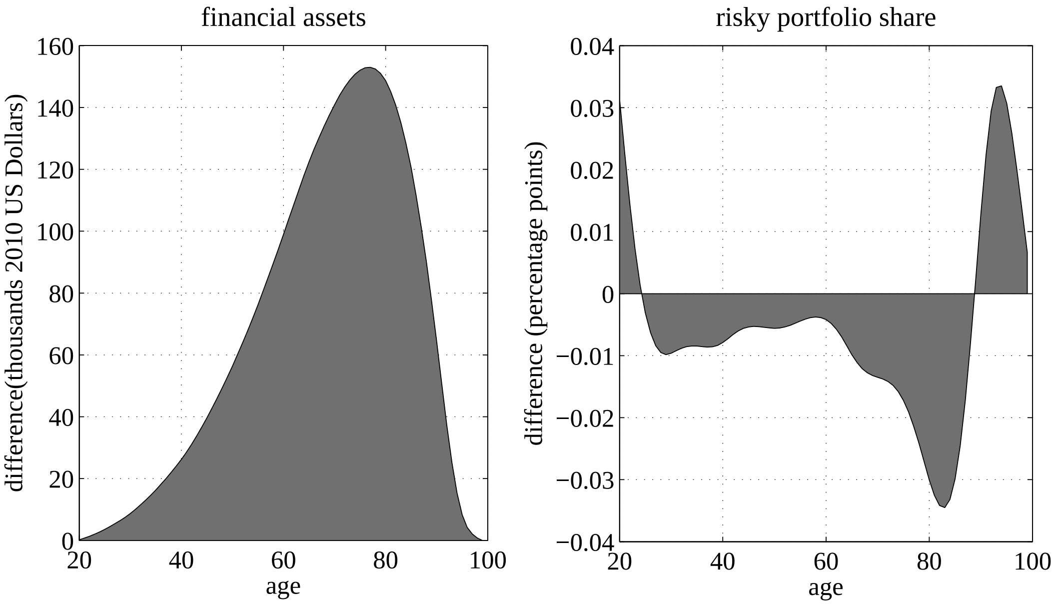

The left panel shows the difference in the life-cycle financial asset holdings in 1950 compared to 2095 at $R = 1.02$ and $rp = 0.04$. The right panel shows the difference in the life-cycle risky share between 1950 and 2095 at the same returns.

change steady state, the other in the post-demographic change steady state. The graphs shows the difference at each age between the two simulation years. We keep the rates of return fixed: the safe return is 1.02 for both simulation years, and the risk premium is 0.04.

Between 1950 and 2095, we observe a strong increase of savings over the life-cycle. Absent adjustments in the rates of return, an increased expected longevity has the effect of making individuals more patient. This is because changes in survival probabilities prolong the share of life that individuals expect to spend in retirement. The difference in asset holdings is particularly strong around the age 80, because the difference in the survival probabilities is the largest here (partly due to the fact that we require survival probabilities to converge to 0 by the age 100). The effect of longevity on intertemporal decisions can also be seen from the Euler equation, Eq. (3.12).

The effect of demographics on portfolio choice over the life-cycle is less uniform. Increases in longevity affect the risky portfolio share of the very young and the very old positively, and the portfolio share of all other age groups negatively. However, the differences are small in absolute terms. The effect works through a combination of changes in savings and changes in human capital, the net present value of future labor and pension income.

Denote human capital of an individual of age n by

$$h_{i,n,t} = \sum_{m=n}^{N^d} \left(\prod_{l=n}^{m} \delta_{l,t-n+l} \frac{\mathbb{E}\left[\hat{y}_{i,m,t+m}\right]}{R^{m-n}} \right) \quad \text{where} \quad \hat{y} = \begin{cases} y & \text{if} \quad m \leq N^r \\ \tilde{y} & \text{if} \quad m > N^r \end{cases} \tag{3.21}$$

Despite being subject to shocks, human capital still resembles a safe asset rather than a

Figure 2.11: Human capital over the life-cycle

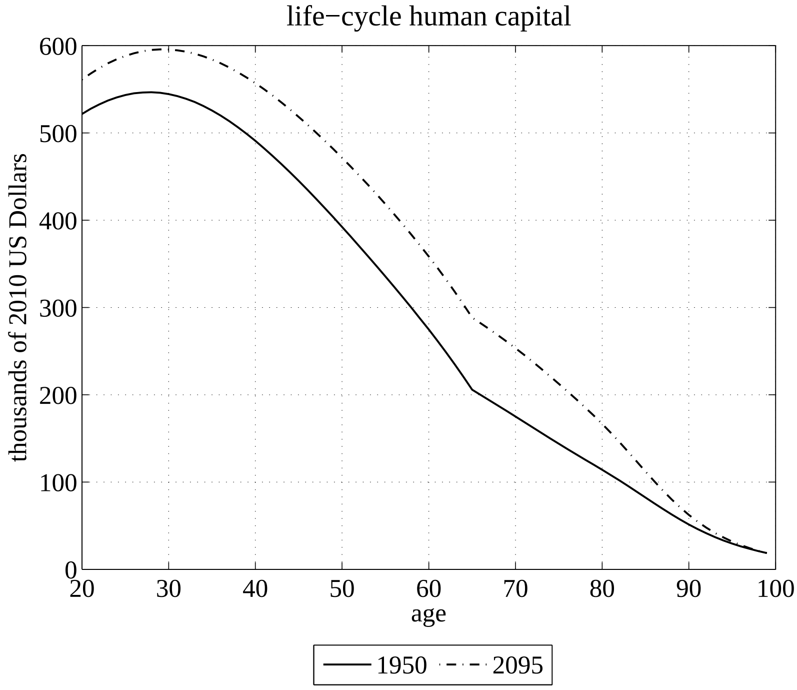

The solid line depicts simulated human capital for 1950, the dashed line for 2095 in the United States for ages 20 to 100.

risky asset. $h_{i,n,t}$ increases in the survival probability, as future non-financial income is more valuable the higher the likelihood of receiving it. This can be seen from Figure 2.11, which depicts the human capital defined in Eq. 3.21: Human capital in 2095 is larger than in 1950 at all ages. The difference is particularly large during the first years in retirement, as the increases in survival probabilities are larger towards the end of life, but discounting is strong at the very end due to our assumption of certain death at age 100.

We revisit a key insight by Merton (1971) on optimal portfolio choice in a model with CRRA preferences and human capital. The fraction of financial wealth held optimally in risky assets is

$$\omega_{i,n,t} = \frac{rp_t}{\vartheta\sigma_\eta}\left(1 + \frac{h_{i,n,t}}{a_{i,n,t}}\right) \tag{3.22}$$

where rp_t is the risk premium assumed constant for now and $a_{i,n,t}$ is total financial wealth as defined in Section 3.4. Individuals always want to hold a constant fraction of their total (financial and non-financial) wealth in risky assets, and adjust their portfolio such that the risky share increases in $h_{i,n,t}$ and decreases in $a_{i,n,t}$. An increase in longevity will increase both human capital and total financial wealth. The relative rate of change over the life-cycle determines the effect of demographic change on the risky share. While the increase in savings is particularly strong from age 60 onwards, thus shifting the portfolio towards safe assets, the increase in human capital is most pronounced around retirement, which is why there is almost zero change in the risky asset share around age 60. For the very young and very old, the difference in $a_{i,n,t}$ is marginal.

Valuation Effect

Both the distribution of cohorts across age groups and the individual life-expectancy will affect the average per capita demand for assets at given rates of return. In contrast, the per capita supply stays constant. Therefore, the returns need to adjust to clear markets. This affects financial income and the asset demands for the next period. We call this the *valuation effect*.

While the distribution effect and the life-cycle effect work to increase asset demand over the demographic transition, the valuation effect is a counteracting force that decreases aggregate asset demand. The market-clearing returns decline strongly over the demographic transition, as we show in Figure 2.5. Thus, the assets held by individuals at the end of the demographic transition have lower yields than at the beginning of the transition.

Further, individuals adapt their savings decisions and portfolio choice to the low-return environment. Policy functions for consumption and the risky asset share at different rates of return can be found in the Appendix, Figures 2.21 and 2.22. Here, we discuss the intuition. Consumption is increasing in the safe asset return, which is due to a low intertemporal substitution. The income effect (higher returns implying higher expected future income, a part of which will already be consumed today) dominates the substitution effect (higher returns make saving more attractive). The risky share decreases in the risk-free rate. Naturally, a higher safe return makes safe assets a more attractive investment at a constant risk premium. An increase in the risk premium induces agents to hold a higher share of their portfolio in risky assets, as these pay out more in expectations. This can be also seen in Eq. (3.22). However, the higher risky share is paralleled with a decrease in consumption at a higher risk premium. Our preference assumption induces agents to be prudent. They want to increase precautionary savings due to the fact that future consumption will be riskier.

Overall, the decline in returns that we predict makes agents want to accumulate more savings. The implication for the risky share is ambiguous.

Channel Decomposition

We now decompose the aggregate simulation results into the *distribution effect*, the *life-cycle effect* and the *valuation effect*, both over time and across age groups. We carry out three counterfactual experiments for the open economy period 1990 to 2095, focusing on the US. The first serves to isolate the life-cycle channel. In this experiment, the aggregate demand for safe and risky assets is calculated using the actual population distribution and evaluating the portfolio at the market-clearing rates of return of Section 5.1. However, agents expect demographics to stay constant at 1990 values. Thus, we do not allow them to internalize their own extended life span relative to 1990. The second experiment serves to isolate the effect of the distribution channel. In this exercise, individuals take their survival probabilities into account correctly and evaluate their portfolio at the market-clearing rates solved for in our baseline model, but the size of each age group is kept fixed at the level of 1990. This identifies the amounts of assets and the risky share demanded if the population structure stayed constant over time. Third, we isolate the valuation channel by solving for the optimal asset demand and the risky share at 1990 returns. Cohort sizes are allowed to vary over time and individuals take changing survival probabilities into account.

Figure 2.12: Channel decomposition for aggregate asset demand, United States

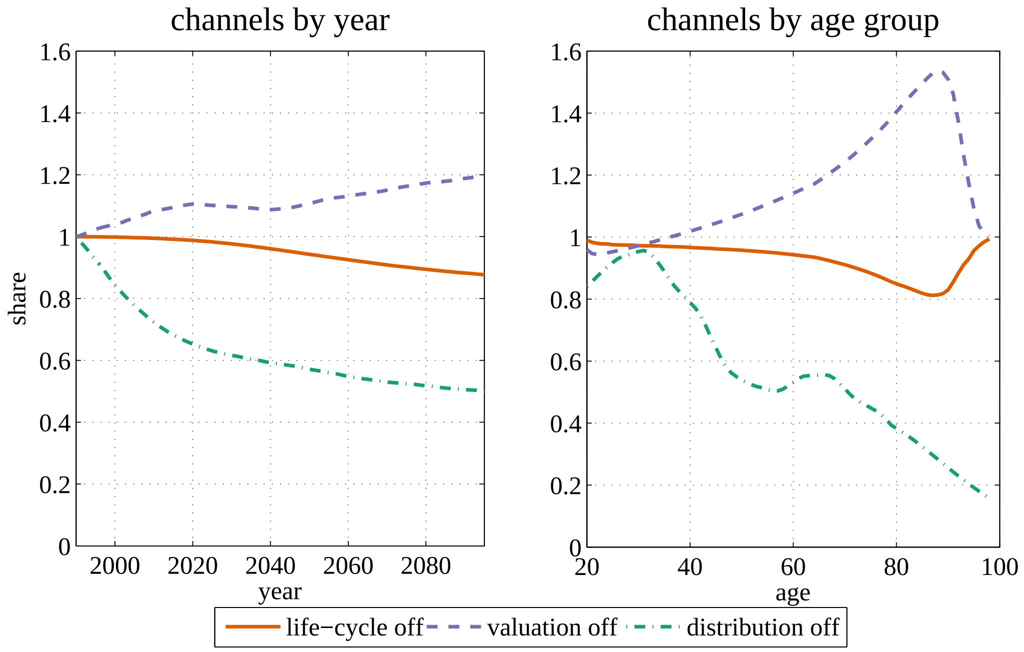

The left panel shows the time path for the share of aggregate asset demand resulting from shutting off the channels one by one. The orange solid line shuts off the life-cycle effect. The lilac dashed line shuts off the valuation effect, the green dashed-dotted line the distribution effect. The right panel shows the same results calculated as average 1990 to 2095 shares of total aggregate assets by age.

The time path of this comparison and the average contribution to the cross-sectional aggregate assets over the same period are shown in Figure 2.12. The smallest aggregate savings response both over time and across age groups is prompted by shutting down the life-cycle channel. By the end of the simulation period, the asset holdings would still be 88 percent of actual holdings had individuals not internalized that they are likely to live longer. The distribution channel has a much larger impact: Only 50 percent of actual asset holdings would result if the size of the population age groups stayed constant over the years. The valuation effect, as described above, works in the opposite direction of the other two, being the equilibrating force that aligns demand and supply. Had there not been a re-valuation of individual portfolios, the total asset demand would have been 120 percent of actual asset demand by 2095.

There are pronounced differences in the cross-section. The distribution channel becomes more important for older cohorts. This is because there are large increases in the number of old people between 1990 and subsequent years. The US baby boomers are also visible in this graph. Given that the reference period for cross-sectional changes is 1990, a time when the baby boomers were of ages between 25 and 45, we see hardly any effects of shutting the distribution channel off for those age groups. The life-cycle effect matters most for those individuals nearing the end of the life-cycle. This is the group where increases in longevity are the largest, as those at younger ages have high survival probabilities throughout the whole simulation period, and we force survival probabilities to converge to 0 by the age of 100. The valuation effect works to offset the life-cycle effect and is the mirror image of the left panel

Figure 2.13: Channel decomposition for aggregate risky share, United States

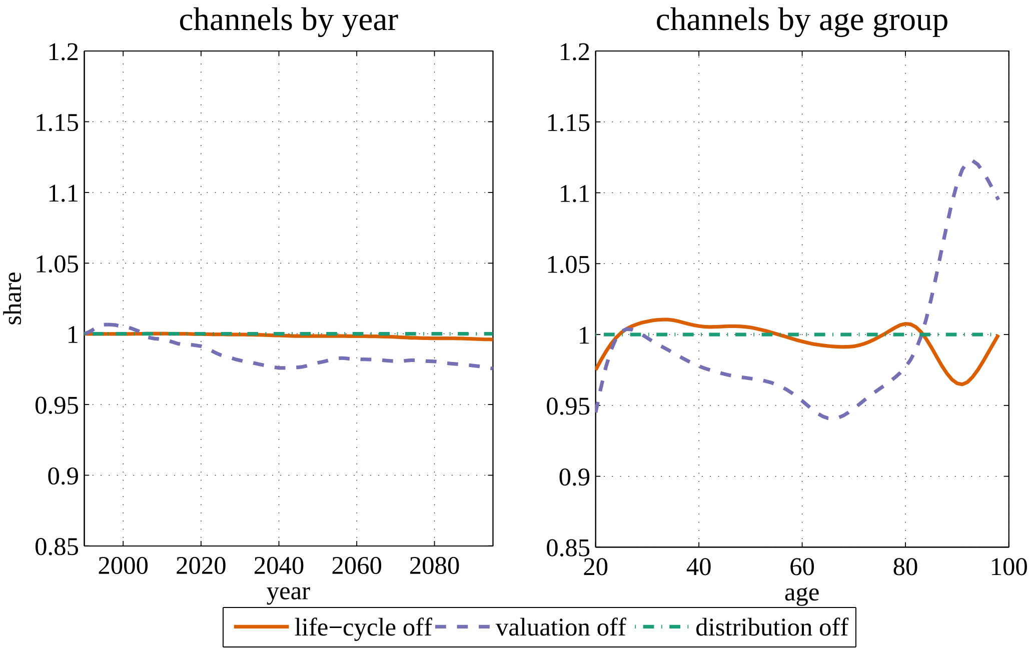

The left panel shows the time path for the share of aggregate risky portfolio share resulting from shutting off the channels one by one. The orange solid line shuts off the life-cycle effect. The lilac dashed line shuts off the valuation effect, the green dashed-dotted line the distribution effect. The right panel shows the same results calculated as average 1990 to 2095 shares of the risky portfolio share by age.

of Figure 2.10. Absent return adjustments, those cohorts at the end of their life span would have been the ones increasing their savings the most, because of the increase in survival probabilities, in combination with a higher number of people at those ages.

Figure 2.13 shows the same decomposition for the risky asset share. As the aggregate risky share corresponds to the average per capita risky share, keeping the population structure fixed does not have any effect. The effect of disabling life-cycle responses also only has minor effects over time, and in the cross-section matters only for the very old. Their reaction is linked to the savings response in Figure 2.12: Higher savings of the very old increase their share of financial wealth in total wealth. In consequence, the portfolio is shifted towards safe assets, as we discussed in Section 5.2. The same reasoning in reverse applies to the effects of disabling the valuation channel. Counterfactually higher savings over time result in a lower desired risky share. In the cross-section, the valuation effect again works to offset the life-cycle effect, mirroring the right panel of Figure 2.10.

Thus, while the distribution effect plays a big role for aggregate asset holdings, the valuation effect is a more important determinant of the risky asset share. We therefore need a complex structural model that takes into account both effects (plus the life-cycle effect) in order to fully capture how demographic change affects the demand for safe and risky assets and ultimately, external positions.

6. Conclusion

Differences in the demographic structure of regions can explain net external asset positions that vary by asset type. We show this by introducing an overlapping generations model with endogenous portfolio choice over the life-cycle. Two fully integrated regions with different demographics trade assets with each other. In the model, a region that is younger than its counterparty has a higher relative demand of risky assets. This is due to a decreasing preference for portfolio risk over the life-cycle. In consequence, a positive net external asset position and a negative net external debt position emerges vis-à-vis the older region. In calibrating the model, we focus on two developed regions: the US and 15 European Union member states. Throughout the simulation period of 1990 to 2095, the US is the younger region. The predicted external positions show an asymmetry in the riskiness of assets held, which reproduces the pattern observed in the data between 1990 and 2015: The US is a net exporter of risky assets and a net importer of safe assets. We identify different channels through which demographics affect aggregate asset demand and portfolio choice. The shift of the age distribution of the population from young cohorts with little savings towards older cohorts with large savings is the most important channel for aggregate asset demand. Changes in market-clearing rates are the most important channel for the risky portfolio share.

The simulation focuses exclusively on the US vis-à-vis 15 European countries. However, the results of this paper hold more generally. The strong association of US relative demographics with US external positions can also be found with respect to other countries. This serves as a further validation of the model. Many young regions that have currently relatively closed capital accounts may integrate with world financial markets over the next decades. We would expect demographic differences to lead to large external imbalances.

In order to keep the role of relative demographics clear, we have abstracted from several important issues. First, the analysis of different pension systems and reforms thereof. As pension systems in their current shape may become unsustainable over the next decades, different types of reforms are discussed on both sides of the Atlantic. Policies meant to neutralize the impact of demographics in one region will have implications for aggregate asset demands, country portfolios and returns. Second, we exclude an endogenous response of asset supply. Incorporating a fully formulated supply side, additional effects of demographic change would become visible: firms are likely to increase the amount of capital used in the production process, and thus the supply of risky assets would rise. While this would make the decline in the rates of return less pronounced, we should at the same time observe wage increases, which affect asset demand. We leave these aspects to future work.

Appendices

A. Data Definitions and Further Results

A.1 Data Sources

Population data: Data on the population structure of the US and the EU is retrieved from the United Nations Population Division (World Population Prospects, 2015 Edition). The data series are: life expectancy at age 60 for both sexes; old-age dependency ratio (Age 65+ / Age 20-64) for both sexes; annual population by age for both sexes; abridged life table (interpolated) for both sexes. For 1950 to 2015, the UN provides data, for 2016 to 2095 projections. For the projections, we choose the medium variant for fertility and rely on normal assumptions for migration. Data and projections on survival probabilities are provided in the the abridged life tables for five-year age groups and until 1990 only until age 80. We interpolate data within the age group bins and for ages 81 to 98 by setting the survival probability to 0 at age 99; interpolation is done using shape-preserving cubic interpolation.

US household data: For cross-sectional holdings of safe and risky assets, we use the Summary Extract data from the Survey of Consumer Finances for survey waves 1989 to 2013 (9 waves in total; in 2013 US Dollars). We drop households at the top 1 percent of the wealth distribution. For the estimation described in Appendix A.3, we calculated moving ges of three consecutive age observations in order to have sufficient observations in each cohort-age-time cell. For the statistical exercise, we interpolate the estimated coefficients and use the trend series after HP-filtering with smoothing parameter 400. Risky and safe assets are defined in Table 2.3. To distinguish between risky and safe assets, we mostly follow the definitions provided by the SCF and common practice, for example as in Chang et al. (2014). Total financial assets are the sum of safe and risky assets.

EU household data: The Household Finance and Consumption Survey (HFCS) has been carried out by the Eurosystem in 2010/11 and 2014/15 for 12 out of the 15 countries in our EU sample, excluding Denmark, the United Kingdom and Sweden. We pool both waves and treat observations for each country as resulting from the same data-generating process. Risky and safe assets are defined so as to correspond closely to the definitions applied to SCF data shown in Table 2.3. However, asset categories are less detailed than in the SCF.

International positions data: For bilateral external asset positions in equity and long-term debt, we use the IMF's Coordinated Portfolio Investment Survey (CPIS). Data is available for 1997 and 2001 to 2015. According to the IMF's Balance of Payments Manual 6, equity securities comprise all instruments usually shares, stocks, participations and similar documents. Long-term debt securities are negotiable instruments serving as evidence of debt, giving the holders the unconditional right to fixed or contractually determined variable payments. The maturity of a long-run debt instrument in the CPIS is more than one year or with no stated maturity. We continue the data series back until 1994 by using the Treasury International Capital (TIC) database (US Department of the Treasury). To conform with the definition in CPIS, we use historical bilateral data from US Residents' Portfolio Holdings of Foreign Securities and Foreign Residents' Portfolio Holdings of US Securities.

We also use TIC to construct bilateral positions on short-term debt. TIC provides data on banking claims on and liabilities to foreigners from the US going back until 1980. Banking claims are defined as short-term securities, excluding equities, further described in Table 2.2.

Due to the longer available series for bilateral banking claims, the safe asset position prior to 1994 comprises predominately the other investment category of the international investment position. Starting in 1994, we supplement TIC data with the data constructed from CPIS, so that the safe position is net claims from the banking sector plus net debt assets (long term only) divided by GDP.

A.2 Definition of Asset Types

Table 2.2: Definition of safe and risky assets in international data

Source	**Risky assets**	**Safe assets**
CPIS	equity and investment fund shares, e.g. shares, stocks, participations or similar documents	short- and long-term debt instruments, e.g. bonds, debentures, treasury bills, negotiable certificates of deposit, commercial papers, bankers' acceptances
TIC		deposits, short-term negotiable securities, US treasury bills, certificates with maturity of one year or less and borrowing and lending

Table 2.3: Definition of safe and risky assets in household data

Source	**Risky assets**	**Safe assets**
SCF	stocks; stock mutual funds; IRAs/Keoghs invested in stock; other managed assets with equity interest (annuities, trusts, managed investment accounts) if invested in stocks; thrift-type retirement accounts invested in stocks; savings accounts classified as 529 or other accounts that may be invested in stocks	transaction accounts; certificates of deposit; bonds (except mortgage-backed); mutual funds invested in bonds; quasi-liquid retirement accounts (IRAs and thrift-type accounts) and individual retirement accounts/Keoghs if invested in bonds; savings bonds; cash value of life insurance; other managed assets (trusts, annuities, managed investment accounts) if invested in bonds
HFCS	Stocks; assets in managed accounts; mutual funds invested in stocks or hedge funds; life insurance and pensions (partly)	sight accounts; savings accounts; federal government or state bonds; money market accounts; life insurance and pensions (partly)

A.3 Estimation of Life-cycle Financial Assets and Risky Asset Share

In our estimations of financial asset holdings and the risky asset share over the life-cycle, we closely follow Fagereng et al. (2017). We apply the Heckman estimator where the first stage is a probit estimation of the decision to participate in the stock-market, P_{inct}, by individual i aged n and belonging to cohort c in year t,

$$prob(P_{inct} = 1|x) = prob\,(\delta_n A_n + \delta_c C_c + \delta_t D_t + \delta_0 Trend + \upsilon Z_{inct} + \eta_{inct} > 0)$$

and the second stage is the estimation of the risky share conditional on participation,

$$\omega_{inct} = \beta_n A_n + \beta_c C_c + \beta_t D_t + \beta_0 Trend + \delta Z_{inct} + \delta_2 \lambda_{inct} + \varepsilon_{inct}.$$

A_n C_c, D_t are age, cohort and time dummies, *Trend* is a linear trend, and Z_{inct} are individual-specific controls of the household head. We include a dummy for whether the individual is married or living with a partner, for home ownership, for whether the household has a high-school degree or a college degree, and the number of children living in the household. λ_{inct} is the inverse Mills ratio from the first-stage regression and η_{inct} and ε_{inct} are error terms.

Applying the Deaton-Paxson method, we introduce the constraints $\sum \beta_t = \sum \delta_t = 0$. The linear time trend accounts for the fact that possibly, there is some trend over time, which would be wrongly captured by the age or cohort proxies if not controlled for.

For life-cycle financial asset holdings the estimation equation is

$$z_{inct} = \gamma_n A_n + \gamma_c G_c + \gamma_t D_t + \psi Z_{inct} + \nu_{inct}$$

where G_c is a proxy for cohort effects, the macroeconomic conditions faced by cohort c at the point of entering into the labor market. It is approximated by the deviation from trend real GNP per capita at the age of 20. All other variables are as defined in the risky share equation. Since the SCF over-samples rich households, we use survey weights in the estimation of financial asset holdings. Ideally, we would do the same for the estimation of the risky share, but this is difficult to implement in Stata. However, since the weighted and unweighted mean of the conditional risky share are similar (0.48 and 0.46, respectively), ignoring the weights is not likely to introduce a large bias to the estimation.

The first-stage result of the Heckman procedure is presented in Figure 2.14 and depicts the familiar hump-shape of age-specific participation in the stock market. Tables 2.4 and 2.5 show the coefficient estimates for the relevant variables.

Table 2.4: Estimation results for life-cycle financial assets

VARIABLES	c	s/e
dummyage_29	-19,067	(15,503)
dummyage_32	-4,937	(15,281)
dummyage_35	12,778	(15,437)
dummyage_38	23,980	(15,261)
dummyage_41	50,448***	(15,059)
dummyage_44	85,612***	(15,277)
dummyage_47	109,589***	(15,389)
dummyage_50	186,724***	(15,543)
dummyage_53	219,442***	(15,804)
dummyage_56	263,626***	(16,260)
dummyage_59	307,681***	(16,668)
dummyage_62	353,126***	(16,979)
dummyage_65	305,959***	(17,177)
dummyage_68	352,932***	(17,826)
dummyage_71	314,525***	(18,262)
dummyage_74	294,168***	(18,943)
dummyage_77	280,443***	(19,317)
dummyage_80	291,289***	(30,594)
real_gnp	-2.247	(16.60)
kids	-3,218	(2,586)
married	-160,002***	(5,884)
hhouses	84,880***	(6,390)
educ_1	-318,017***	(8,086)
educ_2	-252,299***	(6,016)
Observations	187,520	
R-squared	0.064	

Note: The table presents regression outputs for financial assets as described above. Column (1) are the estimated coefficients, column (2) corresponding standard errors. Year dummies are included but not shown. Age dummies are relative to the group of 26-year olds (value 30939.27). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2.5: Estimation results for risky share and participation

VARIABLES	risky share		participation	
	c	s/e	c	s/e
trend			0.00219	(0.0129)
kids	-0.00204**	(0.000949)	-0.0783***	(0.00321)
married	-0.0113***	(0.00313)	-0.488***	(0.00742)
hhouses	0.0496***	(0.00491)	0.782***	(0.00819)
educ_1	-0.101***	(0.00911)	-1.433***	(0.0117)
educ_2	-0.0765***	(0.00421)	-0.755***	(0.00734)
dummyage_26	0.0293**	(0.0120)	-0.155*	(0.0922)
dummyage_29	0.0268**	(0.0130)	-0.108	(0.0893)
dummyage_32	0.0353***	(0.0125)	-0.0736	(0.0846)
dummyage_35	0.0579***	(0.0120)	-0.0322	(0.0800)
dummyage_38	0.0640***	(0.0114)	0.0479	(0.0755)
dummyage_41	0.0521***	(0.0110)	0.137*	(0.0711)
dummyage_44	0.0443***	(0.0108)	0.136**	(0.0669)
dummyage_47	0.0562***	(0.0105)	0.237***	(0.0629)
dummyage_50	0.0554***	(0.0103)	0.242***	(0.0591)
dummyage_53	0.0474***	(0.0103)	0.275***	(0.0556)
dummyage_56	0.0367***	(0.0103)	0.321***	(0.0525)
dummyage_59	0.0231**	(0.0104)	0.306***	(0.0496)
dummyage_62	0.0245**	(0.0106)	0.253***	(0.0471)
dummyage_65	0.0145	(0.0108)	0.159***	(0.0452)
dummyage_68	0.0117	(0.0112)	0.0969**	(0.0440)
dummyage_71	-0.000439	(0.0116)	0.0428	(0.0433)
dummyage_74	0.000418	(0.0121)	0.0189	(0.0435)
dummyage_77	-0.0131	(0.0126)	0.0690	(0.0443)
λ	0.0538***	(0.00894)		
Constant	0.778***	(0.177)	0.192***	(0.0321)
Observations	187,147		187,147	

Note: The table presents regression outputs for the Heckman estimation described above. Column (1) are the estimated coefficients of the second stage, with column (2) corresponding standard errors. Column (3) are estimation coefficients of the first stage with (4) corresponding standard errors. Year and cohort dummies are not shown. Age dummies are relative to the group of 80-year olds (risky share 0.322, participation rate 0.161). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 2.14: Participation by age, United States

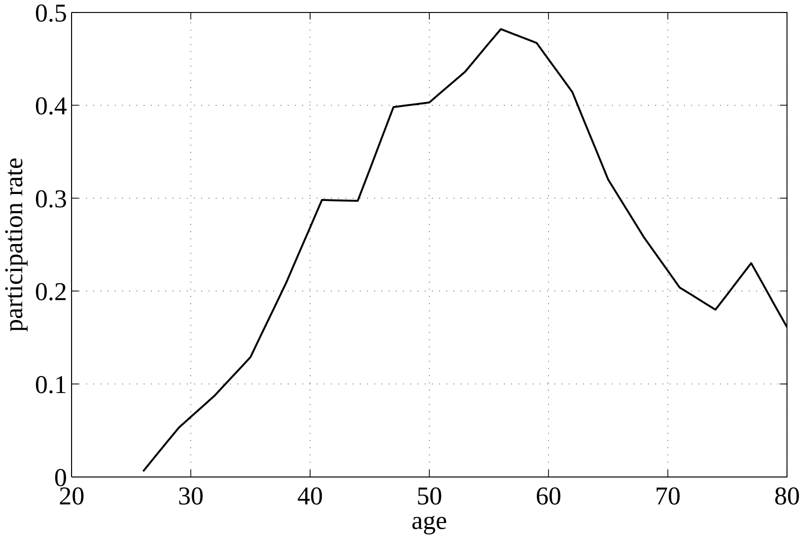

The graph shows the participation rate estimated as the first stage of a Heckman estimation applying the Deaton-Paxson technique for cross-sectional age cohorts from 20 to 80. *Source:* authors calculation using data from SCF for waves 1998 to 2013.

A.4 Additional Data

Figure 2.15: Share of US and EU in world asset holdings

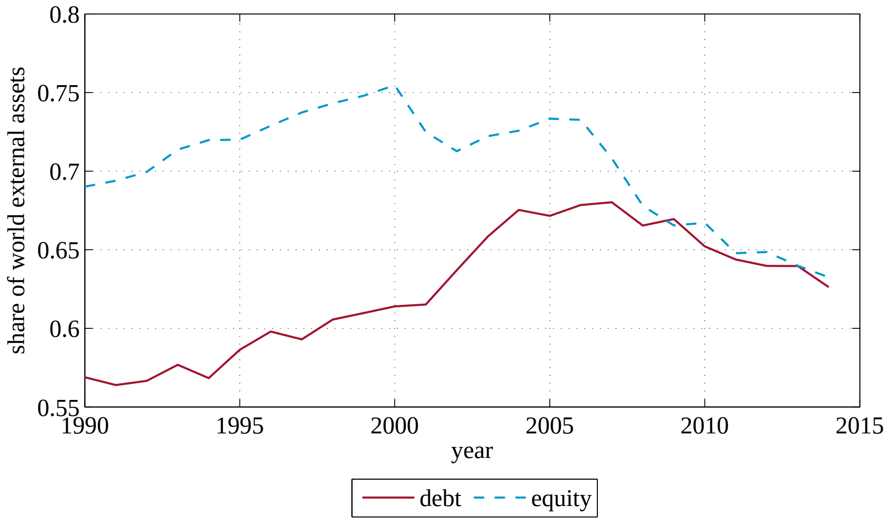

Share of US and EU external asset holdings in total world external asset holdings by asset type. *Source:* Authors' calculations based on Lane and Milesi-Ferretti (2007).

Table 2.6: Share of bilateral positions in total external positions

	EU		**US**	
	debt	equity	debt	equity
OFCs as RoW	0.45	0.44	0.45	0.42
OFCs as bilateral	0.57	0.64	0.59	0.62

Source: CPIS, using a definition of offshore financial centers (OFCs) by the IMF. The lower bound attributes all OFC-US and OFC-EU investments to the rest of the world (RoW), the upper bound attributes them to the respective other region. Numbers are averaged over 1997 - 2015. We restrict our focus on assets, because countries do not report liabilities to the IMF, so that these positions are inferred.

Figure 2.16: Counterfactual aggregate assets by age

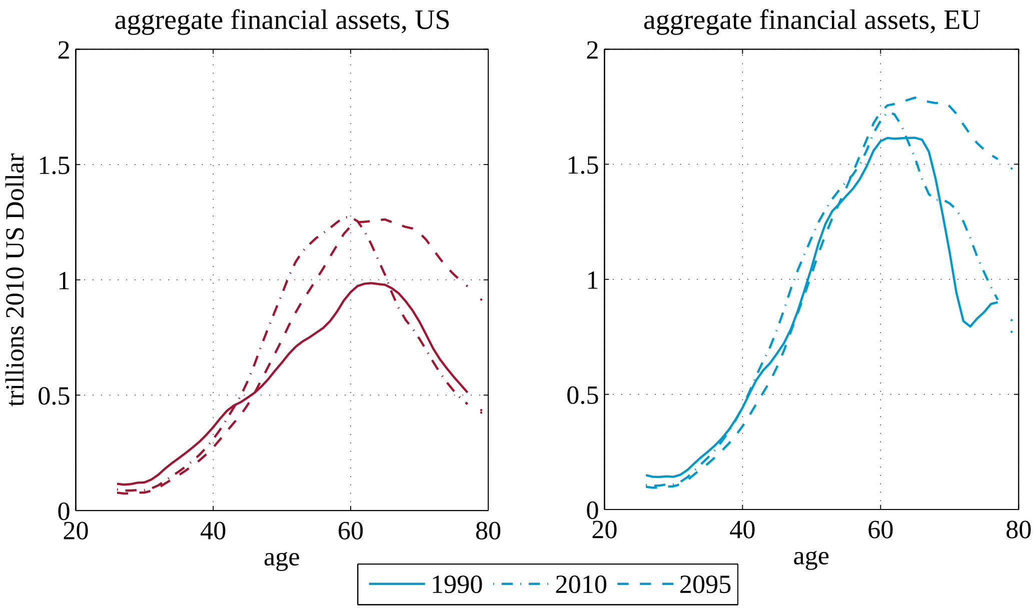

The left panel shows the implied aggregate assets in the US for years 1990, 2010 and 2095. The right panel shows aggregate financial assets for the same years in the EU (both in trillions of 2010 US Dollars). The total population size is held fixed at the level of 2010. *Source:* Authors' calculations using data from SCF and UN Population Prospects.

Figure 2.17: De jure capital account openness

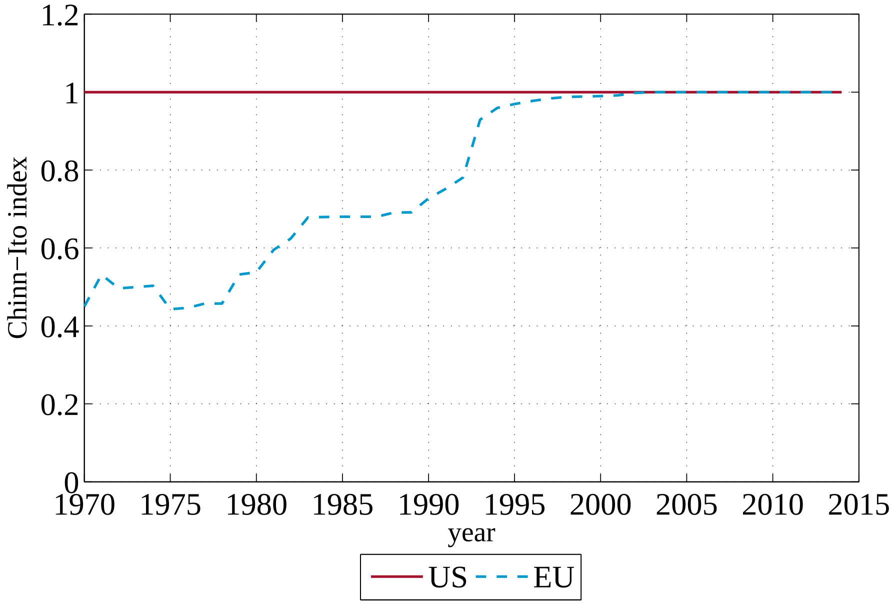

Source: Chinn and Ito (2008). *De jure* capital account openness describes the degree to which capital account openness is allowed for or restricted by law. Based on the IMF's Annual Report on Exchange Arrangements and Exchange Restrictions (AREAER), the Chinn and Ito (2008) index is the first principal component of a principal component analysis on four binary indicators that codify restrictions on cross-border financial transactions. 0 signifies completely closed while 1 signifies completely open.

Figure 2.18: 10-year government bond returns

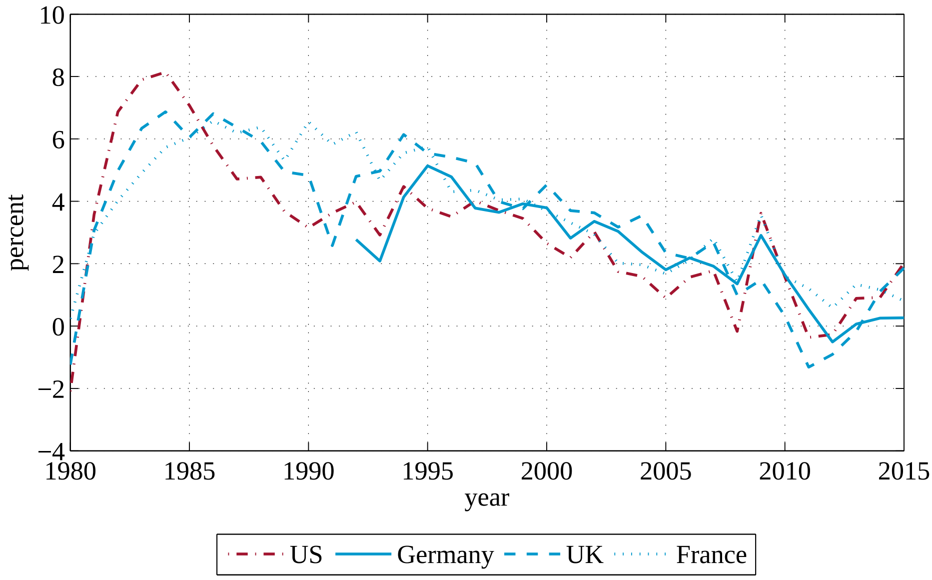

The graph shows 10-year government bond returns in real terms for the US, Germany, France and the UK. *Source:* Federal Reserve Bank of St. Louis.

Figure 2.19: Financial asset holdings and risky share, US vs. EU

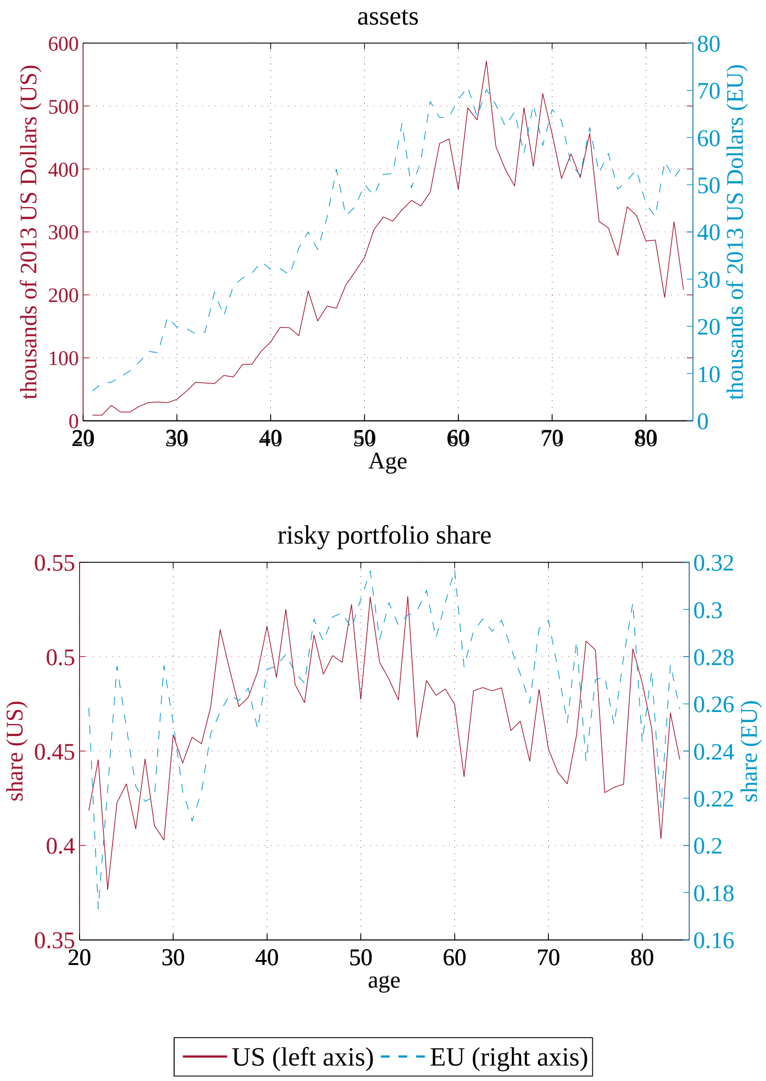

The upper panel compares cross-sectional financial asset holdings for the US (red, solid line) and the EU (blue, dotted line) both in trillions of 2013 US Dollars. The lower panel compares the conditional risky asset share for the US and the EU. *Source:* SCF, pooled waves 1998 to 2013, and HFCS, pooled waves 2010/11 and 2014/15.

A.5 Demography and External Positions in the Rest of the World

The association between demographics and the risk-content of external positions can also be found in a larger set of countries. Figure 2.20 plots countries' old age dependency relative to that of the US against the US net external debt and equity position vis-à-vis those countries. Each dot represents a country-year observation. We show observations for all countries and years that are reported to the CPIS, omitting those that have capital accounts that are less than 50 percent open according to the measure of Chinn and Ito (2008). There are altogether 69 countries, with data available for time spans between 1997 and 2015. The correlation is -0.3 for debt (-0.28 if we also include countries with less open capital accounts) and 0.29 for equity (0.23 when all are included).

This is not just a further validation of our model, but also allows for an outlook on the US external balance sheet in a future where many countries that are so far little financially integrated (for example China) engage more in international asset trades. Some of these countries have a much younger population than advanced regions.

Figure 2.20: US External positions and relative demographics

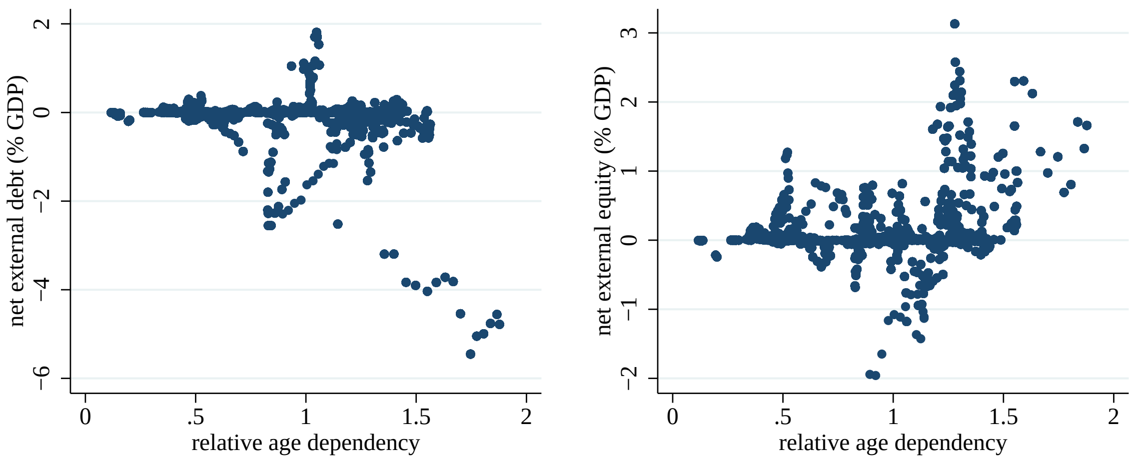

The left panel plots the net external debt position of the US relative to other countries against the relative old-age dependency ratio. The right panel plots net external equity positions against relative old-age dependency ratios. Observations with de jure capital openness of less than 0.5 are omitted. *Source:* CPIS.

A.6 Additional Results

Figure 2.21: Policies at varying safe returns

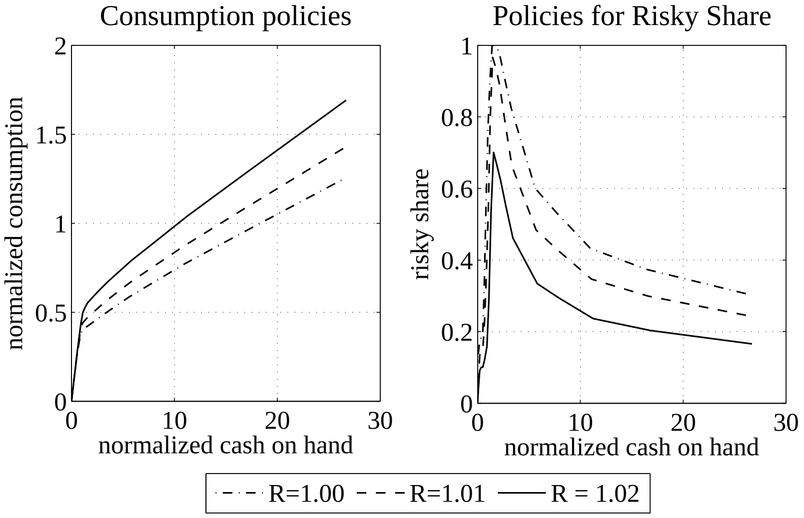

The left panel shows the level of consumption for varying safe returns for a US individual at age 45 in the year 1950. The right panel shows the risky share at varying safe returns for the same individual. The results are qualitatively equivalent for all age groups. The risk premium is kept constant at 0.04.

Figure 2.22: Policies at varying risky returns

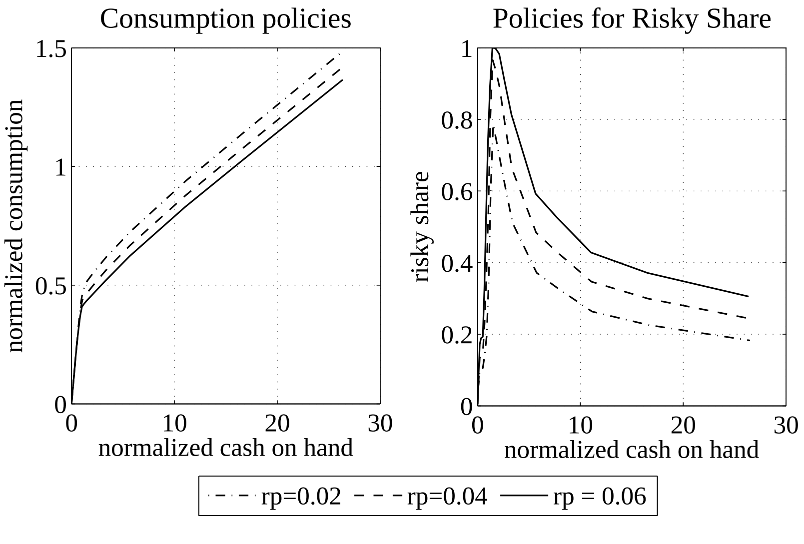

The left panel shows the level of consumption for varying risk premia for a US individual at age 45 in the year 1950. The right panel shows the risky share at varying risk premia for the same individual. The results are qualitatively equivalent for all age groups. The safe rate is kept constant at 1.02.

A.7 Population Structure in the Simulation

Figure 2.23 highlights the difference between the actual (observed or predicted) population size and a population size that would result when starting from the 1950 population structure and updating it annually using birth rates and survival probabilities, as described in the model. The simulated population structure in both regions is lower for two reasons: first, because the simulation ignores immigration, which is generally positive and highest for the youngest age group. Second, because in the simulation we use 1950 as a steady state, and slowly move out of it over the next years (so that the last cohort born into the steady state only dies by 2020). In reality, the 1950 population did not represent a steady state, which means that in most cases, $L_{n,t} \times \Delta\delta^{n-1}_{t-1,t} \neq L_{n,t-1}$. Not taking this into account would lead to unrealistically low population in the simulation, although as far as this effect concerns both regions to the same extent, the difference on external positions would not be large.

Figure 2.23: Life-cycle results comparison

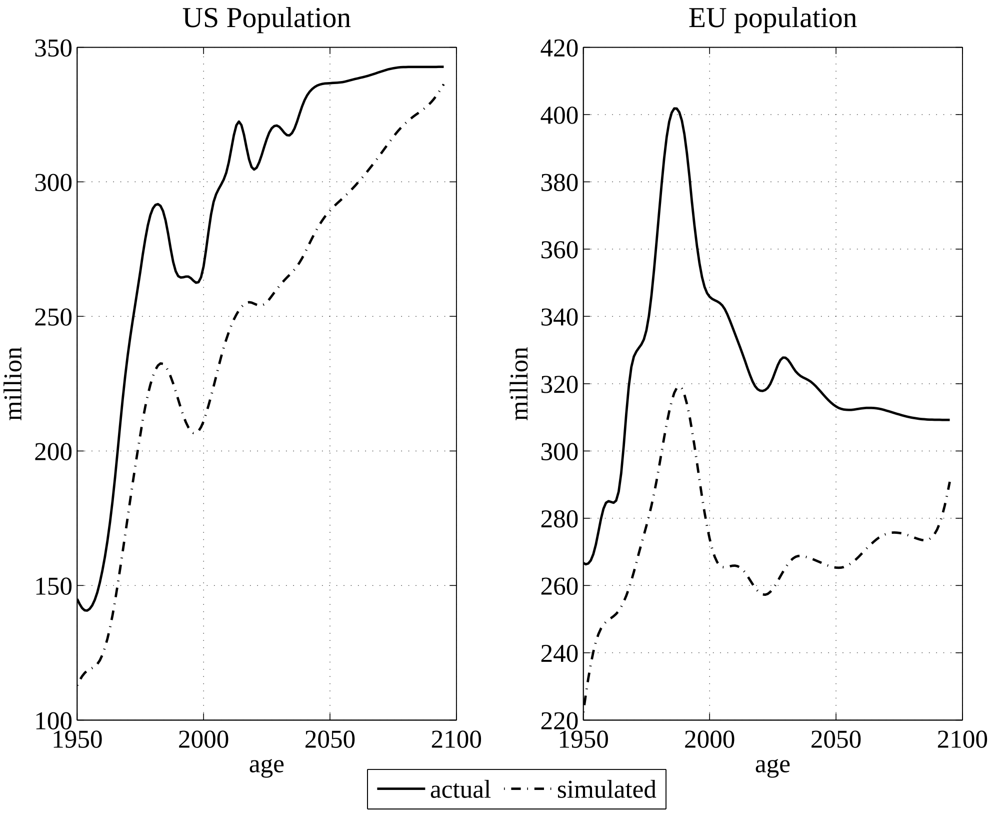

The left panel shows total population between 1950 and 2095 for the US: The solid line is the actual population size, the dashed line is the population size that would result from a simulation using the population structure of 1950 and updating it for subsequent years using survival probabilities and birth rates. The right panel shows actual and simulated population for the EU for the same time period.

B. Appendix: Model Solution

A closed-form solution does not exist for the household's portfolio problem or aggregate asset returns. We describe the numerical solution algorithm below.

B.1 Solution to Individual's Optimization

We can formulate the household's problem recursively dropping time, individual and age subscripts and denoting next period variables a_{t+1} as a'. The Bellman equation is the indirect utility function $V_n(x, z)$ and given as the solution to the following problem for an individual of age n:

$$V_n(x, z) = \max_{c,a',\omega'} \frac{c^{1-\vartheta} - 1}{1 - \vartheta} + \beta\delta_n \mathbb{E} V_{n+1}(x', z'), \tag{B.1}$$

$$\text{subject to} \qquad c + a' = x \tag{B.2}$$

$$\omega' \geq 0 \tag{B.3}$$

$$(1 - \omega') \geq 0, \tag{B.4}$$

where cash on hand $x' = \tilde{R}'\omega' a' + R'(1 - \omega')a' + y'$ is the endogenous state variable that summarizes the available resources for the individual in the next year of life. The vector $z' = [\epsilon'\theta'\eta'\zeta']^\intercal$ summarizes the exogenous states (the shock realizations in the next period).

The individual policy functions are the solution to the individual's problem given in Eq. (B.1) and are derived by backward induction from the terminal period of life. In order to conserve computational time for grid search on a' and x, we use the endogenous gridpoint method (Carroll, 2006; Hintermaier and Koeniger, 2010). The continuous distributions for shocks z' are approximated using the Gauss-Hermite quadrature method (Judd, 1998). Individuals take the aggregate returns as given when making decisions. We solve for the policy functions as decision rules given the endogenous state variable cash on hand and a grid of aggregate asset returns, representing expected future returns. In the sections that follow, we suppress indices for region, individuals, and time, so as to focus on the recursion by age in the solution algorithm. Policy functions are denoted with hat characters.

We construct a triple exponential grid for savings $\mathscr{A} = (a_1, a_2, ...a_j)'$ and a linear grid for the two asset returns, $\mathscr{R} = (R_1, R_2, ...R_k)'$ and $\tilde{\mathscr{R}} = (\tilde{R}_1, \tilde{R}_2, ...\tilde{R}_l)'$ each with $j = k = l = 15$ gridpoints. The return grid for the risky return that excludes the shock realization is

$$\tilde{R}_l - \epsilon = 1 + d + rp_l \in [1.03, 1.10],$$

and the shocks for returns are log-normally distributed with parameters given in Table 2.1. We then form 3-dimensional arrays of grids with dimensions $j \times k \times l$, each element of which corresponds to a unique combination of the elements of $\mathscr{A}$, $\mathscr{R}'$ and $\tilde{\mathscr{R}}'$. We solve for policy functions using the following procedure:

1. For the terminal period of an individual of age N^d, the solution is trivial: $\hat{c}_{N^d} = \hat{x}_{N^d}$, saving is zero and the continuation value of the Bellman equation is $V_{N^d+1} = 0$ such that the value function is $V_{N^d} = u(c)$.

   Working recursively from age $n = N^d - 1$ to age N^b:

2. Use the policy function for next period consumption $\hat{c}_n' = \hat{c}_{n+1}$ from step 1 to find the portfolio allocation $\hat{\omega}_n'$ that satisfies the portfolio optimality condition

$$\beta\delta_n\mathbb{E}\left((\tilde{R}_l{}' - R_k')u'(c_{j,k,l}')\right) = 0 \qquad \forall \quad j,k,l.$$

   The expected future consumption is evaluated using the condition $c_{j,k,l}' = x_{j,k,l}'$ where the indices j,k,l correspond to savings and expected future returns.

3. Given optimal portfolio choices $\hat{\omega}_n'$ from step 2, use the Euler equation to solve for the policy function for consumption $\hat{c}_n$ for each grid point:

$$c_{j,k,l} = \left[\beta\delta_n\mathbb{E}_{x',z'}\left((\tilde{R}_{j,k,l}'\omega_{j,k,l}' + R_{j,k,l}'(1-\omega_{j,k,l}'))(c_{j,k,l}')^{-\vartheta}\right)\right]^{-1/\vartheta}$$

4. Collect the policy functions for age n on grids for savings and asset returns. The corresponding gridpoints for the endogenous state variable $\hat{x}$ arise endogenously from the budget constraint: $\hat{x} = \hat{c} + \mathcal{A}'$, where $\mathcal{A}'$ has the same dimension as $\hat{c}$. This gives triplets $\hat{c}_n$, $\hat{\omega}_n'$ and $\hat{x}_n$.

5. Stepping back to age $n-1$, the pairs from step 4 are used to interpolate $u'(c')$ in the determination of the optimal portfolio allocation as in step 4.

The algorithm provides 4-dimensional arrays of policy functions based on age, savings a, the safe asset return R' and risky asset mean return $\tilde{R}'$ which can be used to simulate the aggregates in the economy and, using the market-clearing conditions, find the equilibrium asset returns.

B.2 General Equilibrium

We carry out a deterministic simulation and solve for the equilibria by explicit aggregation. We set the draws for idiosyncratic labor income shocks equal to their expected values. Since we assume them to be *i.i.d.* and uncorrelated with the aggregate return shock, and because there is a continuum of individuals born in each cohort, idiosyncratic shocks will aggregate by age to expected values. For the aggregate risk, we set $\epsilon_t = 0$ in each period of the simulation. An alternative approach, which we will carry out in future iterations of the work, takes aggregate risk into account by running Monte-Carlo simulations over several time paths of draws of the aggregate shock to the risk premium and calculating the distribution of possible aggregate time paths.

In order to simulate the model for a representative individual at each age, we require the state variable cash on hand, x, as well as future expected returns, which are themselves endogenous. Given the high-dimensional nature of the model, for the time being, we reduce the problem's dimensionality by assuming that individuals set expectations of the risk premium and future safe dividend to current equilibrium values:

$$\mathbb{E}(rp') = rp \qquad\qquad \mathbb{E}(d') = d.$$

Of course, with this assumption, we isolate the agents' forward looking behavior in terms of the exogenous survival probability. An important step for future work will be to provide the benchmark of fully forward looking behavior of individuals, incorporating fully endogenous expectations of future returns.

Equilibrium interest rate solution: For the two regions, we each simulate 305 years and 305 cohorts $(2 \times (N^d - N^b) + 2095 - 1950)$ on the grid of safe and risky return realizations as discussed above. For each year, we aggregate asset positions, consumption and non-financial income over living cohorts as in Eq. (3.14). For simulation years prior to opening in 1990, the relevant market clearing condition for asset holdings is given by Eq. (3.16) and after 1990 by Eq. (3.19). For each gridpoint for safe asset returns, we use interpolation to find the equilibrium risky rate of return as the solutions to:

$$F(R', \tilde{R}') \equiv B - \mathbb{B} = 0 \qquad\qquad G(R', \tilde{R}') \equiv S - \mathbb{S} = 0.$$

Asset supply is found using Eq. (3.15) with calibrated supply parameters given in Table 2.1. This algorithm results in two sets of return pairs that respectively clear bond and stock markets globally. Since combinations of world returns must clear both asset markets, we find the equilibrium pair as the intersection of the two sets of returns pairs. For each simulation year, there is a unique return pair that is contained in each set of returns. These equilibrium returns are shown in Figure 2.5.

Chapter 3

Benign Effects of Automation: New Evidence From Patent Texts

With Lukas Püttmann

We provide a new measure of automation based on patents and study its employment effects. Classifying all U.S. patents granted between 1976 and 2014 as automation or non-automation patents, we document a rise in the share of automation patents from 25 percent to 67 percent. We link patents to the industries of their use and, through local industry structure, to commuting zones. According to our estimates, advances in national automation technology have a positive influence on employment in local labor markets. Manufacturing employment declines, but this is more than compensated by service sector job growth. Commuting zones with more people working in routine occupations fare worse.

1. Introduction

What is the effect of automation technology on employment? The answer to this question is not obvious: While machines may replace workers, new jobs could also be created. For example, if self-driving vehicles become widely used, taxi and truck drivers might lose their jobs. Other sectors such as retail could, however, experience employment growth through lower transport costs.

To identify the employment effects of automation, this paper introduces a new indicator of automation technology. The large literature addressing this question has so far relied on indirect proxies of automation, such as routine task input (Autor, Katz, and Kearney, 2008, Autor, Levy, and Murnane, 2003, Goos and Manning, 2007, Autor and Dorn, 2013), investment in computer capital (Beaudry, Doms, and Lewis, 2010; Michaels, Natraj, and van Reenen, 2014) or investment in robots (Graetz and Michaels, 2015; Acemoglu and Restrepo, 2017). Many of these papers find evidence for job polarization, but the smaller literature on aggregate employment changes reports more ambiguous results. This may be due to difficulties in measuring automation comprehensively.

Our proposed automation indicator relies on patent grant texts. Patents are a natural candidate for measuring technological progress and frequently serve as proxies of innovation. However, few studies examine the consequences of technological progress through patents.

Also, while patent meta-data such as citation counts or the identity of innovators is used regularly (Hall, Jaffe, and Trajtenberg, 2001; Acemoglu, Akcigit, and Celik, 2014; Bell, Chetty, Jaravel, Petkova, and Reenen, 2017), the actual patent texts have not been in the focus so far. We classify patents as automation patents if their texts describe physical inventions (such as robots) or immaterial or conceptual inventions (such as software), which carry out a process independently of human interference.

We extract the texts of all 5 million U.S. patents granted between 1976 and 2014 and train a machine learning algorithm on a sample of 560 manually classified patents to sort patents into automation and non-automation innovations. As a result, we document a strong rise in both the absolute and the relative number of automation patents. As a share of total patents, automation patents have increased from 25 percent in 1976 to 67 percent in 2014. Applying a probabilistic matching that is based on Canadian patents, we link patents to the 956 4-digit SIC industries where they are likely to be used. In this way, we quantify trends in newly available technology at the industry level.

Next, we compare the indicator to established measures of automation. The number of automation patents is positively correlated across industries both with investment in computer capital and with robots shipments. More automation patents have been granted in industries with a larger share of employment in routine occupations in 1960, a result that is in line with the literature on routine-biased technological change. Also, industries with more automation patents were characterized by a rise in non-routine cognitive and non-routine interactive task input and a fall in routine cognitive and routine manual task input.

To estimate the labor-market effects of automation, we transfer our industry-level data to U.S. commuting zones through industry-county employment counts. Commuting zones approximate local labor markets as workers tend to look for jobs within commuting distance from where they live. We obtain a panel dataset of new automation technology across 722 commuting zones over 39 years. Up to the late 1980s, there was a higher density of automation in the Great Lakes region, but automation technology has become less geographically concentrated over time.

Our empirical analysis benefits from the fact that we examine *local* economic outcomes which are impacted by, but unlikely to affect, the innovation activity of industries at the *national* level. Our key assumption is that commuting zone-specific developments in the medium-run do not affect automation innovation in industries that operate there. This is plausible for the following reasons: First, we separate the industries where patents originate from where they are used. Second, many patents belong to foreigners and universities who respond to other incentives than local firms. And third, local industries are small in comparison to national aggregate industries. Our approach thus follows Bartik (1991).

Our main econometric analysis is a fixed effects panel regression for five-year periods. Interpreting the automation index as a flow measure of technology, we assess the relationship between the sum of automation and changes in employment. While we find a positive effect of automation on total employment, this is driven by job growth in the service sector, which compensates for a fall in manufacturing employment. This result is robust to adding a variety of other economic and demographic controls and to weighting patents by the number of citations they received. We also consider separately patents belonging to specific groups of assignees: universities and public research institutes, foreigners and governments. All three should be less responsive to US labor market trends than US companies. Our results hold in the regressions for the subgroups of patentees as well as in an instrumental variable regression. Lastly, we find that automation is associated with more job creation in commuting zones where the share of routine occupations is low.

All in all, our study thus shows automation to be more beneficial for employment than some of the previous literature (Autor et al., 2015; Acemoglu and Restrepo, 2017), which might be due to our broader definition of automation. Our results are in line with Gregory, Salomons, and Zierahn (2016), who show that the detrimental substitution effect of automation on routine jobs is more than compensated by a positive labor demand effect due to larger product demand.

In the final part of our paper, we apply our indicator to replicate two central papers (Autor and Dorn, 2013 and Autor, Dorn, and Hanson, 2015) that study the influence of automation on labor markets using the routine task share of jobs. First, we show that non-college employment rose in commuting zones where more automation patents could be used and where more people worked in routine occupations. Second, we find that automation leads to rises in employment levels even when controlling for Chinese import competition, which stands in contrast to Autor et al. (2015). We provide further evidence that employment increases were driven to a larger extent by flows into the labor force than by a fall in unemployment.

There are strengths and weaknesses to our approach to quantifying automation technology. Text classification is an inherently imprecise activity and we introduce further inaccuracies through probabilistic matchings of patents to industries and commuting zones. Also, we make assumptions on the usefulness of patents and the way they are implemented. On the upside, we have to impose fewer ex-ante assumption on the nature of advances in automation technology, compared to the literature using routine task shares or computer and robot investment. Our indicator allows us to closely track the technology frontier, translating newly granted patents into a fine-grained industry- or commuting zone-level dataset. With these caveats in mind, we consider our indicator a complement to previous measures of automation.

2. New Automation Index

This section introduces the new automation index. We start by arguing why patents are a suitable data source for measuring technological progress and then define automation. We show how we construct the indicator and how the classification algorithm works. Then, we explain how to link patents to industries in which they are likely to be used. The resulting indicator traces the technology frontier across 956 industries and 39 years and displays plausible co-movement with existing indicators of automation such as computer investment, the number of robots used in production and the share of routine tasks across industries.

2.1 Patents as Indicators of Technological Progress

The purpose of patents is to encourage innovation and technological progress by offering a temporary monopoly on an invention. Once granted, no one can re-engineer, create or sell the same object or idea. In return, the text of the patent is made publicly available. The language in the patent text is technical and highly standardized. Applicants have an incentive to provide exact and correct information about their innovation to obtain full protection of their ideas.

Professional patent examiners judge a patent's claims and make changes where appropriate. In return for disclosing the content of the innovation to the public, an intellectual property right is granted for 20 years. To be patentable, an innovation must be *novel*, *non-obvious* and *useful*. The description must further be exact and detailed enough to allow for replication and it must name the invention's most important application. All these characteristics make patents a valuable data source.

Researchers in economics have made frequent use of patents, often in the form of the database established by Hall et al. (2001). Griliches (1990) provides an extensive survey of various issues related to using patents in economics. However, patents are so far usually interpreted as proxies for innovative activity, not as increments of technological progress whose effects can be studied (for an overview of the more recent literature, see Nagaoka et al., 2010). This is related to the fact that existing research almost exclusively uses patents' metadata, such as the location or affiliation of a patentee or a patent's importance.[1]

Magerman, Looy, and Song (2010) note that there is almost no research which uses the actual texts of the patent document, although this has been recommended as early as Griliches (1990). An exception is Bessen and Hunt (2007), who identify software patents by searching patent texts for keywords. Our approach differs as we do not specify a priori which words to search for, but use a state-of-the-art text classification algorithm. Also, we apply the derived measure to study the effects of technology on the labor market, whereas the goal of Bessen and Hunt (2007) is to characterize firms that file software patents.

In other areas of economics, text search has become common, with Gentzkow and Shapiro (2010) and Baker, Bloom, and Davis (2016) being prominent examples of papers that use newspaper articles. However, patent texts hold several advantages for researchers over other document collections: The precise technical language with a high degree of standardization, the incentive to deliver correct information, the additional check through the patent examiners' review and the public access to patent grant texts make patents well suited for text search analysis.

Patent text analysis is common in the private sector for prior art and freedom-to-operate searches by firms and lawyers. However, none of these providers – to the best of our knowledge – offers a comparison of technological trends over time, which leads us to develop our own approach.

2.2 Patent Data

We obtain all 5 million utility patent documents granted in the United States from 1976 to 2014 from Google.[2] While Europe, Japan and increasingly China are also important patent legislations, of the roughly 10.9 million patents effective ("in force") worldwide in 2014, the largest fraction (about one fourth) had been granted in the United States (WIPO, 2016). In addition, the most important innovations are usually patented in all major patent legislations. These properties make U.S. patents a good proxy for the technological frontier in the United States and beyond. Also, given that this paper studies the effect of automation in the United States, U.S. patents are an obvious candidate for how available technology changes.

[1] Patent citations, in particular, are widely applied as indicators of the value of an invention, for example by Bell et al. (2017).

[2] `google.com/googlebooks/uspto-patents.html`

We only consider utility patents, which account for around 90 percent of all patents. Utility patents are "issued for the invention of a new and useful process, machine, manufacture, or composition of matter, or a new and useful improvement thereof" (USPTO, 2015). Other patent types are design, plant and reissue patents and do not track technology that we aim to measure. According to the United States Patent and Trademark Office (USPTO), in the period 1976-2014, 83 percent of all patents granted were owned by firms – mostly large multinational corporations. 15 percent of patents were owned by individuals and less than 2 percent by the U.S. government. About half of all patents are granted to foreign applicants, a share that has increased over time. During the period of our analysis, IBM, Canon and Samsung were the corporations with the largest number of patents granted (USPTO, 2014).

The patent grant document includes the title, patent number, name of the inventor, date, citations of other patents, legal information, drawings, abstract and a detailed description, as well as information on the technology class of the invention. Every patent is assigned one or more technology classification numbers by the patent examiner which describes technological and functional characteristics of a patent and on which we base our link from patents to industries. We exclude chemical and pharmaceutical patents from our classification.[3] The overwhelming majority of these patents do not meet our definition of an automation patent (14 out of 560 manually classified patents were automation patents from those sectors), but including these patents might distor our classification.

2.3 Definition of Automation

We define an automation patent to describe a *device that carries out a process independently*.[4] This broad definition captures technologies such as software, a robot used in a production or the self-driving vehicle mentioned in the introduction. The "device" can be a physical machine, a combination of machines, an algorithm or a computer program. The process it automates may be a production process, but also anything else where an input is altered to generate an output. An important element of the definition is the notion of independence: It works without human intervention, except at the start or for supervision. We require the automation innovation to be a reasonably complete process, product or machine. In addition, we require it to have an at least remotely-recognizable application. This excludes inventions that are minor parts of an automation innovation and highly abstract patents with no obvious application. We make no difference between process and product innovations, so an automation patent could describe either. Table 3.1 displays some examples of automation and non-automation patents.

2.4 Classification of Patents

Based on the definition above, all patents can be classified as either automation or non-automation patents. We use an automated approach. To train a classification algorithm,

[3]Excluded USPC technology numbers: 127, 252, 423, 424, 435, 436, 502, 510-585, 800, 930, 987.

[4]This is a standard definition that can be found in encyclopedias. For example, the Encyclopedia Britannica defines automata as "any of various mechanical objects that are relatively self-operating after they have been set in motion" and adds that "the term automaton is also applied to a class of electromechanical devices—either theoretical or real—that transform information from one form into another on the basis of predetermined instructions or procedures" (Encyclopædia Britannica, 2015).

Table 3.1: Examples of automation and non-automation patents

Patent title	**Patent number**	**Automation patent?**
"Automatic taco machine"	5531156	Yes
"Color measuring method and device"	6362849	Yes
"Coinfusion apparatus"	8857476	Yes
"Hair dye applicator "	6357449	Yes
"Hand-held scanner having adjustable light path"	5552597	No
"Bicycle frame with device cavity"	7878521	No
"Process for making pyridinethione salts"	4323683	No
"Golf ball"	4173345	No

Note: Authors' classifications according to manual coding guidelines. Click on the patent number for the weblink to the patent document.

we need reliable and objective classifications on which we can base the comparison. To this end, we manually classify 560 randomly drawn patents according to rules laid out in manual coding guidelines.[5] Baker et al. (2016) proceed similarly when they manually classify newspaper articles to check the performance of their dictionary-based classification. We aim to minimize coding mistakes and biases by providing a structured classification process, by classifying patents in random order and by reviewing every classification by a second person.

The language in patent texts might have changed over time. But patents from the 1970s read very similar to those from the 2000s and important technological classes such as computers and robots are developed and patented throughout the sample period. The technical nature of the documents and the fact that legal terms change more slowly than other language also makes it less likely that there are short-lived trends that could pose a problem for a classification based on specific terms.

From our sample of patents, we extract word stems, called *tokens*, with the Porter2 stemming algorithm. This shortens "automation", "automated", "automatically", "automatable" to "automat". Table 3.2 summarizes these tokens. A typical title contains about 5 tokens, a typical abstract about 36 and the rest of the patent (the "body") about 500 to 600.

Table 3.2: Tokens in 560 manually classified patents

Part	All tokens	Unique tokens	Mean	Median
Title	2796	1301	4.99	5
Abstract	20781	3971	37.11	36
Body	339366	31499	606.01	506.5

Source: USPTO, Google and own calculations.

In principle, one could now record for all 5 million patents whether they contain one of the roughly 32,000 tokens that we can assign probabilities to. But to keep the computation-intensive data collection feasible and to remove noise features, we use the *mutual information*

[5]See: `http://lukaspuettmann.com/assets/pdf/manual_coding_guidelines.pdf`

Figure 3.1: Words that indicate an automation patent

Note: Token size is proportional to the value of the mutual information criterion in sample 560 classified patents. We show only the 150 highest ranked tokens excluding chemical and pharmaceutical words.
Source: USPTO, Google and own calculations.

criterion to extract those tokens which are most informative about which class a patent belongs to. This is an established statistic for feature selection which prefers tokens that appear significantly more often in one of the classes and punishes tokens that appear rarely overall (Manning et al., 2009). We then pick the highest ranked (according to the mutual information criterion) 50 title tokens, 200 abstract tokens and 500 patent body tokens. The final search dictionary consists of 623 tokens.

Figure 3.1 visualizes the 150 tokens with the highest mutual information criterion. The most important token is unsurprisingly "automat". After that come "output", "execut", "inform", "input" and "detect". Some tokens are indicative of software, such as "microprocessor", "database", "comput", "program" or "transmiss". Others are more likely to appear in descriptions of physical machines, such as "motor", "move", "metal" or "apparatus". The last discernible group of tokens are action verbs that appear in descriptions of a wide range of independently operating devices, such as "distinguish", "command", "respons" or "perform".

Our algorithm emulates how a human being would have classified each patent. We apply the Naive Bayes algorithm which is a supervised learning method which is easy to interpret and which computationally scales well with large amounts of data. The "naive" assumption the probability of a token to appear in a document is independent from the appearance of other tokens. Despite its simplicity it has been shown to perform quite well (Domingos and Pazzani, 1997).[6] One reason for this that the low number of parameters it estimates make it unlikely to overfit (Murphy, 2012).

Manning et al. (2009) explain how this algorithm picks the class c for every document d with maximum a posteriori probability $P(c \mid d)$. In our analysis, the documents d correspond to patent grant texts and the two different classes are automation patents and non-automation

[6]Gentzkow et al. (2017) also recommend this algorithm if the number of observed features (tokens) is much larger than the size of the training sample, as is the case in our analysis. Antweiler and Frank (2004) proceed similarly, as they manually classify 1000 messages and then use the Naive Bayes algorithm to generalize to over 1.5 million other messages.

patents. In the *Bernoulli* Naive Bayes that we use, every document d is represented by a vector e, where entry e_i ($i = 1, \ldots, M$) is 1 if token i appears at least once in the document and 0 if it does not. Patent texts contain matter-of-fact language, where words are often repeated. So the occurrence of a word is more important than the frequency of its appearance and we therefore ignore how often a word appears in a document.

According to this language model, in any document in class c the token e_i occurs with conditional probability $P(e_i \mid c)$. Therefore, the probability of a document d to show up in class c is

$$P(d \mid c) = \prod_{1 \leq i \leq M} P(e_i \mid c), \tag{3.1}$$

and the conditional probability of document d to belong to class c is according to Bayes' rule[7]

$$P(c \mid d) \propto P(c) \prod_{1 \leq i \leq M} P(e_i \mid c). \tag{3.2}$$

We estimate the prior $\widehat{P}(c)$ as the relative frequency of documents in class c in the training set. This is $\widehat{P}(\text{autom}) = \frac{147}{483} = 0.304$, as about a third of eligible patents (i.e., after removing chemical and pharmaceutical patents) were manually labeled as automation patents. We then estimate the conditional probabilities of a certain token to occur in class c, $\widehat{P}(e_i \mid c)$ as

$$\widehat{P}(e_i \mid c) = \widehat{P}(i \mid c) e_i + (1 - \widehat{P}(i \mid c))(1 - e_i), \tag{3.3}$$

where $\widehat{P}(i \mid c)$ is the share of documents with token i in class c. In this way, we calculate posterior probabilities for all 5 million patents to belong to either class and assign each patent to the class with the higher posterior probability.

Table 3.3: Contingency table

		Computerized		
		No	Yes	
Manual	No	323	88	411
	Yes	25	124	149
		348	212	560

"No": not automation patent

Table 3.3 shows how human examiners and how the computer algorithm classified the set of manually investigated patents. Both the manual coding and the algorithmic classification judged around a quarter of patents to be automation patents. In 80 percent of cases ($= \frac{323+124}{560}$) both approaches agreed. The probability of a false positive (type I error) is 21 percent ($= \frac{88}{411}$). The probability of a false negative (type II error) is 17 percent ($= \frac{25}{149}$).

While some share of misclassified patents remains, as long as there is no underlying bias in the classication this should only add noise to our indicator series as we only aim to

[7] $P(c \mid d) = \frac{P(c)P(d|c)}{P(d)} \propto P(c)P(d \mid c)$.

Figure 3.2: Patents, 1976-2014

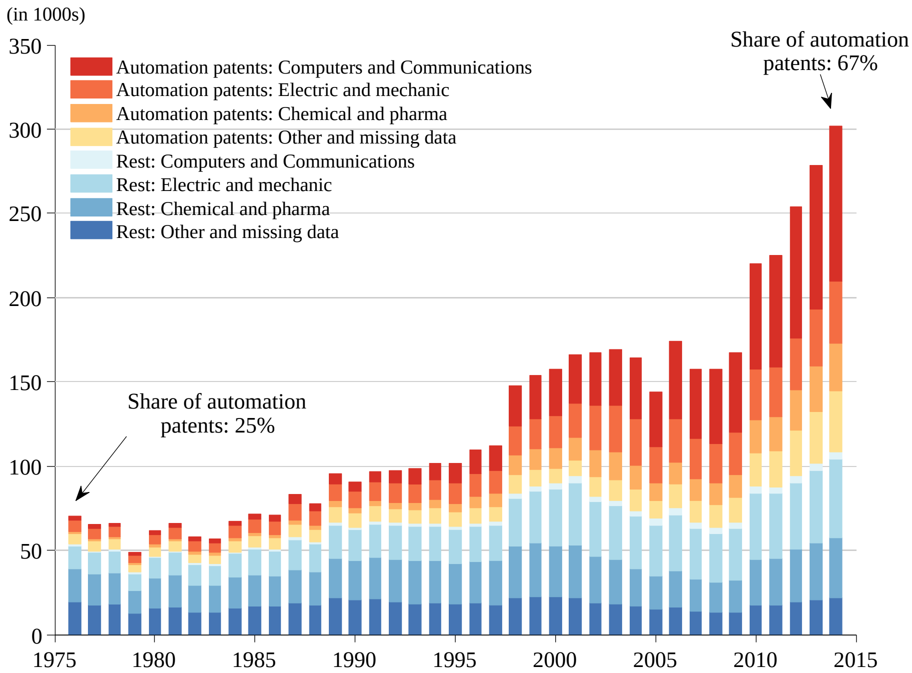

Note: See text for classification of automation patents and assignment of patents to categories.

Source: USPTO, Google, Hall, Jaffe, and Trajtenberg (2001) and own calculations.

approximate trends in technology over time. Any noise should therefore push our empirical results towards zero, making it harder to detect an effect of automation.

A more precise classification might be possible when including patents' other observable characteristics such as their technological class (USPC and IPC numbers), grant years, the origins of inventors or the sector of firms. But we keep the classification into automation and non-automation separate from these observables to allow comparing automation trends across time and industries, without making these associations automatic.

2.5 Aggregate Properties of the Indicator

Figure 3.2 is a graphical representation of all 5 million patents granted in the United States between 1976 and 2014. We show patents by when they were granted, not when applied for, as inventions are unlikely to be shared before they are protected by a patent.

There has been a steady increase from 70,000 granted patents in 1976 to more than 300,000 patents in 2014. Over the whole period, we classify 2.2 million of these as automation patents. The red-shaded parts of the bars show the patents which we classified as automation patents and blue colors signal all other patents. We observe a sharp upward trend in automation patents from 16,000 in 1976 to 180,000 in 2014. The share of patents related to automation also increased, from 25 percent of patents in 1976 to 67 percent of patents in 2014. Table 3.16 in the Appendix provides the yearly numbers.

Figure 3.2 further shows broad categories of patents based on an aggregation method by

Hall et al. (2001) which relies on the technological classification (USPC number) of patents.[8] Patents in the sub-category computers and communication have become much more frequent over the sample period and we mostly classify them as automation. Many of these are likely software patents. Electrical, electronic and mechanical patents also contribute significantly to the stock of automation patents. Robots, for example, fall in this category. By design, most chemical and pharmaceutical patents are not classified as automation patents, but they make up a large portion of the non-automation patents.

The rise in the total number of patents granted is a potential concern for the interpretation of the time-dimension of patent texts. If the nature of patents had changed in parallel with the number, so if the increase in patents is due to something else than an increase in research productivity, the data might not be comparable across time. An increase in the number of automation patents would then not be interpretable as an increase in automation technology. Kortum and Lerner (1999) evaluate different possible explanations for why the number of patent grants has changed: increased patent protection due to patentee-friendly court rulings, regulatory capture by large firms that patent eagerly, new technology fields producing patentable inventions (e.g., information technology, biotechnology and financial intermediation) and more applied research. The authors refute all hypotheses except for the increase in research productivity. This result is in line with an OECD survey (OECD, 2004) in which 94 percent of surveyed firms responded that an increase in the number of inventions was an important or very important driver of their increased patenting activity (66 percent very important). In contrast, changes in patentability played only a minor role. We therefore conclude that the quality of patents granted has not changed over time and that we do not need to worry about any distortive effects of a change in grant numbers. As an additional check, we compute a deflated version of our indicator, for which we divide the number of automation patents in each industry and year by the total number of patents granted in that specific year relative to the number of patents granted in 1990. The resulting measure is an automation count in units of 1990 patents, which takes higher values for earlier years and lower values for later years than the original measure. Our empirical results in section 4 are insensitive to the time deflation.

2.6 From Patents to Industries

Various researchers have proposed matchings of patents to industries. Hall et al. (2001) identify firms filing for patents and Lybbert and Zolas (2014) propose an automated approach that compares descriptions of industries with descriptions of patents' technological classes. The OECD (2011) reviews these techniques in more detail and Griliches (1990) describes the difficulties in matching patents to industries.

However, we are interested in how automation technology affects labor markets. Therefore, we aim to find the industries where automation patents are *used*, not where they originate. These two need not be the same, so that the industry of the patentee is not necessarily the industry we want to assign the patent to. As an example, IBM owns many patents that are not used in the computer industry, but by companies in the manufacturing or in the retail sector. These patents are either sold or licensed out. Attributing them to the computer industry would overstate the automation intensity there, while understating it in the other sectors.

[8]Note that this is a different classification than the one we will employ to match patents to the industries they

Table 3.4: Automation patents across industries of use

Industries	Manufac-turing	Automation patents (1000s)	Share	SICs (1987)
Computers	✓	499	88%	357
Other electronics	✓	250	46%	36*
Measuring instruments; watches	✓	193	60%	38
Telephones and telegraphs	✓	185	68%	3661
Machines	✓	183	40%	35*
Hospitals		137	46%	8062
Househ. audio and video equip.	✓	104	69%	3651
Other services		118	47%	70-89*
Transportation equipment	✓	115	39%	37
Chemicals, rubber, plastics, oil	✓	101	18%	28, 30, 29
Utilities (transport, gas, sanitary)		57	44%	E
Fabricated metal products	✓	51	33%	34
Medical laboratories		37	64%	8071
Construction		34	24%	C
Printing publishing; paper	✓	34	32%	26, 27
Metal, stone, clay, glass, concrete	✓	29	22%	32, 33
Retail and wholesale trade		26	32%	G, F
Agriculture, forestry and fishing		24	33%	A
R&D, management	✓	23	64%	87
Miscellaneous manufacturing	✓	20	38%	39
Public administration; finance		20	47%	J, H
Food, tobacco	✓	19	24%	20, 21
Mining		16	37%	B
Apparel, wood, furniture	✓	15	17%	22-25, 31
total		2,290	46%	

Note: Patents are counted if they can be used in an industry, as described in text. Numbers are sums of patents 1976-2014. Shares are calculated by dividing automation patents by all patents in industry. An asterisk * indicates that some subindustries are shown separately.
Source: USPTO, Google, Silverman (2002) and own calculations.

Linking patents to the industries of their use is difficult. If we wanted to measure the *actual* usage of a specific patent in a certain industry, we would need data on out-licensing. But this information is not available, as firms and research institutions have incentives to keep their licensing agreements private. Interpreting patents more indirectly as a proxy for automation technology rather than a direct measure, we can use information about the areas in which patents can *potentially* be applied. There have been attempts by Schmookler (1966) and Scherer (1984) to manually classify patents and link them to industries of use, but this would not be feasible for a large number of patents. Patent offices themselves usually do not provide information on the link of patents to industries. However, we benefit from an exception to this rule by the Canadian patent office. Between 1978 and 1993, Canadian

are likely to be used in. See section 2.6.

patent officers assigned industries of use for all granted patents. Based on this information, Kortum and Putnam (1997) assembled the "Yale Technology Concordance", a way to link patents through their technological classification to the industries in which they are likely to be used. This is based on the assumption that the pattern linking patents' technological class to industries of use should be similar in Canada and the United States. We use the files provided by Silverman (2002), who calculates empirical frequencies for cross-overs from patent technology classes (IPCs) to 1987 SIC industries using 148,000 patents granted between 1990 and 1993.[9]

This allows for a probabilistic matching. We connect a patent to an industry with the probability of being used in that industry. So if patent A is used in two industries X and Y, then half the patent count is assigned to industry X and half to Y. However, patents are often assigned several IPC technology classifications. In that case, we divide each value for that patent by the number of its IPCs. So if patent A now is assigned another IPC number, then only a quarter of its value will now be attributed to industries X and Y each and the rest to industries in the new IPC. This fractional counting of patents ensures that more general patents that are assigned to several IPCs do not have get more weight than more specialized patents that are assigned to fewer IPCs.[10]

As a result, we obtain an annual dataset of new patents and new automation patents that can be used in 956 industries and over 39 years. Table 3.4 displays all automation patents by industries of use over the whole time period 1976-2014. (The totals differ slightly from Appendix Table 3.16 due to rounding errors and the probabilistic conversion to patent equivalents as described before.)

Out of a total of 2.3 million automation patents, 1.8 million (79 percent) are used in the manufacturing sector (division D in SIC 1987). Half a million automation patents could be used in the production of computers (SIC 357) which includes personal computers, mainframes, storage devices, terminals, billing machines, automatic teller machines and peripheral equipment such as printers, scanners, office equipments or typewriters. The production of electronic devices, sensors and communication equipment also received a large number of automation patents. Outside of the manufacturing sector, hospitals, utilities and medical laboratories were assigned a large number of automation patents. In large parts of the economy – such as agriculture, mining, public administration, finance or retail – only few automation patents were granted. We also calculate the share of patents used in an industry that we classify as automation. This ratio is high for the computer industry or communication-related industries and is low for the chemical industry or "Apparel, wood, furniture".

In our following empirical analysis, we interpret these indices as worker intensities by fully assigning all new (automation) patents in an industry to each person employed in that industry and year. This is equivalent to assuming that patents assigned to an industry will potentially be used by everyone working in that industry. If we considered our indicator narrowly as an exact measure of the use of patents in the production process, this would not be a realistic assumption. But to us, a patent is just one part of an innovation process that will produce many types of outputs. Being a measurable outcome of this process, patents serve

[9]The fact that we use only data for 1990–1993 means that the matching should be most precise during this period, while becoming less exact the further away we move from this period. It helps that this periodis in the middle of our sample, but the fact that patents grow much more near in the later years is some cause for concern.

[10]This also enables us to interpret the resulting indicator as full patent equivalents which we will still refer to simply as "patents" in the remainder of the paper.

Figure 3.3: Comparison with other indicators of automation

(a) Computer investment (NIPA), 1976-2001

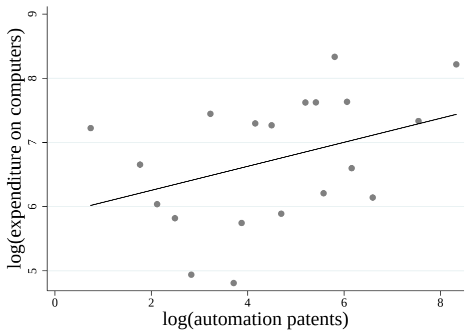

(b) Computer investment (ASM), 2002-2014

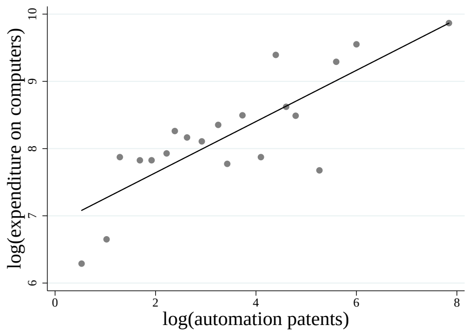

(c) Robots (IFR), 2004-2014

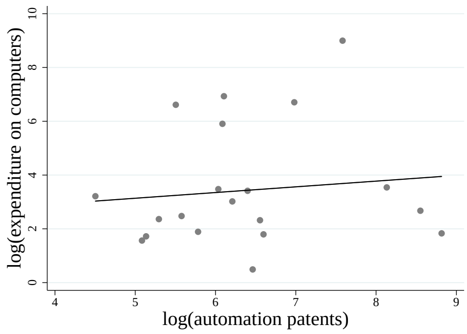

Note: NIPA computer investment is the mean of 1976-2001 in millions of 1996 U.S. dollars, ASM computer investment is the mean of 2002-2014 in thousands of 2009 U.S. dollars. Robots is the mean number of robot shipments in the U.S. over 2003-2014 (U.S. data for 2003-2010 are imputed from North America data). Automation patents are counted for the same time period as the respective comparison data. All three figures show binscatters of log values.
Source: USPTO, Google, Silverman (2002), NIPA, ASM and IRF (2014).

as a proxy for it.In our regressions we will use the total number of automation patents as our main explanatory variable, but we will also control for the amount of all other patents that can be used in an industry.

3. Comparison with Previous Automation Proxies

Next, we analyze how our new industry measure of automation technology is related to established automation indicators. Previous proxies of automation differ from ours along two lines. First, they are indicators of realized automation in the production process, not indicators of automation technology. Second, most capture only one specific facet of automation technology, such as computers or robots, while our indicator incorporates both and even allows delineating it from other kinds of technological progress.

As a measure of computerization, studies use survey data of computer use at the workplace (Autor et al., 2003, Beaudry et al., 2010) or industry-level investment in computer capital (Autor et al., 2003, Michaels et al., 2014). Frey and Osborne (2017) manually assess

Table 3.5: Relationship between automation patents and other automation proxies

	contemp. automation	1st lag of automation	2nd lag of automation	3rd lag of automation
	A. Time fixed effects			
computer invt (ASM)	0.391*** (0.0568)	0.393*** (0.0572)	0.395*** (0.0576)	0.398*** (0.0582)
computer invt (NIPA)	0.194** (0.0948)	0.193** (0.0953)	0.191* (0.0958)	0.189* (0.0963)
robot shipments (IRF)	0.0846 (0.308)	0.121 (0.323)	0.167 (0.342)	0.220 (0.365)
	B. Time and industry fixed effects			
computer invt (ASM)	0.247** (0.125)	0.250** (0.125)	0.257** (0.124)	0.259** (0.128)
computer invt (NIPA)	0.336 (0.252)	0.322 (0.243)	0.272 (0.231)	0.250 (0.220)
robot shipments (IRF)	0.350*** (0.120)	0.335*** (0.107)	0.465*** (0.156)	3.095** (1.440)

Note: ASM: N = 2,524 (14-3 years with max 465 industries); NIPA: N = 1,380 (26-3 years with max 71 industries); IFR: N = 186 (11-3 years with max 24 industries). The table shows results of regressions of various automation proxies on the log of (one plus) the automation measure at the contemporaneous level and various lags. Each coefficient estimate represents a separate regression. Data are annual; industry fixed effects are at the most disaggregate level of industries, but at maximum at the 3-digit SIC level. Regressions include a constant. Industry-clustered standard errors in parenthesis. *** p < 0.01, ** p < 0.05, * p < 0.1.

the probability of computerization of a number of occupations. Akerman et al. (2015) exploit a natural experiment, the introduction of broadband internet in Norway, to study employment effects of automation.

As a proxy for physical automation innovations, Graetz and Michaels (2015), Acemoglu and Restrepo (2017) and Dauth et al. (2017) count the number of robots used in production, a dataset assembled by the International Federation of Robotics. Lewis (2011) applies a more general understanding of automation by looking at adoption rates for new automation technologies, but with limited coverage of industries.

To show how our index relates to some of these measures, Figure 3.3 correlates automation patents with investment in computer capital and shipments of robots. We use two different data sources for investment in computer capital: The National Income and Product Accounts (NIPA), which provides annual data until 2001 for 71 2- and 3-digit SIC industries and the Annual Survey of Manufactures (ASM), which is available annually from 2002 onwards and for 465 4-digit SIC industries, the majority of them being manufacturing industries. As a measure of robots, we use the dataset on robot shipments by the International Federation of Robotics, which is provided at an annual frequency for North America starting from 2004 for 24 SIC industries. All correlations are highly positive, which indicates that our automation measure captures both advances in robotics and in software, which are then

Figure 3.4: Automation patents and routine labor

(a) Routine task share

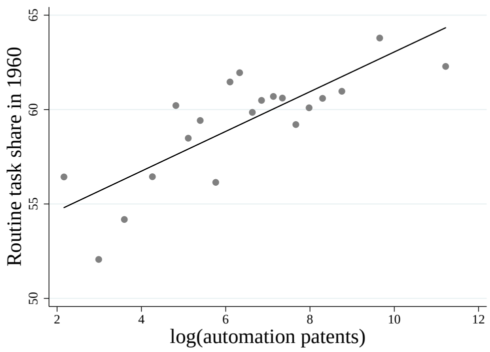

(b) Per decade

	Correlation
1976-1985	0.35
1986-1995	0.34
1996-2005	0.32
2006-2014	0.28

Note: Binscatter of log of total number of automation patents in industries against routine task input share in 1960 across 258 SIC 3-digit industries, 1976-2014.
Source: Autor et al. (2003) and see text.

translated into production and trade of computers and robots.

This positive relationship holds even in a panel regression that controls for time- and/or industry-specific effects. Table 3.5 shows further that the correlations are significant at various lags of automation, accounting for the fact that it may take some time to translate a patented innovation into the actual production of this technology. Although the results across the three variables are not directly comparable due to different time periods and industries covered, the link with computer investment might be slightly stronger than that with robot shipments.

Another way to contextualize our indicator is to evaluate how it relates to the nature of jobs. A large strand of literature, pioneered by Autor et al. (2003), analyzes the labor market effects of automation based on the assumption that automated machines are good at carrying out repeated tasks and fail at complex intellectual or manual tasks. For each occupation, they calculate what share of a job comprises routine (manual or cognitive) tasks. The resulting routine-task index thus measures the outcome of automation given specific – theory- and data-supported – assumptions. Weighing the index by occupation-specific employment, Autor et al. (2003) further create a routine task intensity measure across 140 industries from 1960 to 1998, based on which they show that changes in routine-task intensity are predicted by investment in computer capital: The share of non-routine tasks increases, whereas that of routine tasks decreases as a result of computer investment.

Figure 3.4 plots the routine task share of industries in 1960 against new automation technology patented between 1976 and 2014.[11] The relationship between automation patents and the routine-task index is positive: The larger the routine task share of an industry in 1960, the more automation technology was subsequently invented, patented and potentially used in that industry in the following decades. Our indicator thus seems to be capturing the same phenomenon as described by the literature on routine-biased technological change. The cor-

[11] Data on routine-task intensities at the industry level is obtained from David Autor's website `economics.mit.edu/faculty/dautor` (accessed 14.07.2015). Their dataset is for U.S. Census industries which we translate into SIC industries using a concordance scheme of the U.S. Census Bureau.

Table 3.6: Automation and industry task input

		Outcome: Within-industry change in task input		
		1970-1980	1980-1990	1990-98
Δ Non-routine analytic	Auto Technology	-0.012	0.033***	0.011
		(0.011)	(0.005)	(0.014)
	Constant	0.068***	0.110***	0.139***
		(0.011)	(0.014)	(0.019)
	R^2	0.004	0.019	0.001
Δ Non-routine interactive	Auto Technology	0.017*	0.062***	0.007
		(0.010)	(0.008)	(0.018)
	Constant	0.131***	0.206***	0.279***
		(0.017)	(0.030)	(0.036)
	R^2	0.004	0.016	0.000
Δ Routine cognitive	Auto Technology	-0.032**	-0.066***	-0.031***
		(0.016)	(0.011)	(0.011)
	Constant	-0.081***	-0.185***	-0.254***
		(0.022)	(0.024)	(0.038)
	R^2	0.008	0.027	0.003
Δ Routine manual	Auto Technology	-0.010***	-0.022***	-0.003
		(0.003)	(0.004)	(0.004)
	Constant	0.002	-0.058***	-0.095***
		(0.007)	(0.009)	(0.011)
	R^2	0.008	0.021	0.000

Note: The table presents separate OLS regressions for the subperiods 1970-1980, 1980-1990 and 1990-1998, always using as explanatory variable the average change of new automation patents between 1976 and 1998 (divided by 1000). The dependent variable is the change in industry-level task input as calculated by Autor et al. (2003). Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

relation is strongest in the 1970s to 1980s and declines over time. We interpret this as a sign that the nature of automation technology may have changed: While in the 1970s until 1990s, automation technology mostly replaced routine tasks, it nowadays spreads into other tasks. This could be because many routine jobs have already been replaced by automation, so that additional research in this area is less demanded and less profitable. Another possible explanation is that recent advances in the automation technology frontier affect non-routine workers by being able to replace more complex intellectual or manual tasks. (The self-driving vehicle comes to mind.)

To explore this finding further, we examine the effects of technological change separately for routine manual, routine cognitive, non-routine analytic and non-routine interactive tasks. We regress changes in industry task input within each decade on our measure of new automation technology. This is a replication of a regression analysis by Autor et al. (2003), but we replace investment in computer capital with our index. To stay as close to the analysis of Autor et al. (2003) as possible we calculate the left-hand side variable separately for 1970-1980, 1980-1990 and 1990-1998 whereas on the right-hand side, we use the mean of new automation patents over the whole time period from 1976 to 1998.[12]

Table 3.6 shows that *more* automation patents were granted in industries where routine cognitive and routine manual task inputs *declined* and where the share of non-routine analytic and non-routine interactive tasks *increased*. It is noteworthy that for all four task inputs the effect is strongest in 1980-1990. This differs from Autor et al. (2003) who found that for routine tasks the effect had monotonically increased over time.

4. Labor Market Effects of Automation Technology

In this section, we first motivate our unit of analysis, local labor markets, before explaining how we translate our index from industries to U.S. commuting zones. We show graphically how automation across commuting zones changed over time. Then, we apply the derived measure in our econometric analysis of employment effects. In the regression set-up, we rely on fixed effects five-year overlapping time periods, which we explain in detail before discussing the results. We run regressions for the full sample and separately for manufacturing and non-manufacturing employment.

4.1 Commuting Zones as Level of Analysis

We study the effects of automation on employment at the level of U.S. commuting zones. Tolbert and Sizer (1996) have grouped all counties of the U.S. mainland into 722 commuting zones which each exhibit strong commuting ties within, but weak commuting ties between one another. These regions are meant to approximate local labor markets. In response to a shock to labor demand, most adjustments in the short- and medium-run will take place within the local labor market (Blanchard and Katz, 1992, Moretti, 2011). Workers, when laid off, usually first look for a new job within the same commuting zone. This is particularly true for low-skill workers, who are likely to be affected the most by automation (Notowidigdo, 2011). Therefore, studying the effects of automation on employment on the level of commuting zones gives us a more complete picture of the employment effects of automation than an industry-level analysis, which would neglect worker flows from one industry to another. This is of particular relevance because of the substantial shift of employment from manufacturing to services in the sample period.

We use employment data by the *County Business Patterns* (CBP) to convert patents per industry to worker patent automation intensities on a commuting zone level.[13] To create the commuting zone measure of automation, we first take (one plus) the natural logarithm of industry-level automation patents in order to account for the different levels of patenting across industries: In some industries the pace of technological progress is too fast for patents

[12] Results are very similar when we use the whole period that our indicator covers, 1976-2014. Alternatively, we can count only automation patents of the decade for which the change in task input is calculated. The results stay qualitatively the same. Regression outputs are available from the authors upon request.

[13] In this dataset, employment numbers are reported by county and 4-digit SIC (6-digit NAICS) industry. In contrast to Census data, which is sometimes used for commuting zone analysis, CBP provides annual data for the whole period of analysis. Agriculture (SIC < 1000) and public administration (SIC > 9000) are excluded from CBP. To avoid imprecision due to SIC-NAICS correspondences and missing CBP employment data for some particular industries, we aggregate employment and the automation index on the 3-digit SIC level before matching.

to be a feasible way to protect innovations, while in others, inventors have strategic reasons not to file for a patent. We then divide the employment-weighted sum of automation patents by total employment in the commuting zone. The resulting measure is

$$\underbrace{\text{autoint}_{c,t}}_{\text{automation intensity}} = \frac{\sum_i \ln(1 + \text{automation patents}_{i,t}) L_{i,c,t}}{L_{c,t}}, \tag{3.4}$$

where L is employment, i stands for industry, c for commuting zone and t for time period.

Figure 3.5 shows the number of automation patents per worker across U.S. commuting zones in four subperiods: 1976-85, 1986-95, 1996-2005 and 2006-14. The colors represent four quartiles of the distribution of automation intensity (in levels) in these subperiods: dark red color signals the 25 percent of commuting zones with the most patents, white color signals the 25 percent with the least patents. The map thus indicates which commuting zones have a high or low share of patents *relative* to the rest of the United States in the specific subperiod.[14]

There are pronounced regional patterns in the dispersion of available automation technology. Between 1976 and 1995, the region around the Great Lakes had a large automation patent intensity relative to the rest of the United States. This stems from the conjunction of both a high number of patents in manufacturing industries and a large share of industrial employment in this area. Starting in the mid-1990s, many commuting zones in this region move to a lower quartile as the number of manufacturing employees decreased relative to the number of employees in sectors with fewer patents. But our map of automation density is not simply a reflection of the manufacturing share. In a particular the Southern United States

The commuting zones with the highest automation intensities are more dispersed in the 1990s and 2000s. Commuting zones in Montana, North and South Dakota and Nebraska attract many automation patents per employee. The Rocky Mountain region has a low share of patents throughout the whole sample period. The map therefore reveals substantial geographic variation over time, which we exploit in the regression analysis.

[14]As the legend shows, the absolute number of patents has increased across all quartiles. An individual commuting zone may thus have had its absolute number of patents increase constantly over time, but change from dark red to white because the index increased relatively more slowly than in other commuting zones.

Figure 3.5: Intensity of automation patents across commuting zones, 1976-2014

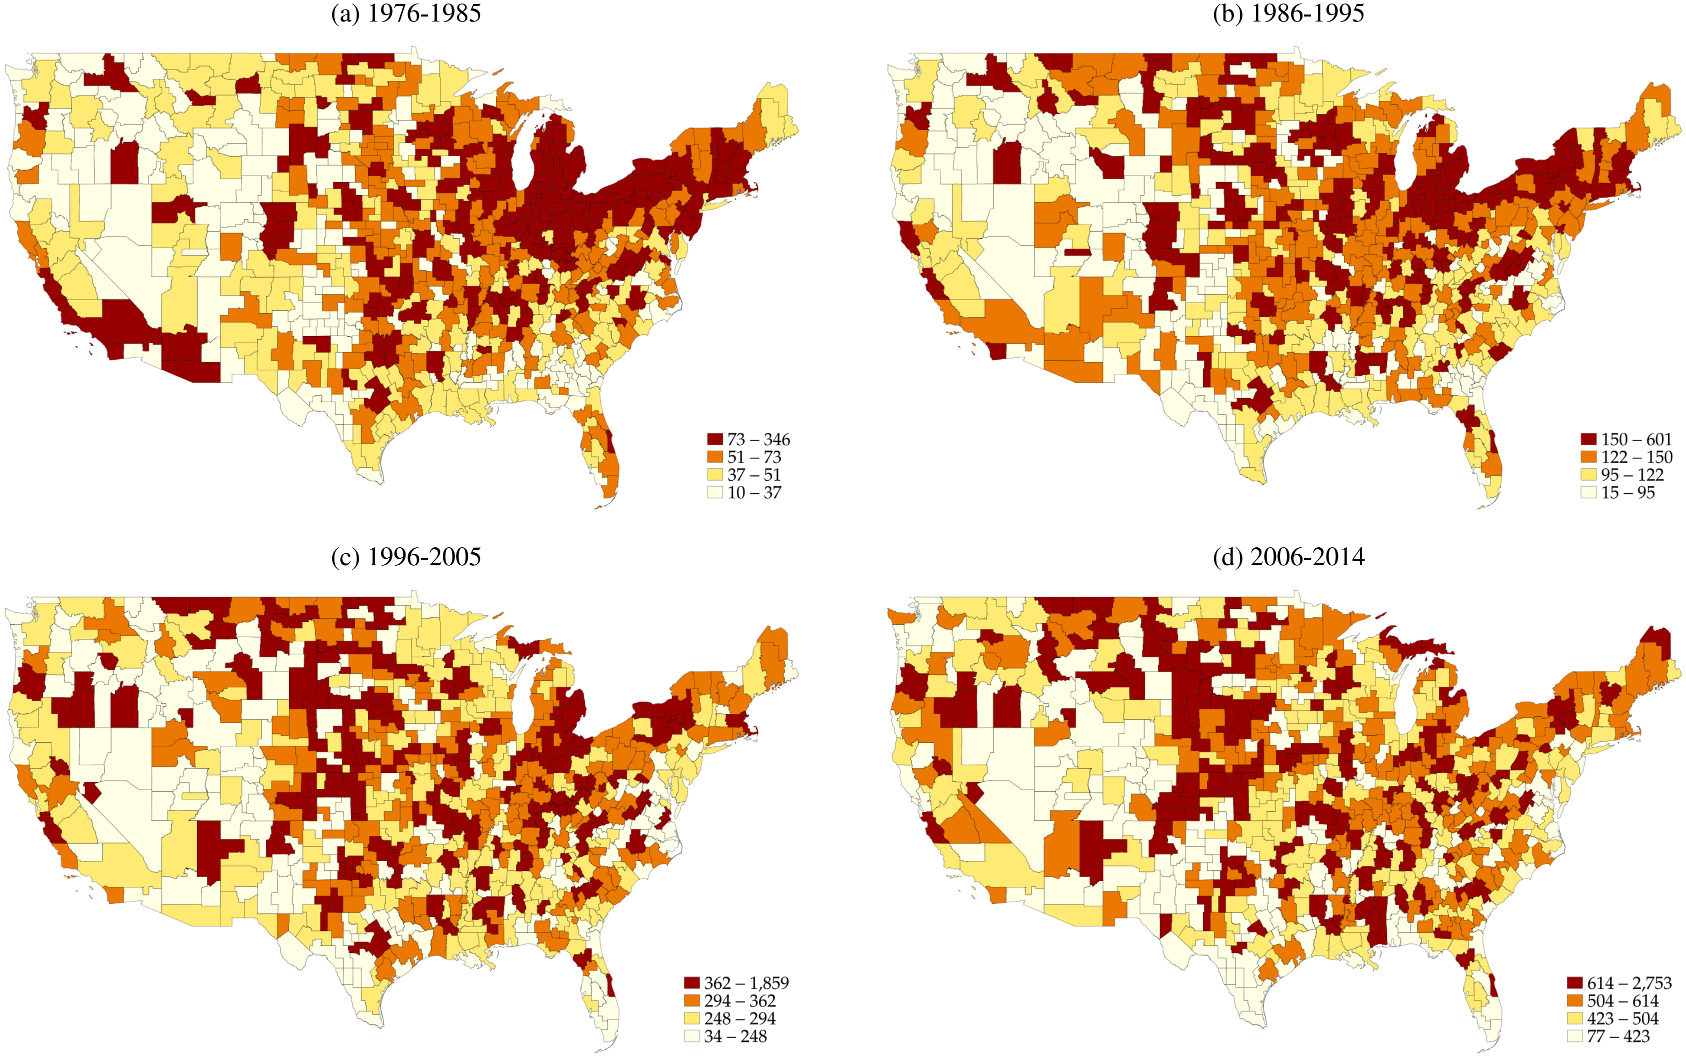

Note: Shows averages of the number of national automation patents that can be used by a single worker.
Source: USPTO, Google, Silverman (2002), CBP and own calculations.

4.2 Empirical Strategy

Our dependent variable is the five-year change in the employment-to-population ratio L_c/pop_c in commuting zone c:

$$\Delta\frac{L_{c,t}}{\text{pop}_{c,t}} = \frac{L_{c,t}}{\text{pop}_{c,t}} - \frac{L_{c,t-5}}{\text{pop}_{c,t-5}},$$

where in contrast to automation, we observe employment directly at the commuting-zone level. We choose a medium term period as new patents might start to be used by firms only with some lags.[15] This also holds the additional benefit of smoothing out business cycle effects.

The main explanatory variable is the five-year sum of the automation intensity in a commuting zone: $\sum_{s=0}^{4} \text{autoint}_{c,t-s}$. By using sums, we interpret patents as a flow measure of technology and therefore, the five-year sum of new patents is the five-year difference in the stock of patents.

In our econometric analysis we ask the following question: What is the impact of newly available nationwide automation technology on changes in the employment structure at the local level? In order to answer this question causally, we need to argue convincingly that our automation measure is exogenous to employment changes. The main potential source of endogeneity is that in their research activity, firms may be reacting to local developments, for example changes in labor costs, regulations or demand, thus introducing a reverse causality bias. There are several reasons why this is less of a concern for us:

Automation by industry of use: Assigning patents to the industries where they are likely to be used, not filed, weakens the danger of reverse causality: The research effort of a firm in one industry is less directly linked to employment trends in another industry than, for example, data on actual investment in automation technology. Additionally, many patents are granted to universities, research institutes or individuals that might follow other objectives than profit maximization, for example intellectual curiosity or an interest in advancing science. These sources of innovation are of relevance, as in year 2000 about 7000 patent licenses to firms were issued by U.S. universities and U.S. public research institutions (OECD, 2003). Further, around half of the patents granted by the USPTO are filed by foreign applicants. This reduces the potential for a feedback from industry wage structure to innovative activity, as a patent from, for example, a manufacturer in Japan is less likely to respond to employment conditions in the manufacturing industry in the United States.

National innovation, local effects: We measure innovations at the level of national industries, whereas we observe employment changes locally. Our constructed commuting zone automation measure is thus a proxy for unobserved locally applicable innovation in the spirit of Bartik (1991), as recently explained by Goldsmith-Pinkham, Sorkin, and Swift (2017). A national industry is unlikely to react to local employment trends in its research activity unless the following conditions hold: First, the specific commuting zone is of key importance to the industry (by hosting a large share of industry employment) and second, the industry is represented strongly in the commuting zone, so that industry trends will translate directly into commuting zone employment trends. These conditions do not drive our findings: In our sample, only two commuting zones are above the 25 percent double threshold (CZ 35002 in

[15] Results are robust to changing the length of a period.

Arizona and CZ 37601 in Nevada, in both of which mining is dominant) and only 34 commuting zones are above the 10 percent double threshold. Excluding these does not significantly change the results.

Fixed industry structure: We fix the employment structure in equation (3.4) to the beginning of each five-year period. This means that in the following five years we assign all patents to a commuting zone according to the initial employment share of relevant industries in this commuting zone. Our indicator thus does not pick up employment changes that happen within the five-year period. A downside of keeping the employment structure fixed is that we potentially do not count all those patents which workers in a commuting zone can use, but might over-represent declining and under-represent growing industries.[16]

Additionally, in Secion 4.6 we exploit information on the owners of patents in order to identify innovations that more likely result from research effort that is unrelated to trends in US labor markets. We show that our baseline regression results hold when focusing only on patents held by foreigners, governments or universities and public research institutes, or when using these as instruments for the patents held by US companies.

4.3 Regression Set-up

We consider changes in overlapping five-year time periods and the sample therefore comprises 34 consecutive five-year periods across 722 commuting zones.[17]

The estimation equation takes the form

$$\begin{aligned}\Delta \frac{L_{c,t}}{pop_{c,t}} = \alpha_k + \gamma_t + \beta_1 \sum_{s=0}^{4} \text{autoint}_{c,t-s} + \beta_2 \sum_{s=0}^{4} \text{non-autoint}_{c,t-s} + \beta_3 \text{routine}_{c,t-5} + \\ \beta_4 \left(\sum_{s=0}^{4} \text{autoint}_{c,t-s} \times \text{routine}_{c,t-5} \right) + X'_{c,t-5} \beta_5 + \varepsilon_{c,t,t-5} \,,\end{aligned} \tag{3.5}$$

where γ_t are time fixed effects and α_k are state fixed effects. $X_{c,t-5}$ are additional control variables. The main variable of interest *autoint* is automation intensity, *non-autoint* is the intensity of any non-automation patents and *routine* is the routine task share which we describe below. To construct the left-hand side variable, we take county level population data from the *Census Population and Housing Unit Estimates* and county-level employment data from CBP. Because the CBP omits employment in some SIC industries for certain years, there are a few large jumps in the outcome variable, which we exclude from the analysis by dropping data below the 1^{th} and beyond the 99^{th} percentile in each year.[18]

In addition to commuting zone intensities of automation patents, we include intensities of non-automation patents (*non-autoint*) in the regression, computed analogously to equation

[16] The results are however robust to using an adaptive industry structure.

[17] The overlapping data structure generates serial correlation. We correct the standard errors by using the Driscoll and Kraay (1998) estimator, which corrects both for serial and spacial correlation. An alternative would be to use non-overlapping time periods. But not only would this mean losing a considerable amount of observations (and thus precision), but it would also require us to choose cut-off points for the five-year intervals, which would always be to some extent arbitrary. As shown in the appendix, all main results go through using this more standard estimation procedure instead.

[18] For details, see `census.gov/program-surveys/cbp/technical-documentation`. The number of commuting zones in each year falls to 708.

Table 3.7: Summary statistics of main variables in baseline regression

Variable	**Mean**	**Overall Std. Dev.**	**Between Std. Dev.**	**Within Std. Dev.**	**Min**	**Max**
Δ emp/pop	1.19	2.71	0.710	2.62	-9.40	13.2
Δ manu emp/pop	-0.342	1.08	0.457	0.977	-5.35	4.33
Δ non-manu emp/pop	1.53	2.19	0.542	2.12	-8.63	12.9
autoint	16.4	3.02	1.23	2.76	7.63	28.6
non-autoint	18.8	1.89	1.38	1.29	8.88	26.7
routine	34.4	5.32	4.25	3.20	8.51	56.3

Note: Variables are as defined in the text.

(3.4). This variable controls for the effect of technological change other than in automation technology. Given that some industries generally patent more, it is likely that the number of automation patents and non-automation patents granted annually are correlated across industries and commuting zones. At the same time, non-automation inventions may also have an independent effect on employment. In particular, they may be interpreted as an indicator for local growth potential, which we might otherwise suspect to be accountable for correlations between automation and employment: If growing industries increase their workforce as well as invest more in R&D, this should be reflected by the coefficient on *non-autoint*.

As described in Section 3, an often-used measure of susceptibility to automation is the routine-task index by Autor et al. (2003). The different construction of this measure from ours creates the opportunity to explore how the effects of these two are related and to ask the question: How does the effect of automation depend on the routine task share of a commuting zone? We therefore include the initial ($t-5$) routine task share (*routine*) in the regression as well as an interaction term between this measure and the variable for automation intensity.

We further include the initial share of manufacturing employment in total employment (CBP) to capture structural change in the economy. Automation patents occur to a larger extent in the manufacturing sector than in the service sector, so an increase in the automation index may parallel a decline in the manufacturing industry for other reasons, such as the cheap import of manufactured goods from abroad or changes in the demand for goods. If not included as a control, any effect stemming from non-automation-related structural change might be attributed to automation technology.

Similar to Acemoglu and Restrepo (2017), our set-up also includes the log of initial commuting zone population because employment in larger and smaller commuting zones – in particular when interpreting this as a proxy for urban vs. rural areas – might react differently to automation. We also control for the share of non-white citizens in the commuting zone population and for the (log of) per capita level of personal income. Data on the demographic variables are taken from the *Census Population and Housing Unit Estimates*, data on income come from the Bureau of Economic Analysis' *Regional Economic Information System* (REIS), which exploits county-level data from administrative records and censuses.

Table 3.7 summarizes the main variables of interest. Employment per population grew on average over the sample period.[19] Employment changes were negative on average for the manufacturing sector and positive for the non-manufacturing sector with more within and across variation for the latter.[20] Our automation intensity measure *autoint* takes the value 16.4 on average across years and commuting zones. This value is equivalent to a commuting zone with a flat industry structure (i.e., all 377 SIC 3-digit industries having the same employment share) where 25 new automation patents are granted every year in all industries. Because patents are skewed across industries, this number will be larger for most industries.

4.4 Estimation Results: Total Employment

Table 3.8 presents the baseline results. Throughout almost all specifications, *autoint* has a significantly positive coefficient in the range of around 0.10 to 0.23 percentage points. So new automation technology per worker is significantly related to employment gains in the same commuting zone. This result is robust to controlling for several economic and demographic variables.

Column (1) shows the positive association between automation and employment when no further controls but time and industry fixed effects are included. The relationship becomes more pronounced when we control for other non-automation patents in column (2). Columns (3) shows our preferred regression specification. The coefficient on *autoint* in column (3) can be interpreted such that a one-unit increase in the automation intensity leads to a 0.178 percentage point increase in the employment-to-population ratio. As laid out in Table 3.7, this is about one sixth of the average five-year increase across all observations. The within-year interquartile range of *autoint* lies between 1.23 and 2.15, so a one-unit increase is well within the range of variation of the sample. In terms of the actual number of new patents that this implies, a one-unit increase in *autoint* around its mean is equivalent to the number of new automation patents in a commuting zone with a flat industry structure rising from 23 to 29 per year.

A particularly interesting result is how automation technology interacts with the routine task share. In the setup with both variables in column (4), the coefficients on automation and on routine-intensity become insignificant. This is likely due to the fact that the variables measure overlapping concepts, as argued in Section 3. However, both coefficients are significant when we include the interaction between the two variables. The negative coefficient on the interaction shows that the magnitude of the effect of automation on employment varies with the level of the routine task share: In commuting zones with more routine labor, automation technology has a less positive effect. The total effect of automation in column (5) turns negative for commuting zones with a routine task share larger or equal to 54.5 percent. The mean of *routine* is 34.4 and in only 0.1 percent of all observations it exceeds 54.5 percent. So, the total effect of automation is positive in the overwhelming majority of commuting zones.

Non-automation patents are not associated with changes in employment. This might be

[19]This is mainly driven by increases in female labor market participation, which rose from 47 percent in 1976 to 57 percent in 2014, peaking at 60 percent in 1999. (See the BLS series LNS11300000, LNS11300001 and LNS11300002.) Male participation rates fell quite monotonously from 78 percent in 1976 to 69 percent in 2014. We take care of these structural long-run changes in the labor market not related to automation through time fixed effects.

[20]We will use "non-manufacturing" and "services" interchangeably, but "non-manufacturing" also includes mining and construction.

Table 3.8: Labor market effects of automation, five-year overlapping time periods

	Outcome: Employment-to-population				
	(1)	(2)	(3)	(4)	(5)
autoint	0.105***	0.222***	0.178**	0.144	0.563**
	(0.0363)	(0.0783)	(0.0853)	(0.0886)	(0.214)
non-autoint		-0.120	-0.0245	0.0249	-0.0170
		(0.0997)	(0.0931)	(0.0920)	(0.0989)
manufacturing			-1.782*	-1.211	-1.177
			(1.016)	(1.082)	(1.121)
population			0.0875	0.0745	0.0525
			(0.114)	(0.108)	(0.102)
income			-1.319***	-1.284***	-1.232***
			(0.351)	(0.347)	(0.338)
non-white			-1.222***	-1.256***	-1.383***
			(0.259)	(0.273)	(0.283)
routine				-0.0257	0.143*
				(0.0161)	(0.0787)
autoint × routine					-0.0109**
					(0.00468)
Observations	24,064	24,064	24,064	24,064	24,064
R^2	0.42	0.42	0.43	0.43	0.43

Note: The table presents fixed effects regressions using five-year changes in employment as percent of commuting zone population as the dependent variable. *autoint* and *non-autoint* are five-year sums of new automation and non-automation technology. *routine* is the initial percentage of routine tasks in commuting zone employment. The initial manufacturing share, the log of initial commuting zone population, the log of initial per capita income and the initial share of non-white citizens in the population are further controls . All regressions include state and year fixed effects and a constant. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

driven by the nature of these innovations. Many non-automation patents are chemical or pharmaceutical and some are patents without any clear applications. In contrast, automation patents are required by our definition to have at least a distantly recognizable application.

The initial manufacturing share has a mildly significant negative coefficient in our baseline setup of column (3), which might capture the part of the secular trend from manufacturing to services that takes place in the five-year periods we study. The population size is not significantly related to employment changes. A higher per capita income negatively predicts employment changes across all specifications. The employment level is generally higher in commuting zones with a higher per capita income. This could be a sign of convergence in employment shares across commuting zones, but could also reflect a reversely causal effect: as personal income is composed to a large extent of labor income, there could be slower employment growth in commuting zones with a higher wage level, because it is more costly to create jobs. A higher share of the non-white population is negatively associated with employment changes.

Our findings thus paint a more positive picture of the net employment effects of automation than Autor et al. (2015), Graetz and Michaels (2015) and Acemoglu and Restrepo (2017), who found negative or insignificant effects of automation on jobs.[21] It is, however, in line with the findings by Gregory et al. (2016), who show that next to a substitution effect on routine-task jobs, automation lower the production costs. Declining goods prices boost product demand, and so new (non-routine) jobs are created. The positive product demand effect trumps the negative substitution effect. Both the positive level effect of automation and the negative coefficient on the interaction term with the routine task share in our regressions support this explanation. By using a broader measure of automation, we can thus extend the knowledge on its employment effects beyond the findings of a literature that focuses on specific types of automation.

4.5 Estimation Results: Sectoral Employment

We further study the effect of automation on different types of employment separately. Table 3.9 shows pointedly different effects of automation technology on manufacturing and non-manufacturing employment.

Panel A consistently shows that manufacturing employment falls when the automation intensity increases. The effect is significant in our preferred specification (3) and when adding the routine task share in column (4). In contrast to the total US population, the group of manufacturing workers experiences job losses - even when controlling for the initial manufacturing share, which itself has a significantly negative effect. The negative employment effect of automation is more pronounced in commuting zones with a higher routine task share, as the interaction term shows. It turns positive only for commuting zones with a routine task share below 20.9 percent. This is only the case for 115 out of 24,058 observations. Panel B paints a very different picture. In non-manufacturing industries, automation has a very robust job-creating effect. The coefficients are twice as large as in Table 3.8. Non-manufacturing occupations are clear beneficiaries from automation in terms of employment numbers. In contrast to Panel A, the routine task share in the commuting zone does not play a significant role for the size of the automation effect.

Related to this, the coefficient on the routine task share also reveals strong differences between manufacturing and non-manufacturing employment. Commuting zones with a lot of routine labor lose more manufacturing jobs, but this is not the case for non-manufacturing employment. This is likely due to the larger share of routine tasks in the manufacturing than in the service sector. These findings may explain why Acemoglu and Restrepo (2017), in their analysis of the impact of robot use on employment, found automation to be harmful for employment and why Graetz and Michaels (2015), using the same dataset, found evidence for skill polarizing effects of robots: Robots are mainly used in the manufacturing sector and indeed 19 out of the 24 industries covered by IRF robot data are manufacturing industries. Other types of automation innovations, in particular those that can be used in the non-manufacturing sector, may have a more positive effect on employment than industrial robots. Indeed, Acemoglu and Restrepo (2017) show that the effect of robots is less negative or even positive in non-manufacturing industries. They also find that computer usage tends to increase the demand for labor.

[21] Section 4.5 sheds light on why this is the case.

Table 3.9: Labor market effects of automation for manufacturing and non-manufacturing employment, fixed employment structure

	(1)	(2)	(3)	(4)	(5)
A. Outcome: Manufacturing employment-to-population					
autoint	-0.0169	-0.0480	-0.173***	-0.200***	0.144
	(0.0176)	(0.0665)	(0.0300)	(0.0300)	(0.0911)
non-autoint		0.0317	0.235***	0.275***	0.240***
		(0.0747)	(0.0299)	(0.0296)	(0.0218)
manufacturing			-2.581***	-2.142***	-2.127***
			(0.587)	(0.617)	(0.656)
population			-0.0335**	-0.0437***	-0.0608***
			(0.0133)	(0.0128)	(0.0149)
income			-0.739***	-0.712***	-0.668***
			(0.206)	(0.206)	(0.201)
non-white			-0.122	-0.150	-0.259
			(0.238)	(0.232)	(0.214)
routine				-0.0200***	0.119**
				(0.00247)	(0.0437)
autoint × routine					-0.00898***
					(0.00243)
Observations	24,058	24,058	24,058	24,058	24,058
R^2	0.21	0.21	0.25	0.25	0.26
B. Outcome: Non-manufacturing employment-to-population					
autoint	0.113***	0.278***	0.372***	0.370***	0.420***
	(0.0344)	(0.0984)	(0.0768)	(0.0799)	(0.147)
non-autoint		-0.169	-0.293***	-0.290***	-0.296***
		(0.112)	(0.0870)	(0.0840)	(0.0894)
manufacturing			0.852	0.883	0.887
			(0.728)	(0.719)	(0.726)
population			0.118	0.117	0.115
			(0.109)	(0.103)	(0.101)
income			-0.612**	-0.610**	-0.604*
			(0.291)	(0.299)	(0.298)
non-white			-1.105***	-1.107***	-1.122***
			(0.178)	(0.188)	(0.194)
routine				-0.00136	0.0186
				(0.0173)	(0.0384)
autoint × routine					-0.00129
					(0.00256)
Observations	24,067	24,067	24,067	24,067	24,067
R^2	0.38	0.39	0.39	0.39	0.39

Note: The table presents fixed effects regressions for five-year changes in manufacturing employment-to-population and non-manufacturing employment-to-population. The control variables are defined as in Table 3.8. All regressions include state and year fixed effects and a constant. Driscoll-Kraay standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

We add to the existing literature by documenting different effects of automation on manufacturing and non-manufacturing employment: Next to a polarization in skills and tasks, automation has lead to a sectoral shift. Manufacturing sector jobs win, while non-manufacturing jobs lose from automation.

The results presented in this and the previous section are robust to weighing patents by how often they have been cited. Patent citations are sometimes used as an indicator of the value of an invention and therefore, giving stronger weight to highly cited patents might paint a more realistic picture of the degree to which a patent is used in the production process. In Tables 3.19 and 3.20 we replicate the regressions presented in Tables 3.8 and 3.9 using a citations-weighted measure of automation, which we explain further in the Appendix. While our sample is thus shortened by several years, we still find a mildly positive effect of automation for total employment and a pronounced disparity between manufacturing and non-manufacturing.

4.6 Effects of Automation by Assignees

Patents contain information on who owns (or "is assigned") a patent. This information is valuable, because it hints on how closely a patentee's research activities are linked to developments in US labor markets. Innovation activity by entities that do not have business interests in US markets is less likely to be influenced by developments on US labor markets. By focusing on new automation technologies that are originating from such groups, we therefore get a cleaner identification.

To classify the patents, we use data by Lai, D'Amour, Yu, Sun, Doolin, and Fleming (2011), who extract the names of assignees from 1976 until 2012 and provide a host of other information about patents and their owners. We focus on patents held by three groups of assignees, who we believe to be less directly responsive to US labor market trends than US companies: foreigners (these can be companies, individuals or public entities), government bodies (US or foreign) and universities and public research institutes.[22]

Research by foreigners can be assumed to respond to developments in their home country rather than in the United States, as long as the following two conditions are met: The company does not operate on a large scale in the United States, and the domestic labor market trends are not linked to US trends. We do not observe if these conditions hold, so the group of foreigners is the most endogenous of the three. Universities and public research institutes conduct more basic research than corporations, so for them, the immediate applicability or profit maximization might only be a distant motivation. Government patents are also unlikely to be motivated by labor market developments, but should rather respond to military buildups, the needs of certain ministries or cycles in budgetary planning.

Table 3.10 shows summary statistics for patents by the different groups of assignees. US firms are the largest group with around 1.9 million patents. The second largest group are foreigners, who hold 1.8 million patents. Based on the classification by Lai et al. (2011), we identify 45 thousand patents that are assigned to governments. The most important assignees in this category are the US Navy with 10,922 patents, the US Army with 6,217 patents, the US Department of Energy with 4,416 patents, the US Air Force with 3448 patents and

[22]These groups are mostly mutually exclusive, but we count foreign governments (a small group) in both the "foreign" and the "governments" category and foreign universities also show up in the foreigners category.

Table 3.10: Assignee summary statistics, 1976-2012

Assignee	Patents (1000s)	Automat (1000s)	Share	Cit.	Cit. (weighted)	Excl.	Length
US firm	1875.7	948.3	51%	12.2	1.24	14%	1012.3
foreigners	1827.8	777.1	43%	7.1	0.78	12%	831.5
universities	115.1	67.0	58%	10.4	1.02	41%	1435.8
governments	44.8	19.1	43%	8.6	0.74	17%	700.8
missing	609.9	187.5	31%	9.7	0.91	9%	653.8

Note: "Automat" are automation patents as described in text. "Cit." are the average number of citations, "Cit. (weighted)" are the number of citations after removing time-subclassification (HJT) means, where subgroups correspond to those of Table 3.17. "Excl." is the share of excluded patents due to being pharmaceutical and chemical patents. "Length" is the average number of lines in a patent document.
Source: Lai et al. (2011) and own calculations.

NASA with 2,823 patents. The largest foreign government institutions owning US patents are French nuclear energy and aviation commissions and the British and Canadian defense ministries. To identify patents assigned to universities, we inspected the 10,000 assignees with the most patents and determined whether they are an university or a public research institute. There are 581 such entities holding a total of 115 thousand patents. The most productive are the University of California (5,400 patents), the Industrial Research Institute of Taiwan (4,289 patents), the Massachusetts Institute of Technology (3897 patents), the Electronics and Telecommunications Research Institute from South Korea (3,606 patents) and the French Institute of Petroleum (2,471 patents). For the remaining 610 thousand patents, we do not know the assignee, as this information is missing in Lai et al. (2011). A casual inspection of these patents suggests that most of these also belong to US firms or individuals.

The automation patents assigned to foreigners, universities or governments may be of a different nature than those held by US firms – not just for their less direct link to economic developments in the United States, but for reasons related to their applicability. We might see different effects of automation on employment if they were not representative of the technology frontier in automation. Table 3.10 shows that patents held by US firms are characterized by a larger share of automation patents and are more widely cited than those held by other patentees. However, automation patents are highly correlated across groups at the industry level, as Table 3.11 shows. Automation innovations by governmental, foreign and university patentees seem to be applicable in similar industries as automation innovations patented by US firms or individuals. This is not the case when considering all patents. So while it is reasonable to assume that patented automation technology is similar across assignee groups, this is not the case for technology in general.

Indeed, the types of patented innovations differ across technology subgroups. As Table 3.17 shows, US firms hold a particularly high share of "Communication & Computer" patents, which contain a large number of automation patents. Foreigners hold fewer pharmaceutical patents, but many mechanical patents and their patents are cited least often. The column "Cit. (weighted)" in Table 3.10 shows that this holds even after controlling for time and subgroup fixed effects. Universities hold many chemical and pharmaceutical patents and few in the "Communication & Computer" category. These patents are also particularly

Table 3.11: SIC-level correlation of patents in assignee subcategories with US companies

Assignee	Patents		Automation	
	year	year & SIC	year	year & SIC
foreigners	0.33	0.33	0.94	0.95
universities	0.35	0.36	0.88	0.88
governments	-0.45	-0.43	0.02	0.04

Note: Numbers show correlations of subcategories with the categories of US firms and missing assignees. "year" indicates that year trends are taken out, "year & SIC" indicates that year and industry trends are taken out.

lengthy. In contrast, governments hold many patents on electric and electronic innovations, and the corresponding patent texts are shorter than those from other assignees.

We replicate our empirical analysis from the previous section in two ways. First, we repeat the panel data regressions of Table 3.8 and Table 3.9, but for *autoint* and *non-autoint* we use the intensities computed from either only university patents, foreign patents or government patents. Second, we use all three automation sub-indicators as instrumental variables for possibly more endogenous category of US companies and non-identified assignees. The purpose of this exercise is to extract only the component of automation that is unrelated to US labor market developments. As we only have assignee data until 2012, we limit our analysis to the period 1976 to 2012.

For university patents, we document positive net effects of automation on employment. The same holds when using all three groups of automation patents as instruments in column (4). It is striking that again none of the effects of automation on total employment is negative. The size of the coefficient in Table 3.8 lies in the middle of the new estimates. Table 3.13 reports separate results for manufacturing and non-manufacturing employment. We find negative effects of automation on manufacturing employment for all assignee groups apart from university patents. All types of patented automation technology lead a rise in non-manufacturing employment. The magnitude of the coefficients again frame the previous estimates. The findings strongly support the results from our baseline analysis and thus show that the earlier findings were likely not biased by endogeneity of the regressors.

While having roughly the same effects on employment, we can detect slight differences between the patent assignee categories. Automation technology patented by universities and public research institutes has the most strongly positive effects on employment and even the manufacturing sector does not lose from this type of technology. The negative employment effects of automation on the manufacturing sector are strongest when we consider only government patents. Why could this be the case? Universities hold many chemical and pharmaceutical patents, while governments patent many electrical and mechanical patents (Table 3.17). But as explained before, we exclude most chemical and pharmaceutical patents and the classification algorithm further extracts only a relevant subset of patents. As Table 3.18 shows, the makeup of the final automation patents does not differ much between those two groups of assignees. Pharmaceutical patents make for 4 percent of university automation patents and 1 percent of government university patent. A more likely explanation is that the innovations by universities and governments differ along other dimensions that we do not measure.

Table 3.12: Labor market effects of automation, various assignee groups

	Outcome: Employment-to-population			
	(1) university	(2) foreign	(3) gov't	(4) IV
autoint	0.410*	0.153	-0.108	0.128*
	(0.217)	(0.128)	(0.223)	(0.0717)
non-autoint	-0.332	0.0145	0.379	0.0344
	(0.238)	(0.144)	(0.252)	(0.0756)
manufacturing	-0.769	-2.017	-2.061*	-1.961***
	(1.217)	(1.203)	(1.058)	(0.377)
population	0.121	0.110	0.113	0.119***
	(0.114)	(0.112)	(0.116)	(0.0232)
income	-1.225***	-1.393***	-1.358***	-1.358***
	(0.342)	(0.369)	(0.370)	(0.192)
non-white	-1.277***	-1.256***	-1.301***	-1.255***
	(0.233)	(0.256)	(0.250)	(0.255)
Observations	22,648	22,648	22,648	22,648
R^2	0.41	0.42	0.42	0.42

Note: All columns replicate column (3) of Table 3.8. In columns (1) - (3), the full automation measure is replaced by automation by universities, foreigners and governments, respectively. The non-automation measure is constructed accordingly. The last column represents an IV regression, where university, foreign and government (automation) patents are used as instruments for the remaining (automation) patents. Standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

5. Reassessing the Literature

With our new dataset we revisit findings from two important papers of the literature on the local labor market effects of automation. We investigate whether our measure of automation predicts different effects for the growth of non-college service sector jobs (Autor and Dorn, 2013) and how the effects of automation compare with those from China import competition (Autor et al., 2015).[23] Apart from gaining additional insights through our new indicator, this allows comparing our results to the findings from the literature using the established routine-share measure.

5.1 Revisiting Autor and Dorn (2013): The Non-College Service Sector and Employment Polarization

Autor and Dorn (2013) address the issue why there has been an increasing polarization in both employment and wages in 1980-2005. They focus on non-college service sector jobs (e.g., cleaners or security guards), which have grown more rapidly than other less-educated and low-paying occupations (such as factory work) and which have experienced wage increases.

[23]Data and replication files for both papers are from David Dorn's website, `ddorn.net/data` (accessed 10.02.2017).

Table 3.13: Labor market effects of automation for manufacturing and non-manufacturing employment, various assignee groups

	(1) university	(2) foreign	(3) gov't	(4) IV
A. Outcome: Manufacturing employment-to-population				
autoint	-0.120	-0.208***	-0.435***	-0.216***
	(0.114)	(0.0314)	(0.128)	(0.0329)
non-autoint	0.157	0.286***	0.518***	0.303***
	(0.145)	(0.0375)	(0.169)	(0.0331)
manufacturing	-1.796**	-2.827***	-2.441***	-2.807***
	(0.672)	(0.652)	(0.693)	(0.171)
population	-0.0399***	-0.0429***	-0.0419***	-0.0321***
	(0.0137)	(0.0138)	(0.0140)	(0.00941)
income	-0.807***	-0.793***	-0.862***	-0.724***
	(0.213)	(0.234)	(0.213)	(0.0746)
non-white	-0.287	-0.130	-0.310	-0.0937
	(0.264)	(0.257)	(0.270)	(0.125)
Observations	22,642	22,642	22,642	22,642
R^2	0.24	0.25	0.25	0.25
B. Outcome: Non-manufacturing employment-to-population				
autoint	0.518***	0.380***	0.354***	0.374***
	(0.170)	(0.112)	(0.125)	(0.0598)
non-autoint	-0.479**	-0.304**	-0.175	-0.314***
	(0.204)	(0.135)	(0.148)	(0.0637)
manufacturing	0.912	0.897	0.337	0.963***
	(0.849)	(0.842)	(0.605)	(0.310)
population	0.157	0.150	0.150	0.147***
	(0.108)	(0.108)	(0.110)	(0.0200)
income	-0.431	-0.611*	-0.506	-0.661***
	(0.304)	(0.311)	(0.314)	(0.167)
non-white	-0.989***	-1.143***	-0.978***	-1.188***
	(0.138)	(0.131)	(0.146)	(0.207)
Observations	22,650	22,650	22,650	22,650
R^2	0.37	0.37	0.37	0.37

Note: All columns replicate column (3) of Table 3.9. In columns (1) - (3), the full automation measure is replaced by automation by universities, foreigners and governments, respectively. The non-automation measure is constructed accordingly. The last column represents an IV regression, where university, foreign and government (automation) patents are jointly used as instruments for the remaining (automation) patents. Standard errors in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

The authors hypothesize that this is due, among other things, to an increase in automation technology: Automation has reduced the demand for routine manual tasks, while increasing the demand for non-routine manual tasks, thus benefiting non-college service sector jobs at the expense of non-college production jobs.

In their empirical analysis, Autor and Dorn (2013) use the routine-task share as a proxy for automation and show that in commuting zones where initially more people worked in routine occupations, there was a larger increase in non-college service employment. In Table 3.14, column (1), we reproduce their finding to the letter.

Table 3.14: Automation and non-college service employment, 1980-2005

	Outcome: 10 × annual change in share of non-college employment in service occupations			
	(1)	(2)	(3)	(4)
routine	0.105***		0.105***	-0.336
	(0.0320)		(0.0284)	(0.230)
autoint		-0.00100	-0.000990	-0.00533**
		(0.000688)	(0.000645)	(0.00227)
routine × autoint				0.0139*
				(0.00695)
Constant	-0.00632	0.0568***	0.0241	0.161**
	(0.0104)	(0.0210)	(0.0202)	(0.0740)
R^2	0.179	0.171	0.185	0.188

Note: 2,166 observations (3 time periods × 722 commuting zones); robust standard errors in parentheses; all models include state fixed-effects and period fixed effects and are weighted by start of period commuting zone share of national population.
*** p<0.01, ** p<0.05, * p<0.1
Source: Own calculations following Autor and Dorn (2013), Table 5.

We then add *autoint*, our new automation intensity measure. The interaction term in column (4) between *autoint* and *routine* is positive and significant: Non-college service jobs rise in commuting zones with a high routine-task share initially *and* where many new automation patents could be used. This is consistent with the model presented by Autor and Dorn (2013) and highlights an important piece of evidence: the presence of those routine jobs that can be easily automated is necessary for the shift of low-skilled employment into the service sector, not the availability of automation technology by itself.

However, the total effect of automation changes from negative to positive only at a routine-task share of 0.38, a number reached by just 2 out of 2,166 observations and the coefficient on *autoint* in columns (2) and (3) is insignificant. So although we found in Section 4.5 that automation creates non-manufacturing jobs, the rise in non-college service jobs depends crucially on the mix between automation and the existence of routine jobs.

5.2 Revisiting Autor, Dorn, and Hanson (2015): Automation vs. Chinese Trade Exposure

Since the 1990's, there has been a strong rise in trade between the United States and China. A number of papers, such as Autor et al. (2013), Acemoglu et al. (2016) and Pierce and Schott (2016), argue that Chinese import competition is responsible for employment losses in those regions where firms reside that are most exposed to it. Autor et al. (2015) investigate

Table 3.15: Labor market effects of automation patents, routine employment share and exposure to Chinese import competition, 1990-2007

	(1)	(2)	(3)	(4)	(5)	(6)
	A. Outcome: Share of employed in workage population					
routine	-0.0481		-0.0369	-0.207		-0.185
	(0.224)		(0.233)	(0.254)		(0.260)
autoint		0.215***	0.206***		0.331***	0.297***
		(0.0670)	(0.0748)		(0.0757)	(0.0792)
Δ (Imports from China to US)/Worker				-0.831***	-0.832***	-0.942***
				(0.215)	(0.181)	(0.221)
	B. Outcome: Share of unemployed in workage population					
routine	-0.0144		-0.0247	-0.00513		-0.0104
	(0.0616)		(0.0653)	(0.0702)		(0.0728)
autoint		-0.0579**	-0.0645**		-0.0926***	-0.0914***
		(0.0255)	(0.0282)		(0.0222)	(0.0285)
Δ (Imports from China to US)/Worker				0.186***	0.249***	0.221***
				(0.0527)	(0.0676)	(0.0612)
	C. Outcome: Share of not in labor force in workage population					
routine	0.0624		0.0616	0.213		0.195
	(0.172)		(0.178)	(0.194)		(0.197)
autoint		-0.158***	-0.141**		-0.239***	-0.206***
		(0.0538)	(0.0608)		(0.0667)	(0.0672)
Δ (Imports from China to US)/Worker				0.645***	0.583***	0.721***
				(0.188)	(0.155)	(0.190)

Note: The table is based on Autor et al. (2015), Table 1, juxtaposing the effect of Chinese import competition and routine biased technological change on 10-year equivalent changes in the employment status of the working-age population. N = 1444 (2 time periods 1990-2000, 2000-2007, 722 commuting zones). All regressions control for the start of period levels of share of employment in manufacturing, share of population that is college educated, share of population that is foreign born, employment rate among females and Census division dummies. Robust standard errors in parentheses are clustered on state. Models are weighted by start of period commuting zone share of national population. *** p < 0.01, ** p < 0.05, * p < 0.1.

whether this “China shock" or automation has a larger impact on U.S. labor markets. They find that while import competition reduces employment in local labor markets, automation – as measured by the routine task share – is not related to employment changes.

We revisit this finding with our dataset. Table 3.15 replicates the baseline analysis of Autor et al. (2015), Table 1, in which the authors regress 10-year equivalent changes in the employment-to-population ratio, unemployment-to-population ratio and non-participation rate among working age adults. The two main variables of interest are the contemporaneous change in Chinese import exposure per worker and the start-of-decade employment share in routine occupations, both of which are being instrumented.[24]

[24]The instrument for the trade variable is imports from China to other advanced economies. For the initial

Columns (1) and (4) of Table 3.15 are exact replications of columns (1) and (3) of Autor et al. (2015), one containing only the initial routine share, the other one both the routine share and the China shock as explanatory variables. In columns (2) and (5), we replace the routine share by our commuting zone automation intensity. While the coefficient on the routine share is always insignificant, our automation measure has a significantly positive effect on the employment share and a significantly negative effects on both share of unemployed workers and the share of workers that are not in the labor force. This even holds when including both *autoint* and the routine task share. Automation patents have positive effects by reducing the unemployment rate and the number of people outside of the labor force, with a larger effect on the latter group.

An additional finding is that while the effect of the routine task share stays insignificant when including the China shock in column (5), the estimates become even more strongly positive when using our automation indicator. The coefficient on the China shock change little when using *autoint* (column (5)) instead of the *routine* (column (4)). This lends further support to the findings of Autor et al. (2015) on the detrimental effect of Chinese import competition, while automation is playing a more positive role now.

6. Conclusion

This paper makes two contributions: First, it provides a new indicator of automation by applying a text classification algorithm to the universe of U.S. patents granted since 1976. Linking patents to their industry of use and, ultimately, to commuting zones, we construct geographical intensities of newly available automation technology. The second contribution is a fresh assessment of the labor market effects of automation. In an econometric analysis, we show that in commuting zones where more newly-invented automation technology becomes available, the employment-to-population ratio increases. At the same time, there is a shift from routine manufacturing jobs towards non-routine service sector jobs. These results hold when we study only patents by universities, governments or foreigners, which are likely less responsive to developments in US labor markets than domestic firms.

While rising employment ratios in response to automation technology are good news, the benefits of automation may be unevenly distributed. We hope that future research will provide more insights in this respect. A more general contribution of this paper is that it pioneers a way of extracting trends in innovation which can also be used to study the effects of other technologies on the economy.

routine task share, Autor et al. (2015) use its 1950 value in all states but the one that contains the commuting zone, weighted by 1950 employment shares. They argue that in this way, they can isolate the stable, long-run differences in the production structure across commuting zones.

Appendices

A. Additional Tables

Table 3.16: Yearly automation and non-automation patents

	#A	#P			#A	#P			#A	#P	
1976	16279	70194	(25%)	*1989*	27928	95565	(35%)	*2002*	77267	167400	(54%)
1977	15433	65215	(26%)	*1990*	25925	90421	(34%)	*2003*	82017	169077	(56%)
1978	15412	66087	(26%)	*1991*	28037	96561	(35%)	*2004*	84372	164384	(58%)
1979	11721	48840	(28%)	*1992*	29165	97472	(36%)	*2005*	69602	143891	(54%)
1980	14937	61815	(28%)	*1993*	30439	98385	(38%)	*2006*	91201	173822	(59%)
1981	15885	65770	(28%)	*1994*	33699	101695	(39%)	*2007*	83196	157331	(60%)
1982	15092	57877	(31%)	*1995*	35135	101431	(41%)	*2008*	86705	157788	(62%)
1983	14546	56863	(31%)	*1996*	40411	109654	(44%)	*2009*	92843	167463	(62%)
1984	17665	67212	(31%)	*1997*	40217	112019	(44%)	*2010*	121163	219835	(62%)
1985	19415	71668	(32%)	*1998*	57293	147577	(46%)	*2011*	126328	224871	(63%)
1986	19515	70867	(32%)	*1999*	58464	153591	(45%)	*2012*	147550	253633	(65%)
1987	24359	82963	(34%)	*2000*	61273	157595	(45%)	*2013*	163112	278507	(66%)
1988	22006	77938	(33%)	*2001*	64796	166158	(46%)	*2014*	178422	301643	(67%)
								total	**2158825**	**4971078**	**(43%)**

Note: #A: number of automation patents as classified by own algorithm; the patent totals #P are reported as counted by us in the patent files. The USPTO reports slightly different numbers for total patent counts on its website, but the difference is below 0.5% in all years.
Source: USPTO, Google and own calculations.

Table 3.17: Assignee's patents across technological categories, 1976-2012

Assignee	Patents (1000s)	Chem-ical	Comm., Comput.	Drugs, Med.	Electr., Electron.	Mech-anical	Miss-ing	Oth-ers
US firm	1875.7	17%	22%	11%	15%	13%	10%	12%
foreigners	1827.8	16%	20%	7%	19%	18%	10%	9%
universities	115.1	23%	12%	31%	17%	6%	6%	5%
governments	44.8	21%	14%	11%	21%	14%	11%	8%
missing	609.9	11%	7%	11%	9%	21%	11%	29%

Note: Technological classifications are based on USPC numbers and aggregated using the scheme by Hall et al. (2001).
Source: Lai et al. (2011), Hall et al. (2001) and own calculations.

Table 3.18: Share of automation patents after excluding patents

Assignee	Patents (1000s)	Chem- ical	Comm., Comput.	Drugs, Med.	Electr., Electron.	Mech- anical	Miss- ing	Oth- ers
US firm	1875.7	2%	21%	2%	8%	5%	6%	3%
foreigners	1827.8	1%	17%	1%	7%	7%	6%	2%
universities	115.1	2%	11%	4%	10%	2%	4%	2%
governments	44.8	2%	11%	1%	10%	4%	6%	2%
missing	609.9	1%	6%	2%	4%	6%	4%	5%

Note: Technological classifications are based on USPC numbers and aggregated using the scheme by Hall et al. (2001). This table excludes all patents based on the selected pharmaceutical and chemical industries as explained in text.
Source: Lai et al. (2011), Hall et al. (2001) and own calculations.

B. Further Robustness Checks

B.1 Patent citations

Not all patents are of the same importance. Scherer and Harhoff (2000) show that the returns on innovation are highly concentrated, with the 10 percent most valuable patents accounting for around 80 percent of realized value. While Griliches (1990) argues that using a large number of patents partly addresses this concern, we can count how often a patent was cited by other patents as an indicator of its value. We use the patent citations files by Lai et al. (2011) until 2009. The number of citations per patents follow a well-known hump-shape, as newer patents are cited less frequently, but the propensity to cite has risen. Also, some industries (such as pharmaceutical and chemical patents) cite many more patents than others (such as electronics). To control for this, we demean citations across years and the broad technology classes defined by Hall et al. (2001). This is the "fixed effect" method proposed by Hall et al. (2001).

We then weight patents by how often they were cited and replicate our analysis. The analysis shows similar results: Manufacturing employment falls and service employment rises when more (citation-weighted) automation patents become available. The baseline effect on all employment becomes insignificant in this specification, but the interaction between automation and routine task share is still significant.

Table 3.19: Labor market effects of citations-weighted autoation patents

	Outcome: Employment-to-population				
	(1)	(2)	(3)	(4)	(5)
autoint	0.0896**	0.177*	0.0337	-0.0226	0.456**
	(0.0339)	(0.0990)	(0.0748)	(0.0834)	(0.212)
non-autoint		-0.0917	0.104	0.182*	0.106
		(0.124)	(0.0885)	(0.0949)	(0.107)
manufacturing			-2.391**	-1.458	-1.264
			(0.887)	(1.080)	(1.171)
population			0.192**	0.171**	0.146*
			(0.0846)	(0.0800)	(0.0771)
income			-1.337***	-1.285***	-1.222***
			(0.389)	(0.380)	(0.378)
non-white			-1.374***	-1.420***	-1.559***
			(0.129)	(0.136)	(0.123)
routine				-0.0417**	0.142*
				(0.0152)	(0.0772)
autoint $\times$ routine					-0.0117**
					(0.00424)
Observations	20,524	20,524	20,524	20,524	20,524
R^2	0.32	0.32	0.34	0.34	0.34

Note: The table replicates the regressions of Table 3.8 but using citations-weighted five-year sums of new automation and non-automation technology. We only include observations until 2009. Citations are adjusted with the Hall et al. (2001) fixed effect method. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.20: Labor market effects of citations-weighted automation patents for manufacturing and non-manufacturing employment

	(1)	(2)	(3)	(4)	(5)
A. Outcome: Manufacturing employment-to-population					
autoint	-0.0198	-0.110	-0.262***	-0.292***	0.0443
	(0.0186)	(0.0762)	(0.0373)	(0.0425)	(0.136)
non-autoint		0.0948	0.332***	0.375***	0.323***
		(0.0884)	(0.0418)	(0.0490)	(0.0488)
manufacturing			-2.880***	-2.374***	-2.237***
			(0.536)	(0.590)	(0.613)
population			-0.0310**	-0.0425***	-0.0595***
			(0.0151)	(0.0146)	(0.0169)
income			-0.792***	-0.765***	-0.721***
			(0.169)	(0.162)	(0.169)
non-white			-0.167	-0.192	-0.287
			(0.224)	(0.216)	(0.175)
routine				-0.0228***	0.106*
				(0.00440)	(0.0592)
autoint $\times$ routine					-0.00826**
					(0.00320)
Observations	20,520	20,520	20,520	20,520	20,520
R^2	0.19	0.19	0.24	0.24	0.26
B. Outcome: Non-manufacturing employment-to-population					
autoint	0.103***	0.295***	0.315***	0.297***	0.405**
	(0.0313)	(0.0722)	(0.0615)	(0.0705)	(0.157)
non-autoint		-0.202**	-0.258***	-0.233***	-0.250***
		(0.0832)	(0.0741)	(0.0799)	(0.0882)
manufacturing			0.519	0.818	0.861
			(0.673)	(0.776)	(0.829)
population			0.218**	0.211**	0.205**
			(0.0832)	(0.0782)	(0.0757)
income			-0.580*	-0.563	-0.549
			(0.332)	(0.334)	(0.336)
non-white			-1.275***	-1.289***	-1.321***
			(0.110)	(0.105)	(0.0997)
routine				-0.0132	0.0281
				(0.0169)	(0.0413)
autoint $\times$ routine					-0.00263
					(0.00296)
Observations	20,529	20,529	20,529	20,529	20,529
R^2	0.26	0.26	0.28	0.28	0.28

Note: The table presents fixed effects regressions for five-year changes in manufacturing employment-to-population and non-manufacturing employment-to-population. Automation and non-automation are citations-weighted. We only include observations until 2009. Citations are adjusted with the Hall et al. (2001) fixed effect method. The other control variables are defined as in Table 3.9. All regressions include state and year fixed effects and a constant. Driscoll-Kraay standard errors in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

B.2 Non-Overlapping Five-Year Periods

Table 3.21: Labor market effects of automation, five-year non-overlapping time periods

	Outcome: Employment-to-population				
	(1)	(2)	(3)	(4)	(5)
autoint	0.154***	0.324***	0.258**	0.246*	0.611***
	(0.0334)	(0.0892)	(0.126)	(0.134)	(0.162)
non-autoint		-0.173**	-0.0776	-0.0610	-0.0825
		(0.0740)	(0.125)	(0.137)	(0.133)
manufacturing			-1.191*	-1.031*	-0.886
			(0.616)	(0.601)	(0.595)
population			0.107***	0.102***	0.0908***
			(0.0256)	(0.0236)	(0.0241)
income			-0.644***	-0.627***	-0.601***
			(0.228)	(0.224)	(0.223)
non-white			-1.215***	-1.232***	-1.281***
			(0.447)	(0.444)	(0.427)
routine				-0.00751	0.132**
				(0.0136)	(0.0507)
autoint*routine					-0.00969***
					(0.00356)
Observations	5,663	5,663	5,663	5,663	5,663
R^2	0.40	0.40	0.41	0.41	0.41

Note: The table presents fixed effects panel data regressions using non-overlapping five-year equivalent changes in employment as percent of commuting zone population as the dependent variable. *autoint* and *non-autoint* are five-year sums of new automation technology and non-automation technology, as defined in the text. *routine* is the initial percentage of routine tasks in commuting zone employment. Further controls are the initial manufacturing employment share, the log of the initial commuting zone employment, the log of initial per capita income and the initial share of non-white citizens in the population. All regressions include state and year fixed effects and a constant. Standard errors clustered at the state level in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

As an alternative to the five-year overlapping regressions presented in the main part of the paper, we show regression results for non-overlapping periods. These are 1977-1981, 1982-1986, 1987-1991, 1992-1996, 1997-2001, 2002-2006, 2007-2011 and 2012-2014, for which we compute five-year equivalents for the last period that covers only three years. The panel therefore comprises 8 time periods and 708 commuting zones. The results are similar to those presented in the main text. The coefficients in Table 3.21 are slightly larger and more significant than those presented in Table 3.8. The effects of automation for the two employment groups of Table 3.22 are also each slightly more positive than those of Table 3.9, but the finding of the contrary effect of automation is strongly supported.

Table 3.22: Labor market effects of automation for manufacturing and non-manufacturing employment

	(1)	(2)	(3)	(4)	(5)
A. Outcome: Manufacturing employment-to-population					
autoint	0.00382	-0.0365*	-0.110***	-0.137***	0.255***
	(0.0102)	(0.0216)	(0.0269)	(0.0297)	(0.0653)
non-autoint		0.0409**	0.164***	0.205***	0.179***
		(0.0187)	(0.0285)	(0.0328)	(0.0390)
manufacturing			-1.588***	-1.209***	-1.055***
			(0.222)	(0.223)	(0.231)
population			-0.00100	-0.0118	-0.0239*
			(0.0117)	(0.0128)	(0.0125)
income			-0.710***	-0.666***	-0.645***
			(0.139)	(0.143)	(0.140)
non-white			-0.149	-0.191	-0.237
			(0.198)	(0.208)	(0.190)
routine				-0.0181***	0.133***
				(0.00402)	(0.0262)
auto*routine					-0.0104***
					(0.00184)
Observations	5,660	5,660	5,660	5,660	5,660
R^2	0.21	0.21	0.23	0.24	0.26
B. Outcome: Non-manufacturing employment-to-population					
autoint	0.137***	0.363***	0.368***	0.382***	0.260
	(0.0317)	(0.0778)	(0.113)	(0.124)	(0.160)
non-autoint		-0.230***	-0.253**	-0.274**	-0.267**
		(0.0612)	(0.104)	(0.121)	(0.123)
manufacturing			0.321	0.115	0.0700
			(0.425)	(0.419)	(0.424)
population			0.105***	0.111***	0.114***
			(0.0246)	(0.0207)	(0.0210)
routine				0.00957	-0.0371
				(0.0138)	(0.0390)
autoint*routine					0.00324
					(0.00239)
income			0.0393	0.0176	0.0101
			(0.172)	(0.178)	(0.176)
non-white			-1.032***	-1.013***	-0.992***
			(0.229)	(0.219)	(0.219)
Observations	5,662	5,662	5,662	5,662	5,662
R^2	0.36	0.36	0.37	0.37	0.37

Note: The table presents fixed effects panel data regressions for non-overlapping five-year equivalent changes in manufacturing employment-to-population and non-manufacturing employment-to-population. See Table 3.8 for variable definitions. All regressions include state and year fixed effects and a constant. Standard errors clustered at the state level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Bibliography

Abel, A. B. (2001). Will bequests attenuate the predicted meltdown in stock prices when baby boomers retire? *Review of Economics and Statistics 83*(4), 589–595.

Abel, A. B. (2003). The effects of a baby boom on stock prices and capital accumulation in the presence of social security. *Econometrica 71*(2), 551–578.

Acemoglu, D., U. Akcigit, and M. A. Celik (2014). Young, restless and creative: Openness to disruption and creative innovations. NBER Working Paper No. 19894.

Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2016). Import competition and the great US employment sag of the 2000s. *Journal of Labor Economics 34*(1), 141–198.

Acemoglu, D. and P. Restrepo (2017). Robots and jobs: Evidence from US labor markets. NBER Working Paper No. 23285.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *Quarterly Journal of Economics 130*(4), 1781–1824.

Alesina, Alberto, V. G. and G. M. Milesi-Ferretti (1994). The political economy of capital controls. In L. Leiderman and A. Razin (Eds.), *Capital Mobility: The Impact on Consumption, Investment and Growth*, pp. 289–328. Cambridge University Press.

Alfaro, L., S. Kalemli-Ozcan, and V. Volosovych (2014). Sovereigns, upstream capital flows, and global imbalances. *Journal of the European Economic Association 12*(5), 1240–1284.

Algan, Y., O. Allais, and W. J. D. Haan (2008). Solving heterogeneous-agent models with parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control 32*(3), 875 – 908.

Allen, H. F. and D. Gale (2000). *Comparing Financial Systems*. MIT Press.

Ameriks, J., J. S. Briggs, A. Caplin, M. D. Shapiro, and C. Tonetti (2015). Long-term care utility and late in life saving. NBER Working Papers 20973, National Bureau of Economic Research.

Ameriks, J. and S. P. Zeldes (2004). How do household portfolio shares vary with age. University of Columbia Working Paper.

Antweiler, W. and M. Z. Frank (2004). Is all that talk just noise? The information content of internet stock message boards. *Journal of Finance 59*(3), 1259–1294.

Attanasio, O., S. Kitao, and G. L. Violante (2007). Global demographic trends and social security reform. *Journal of Monetary Economics 54*(1), 144–198.

Auerbach, A. J. and L. J. Kotlikoff (1987). *Dynamic fiscal policy*. Cambridge University Press.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review 103*(5), 1553–1597.

Autor, D. H., D. Dorn, and G. H. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6), 2121–2168.

Autor, D. H., D. Dorn, and G. H. Hanson (2015). Untangling trade and technology: Evidence from local labor markets. *Economic Journal 125*(584), 621–646.

Autor, D. H., L. F. Katz, and M. S. Kearney (2008). Trends in US wage inequality: Revising the revisionists. *Review of Economics and Statistics 90*(2), 300–323.

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics 118*(4), 1279–1333.

Backus, D., T. Cooley, and E. Henriksen (2014). Demography and low-frequency capital flows. *Journal of International Economics 92*(S1), 94–102.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636.

Balassa, B. (1964). The purchasing-power parity doctrine: a reappraisal. *Journal of Political Economy 72*(6), 584–596.

Barany, Z., N. Coeurdacier, and S. Guibaud (2015). Fertility, longevity and international capital flows. Unpublished working paper.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

Beaudry, P., M. Doms, and E. Lewis (2010). Should the personal computer be considered a technological revolution? Evidence from US metropolitan areas. *Journal of Political Economy 118*(5), 988–1036.

Beck, T., A. Demirgüç-Kunt, and R. Levine (2000). A new database on the structure and development of the financial sector. *World Bank Economic Review 14*(3), 597–605.

Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. V. Reenen (2017). Who becomes an inventor in America? The importance of exposure to innovation. NBER Working Paper No. 24062.

Benigno, G., N. Converse, and L. Fornaro (2015). Large capital inflows, sectoral allocation, and economic performance. *Journal of International Money and Finance 55*, 60–87.

Bernanke, B. (2005). The global saving glut and the US current account deficit. Technical report, Board of Governors of the Federal Reserve System (US).

Bernanke, B. S. and M. Gertler (1989). Agency costs, net worth, and business fluctuations. *American Economic Review 79*(1), 14–31.

Bessen, J. and R. Hunt (2007). An empirical look at software patents. *Journal of Economics and Management Strategy 16*(1), 157–189.

Blanchard, O. J. and L. F. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 1992*(1), 1–75.

Bonfiglioli, A. (2008). Financial integration, productivity and capital accumulation. *Journal of International Economics 76*(2), 337–355.

Brooks, R. (2002). Asset-market effects of the baby boom and social-security reform. *American Economic Review 92*(2), 402–406.

Brooks, R. (2004, August). The Equity Premium and the Baby Boom. Econometric Society 2004 North American Winter Meetings 155, Econometric Society.

Börsch-Supan, A., A. Ludwig, and J. Winter (2006). Ageing, pension reform and capital flows: A multi-country simulation model. *Economica 73*(292), 625–658.

Buera, F. J., J. P. Kaboski, and Y. Shin (2011). Finance and development: A tale of two sectors. *American Economic Review 101*(5), 1964–2002.

Bussiere, M. and M. Fratzscher (2008). Financial openness and growth: Short-run gain, long-run pain? *Review of International Economics 16*(1), 69–95.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008). An equilibrium model of "global imbalances" and low interest rates. *American Economic Review 98*(1), 358–93.

Campbell, J., J. Cocco, F. Gomes, and P. Maenhout (2001). Investing retirement wealth: A life-cyle model. In John Campbell and Martin Feldstein (Ed.), *Risk Aspects of Social Security Reform*, pp. 439–482. University of Chicago Press.

Carroll, C. D. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity 1992*(2), 61–156.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters 91*(3), 312–320.

Carvalho, C., A. Ferrero, and F. Nechio (2016). Demographics and real interest rates: inspecting the mechanism. *European Economic Review 88*, 208–226.

Caselli, F. and J. Feyrer (2007). The marginal product of capital. *Quarterly Journal of Economics 122*(2), 535–568.

Cetorelli, N. and P. E. Strahan (2006). Finance as a barrier to entry: Bank competition and industry structure in local US markets. *Journal of Finance 61*(1), 437–461.

Chang, Y., J. H. Hong, and M. Karabarbounis (2014). Labor-market uncertainty and portfolio choice puzzles. RCER Working Papers 582.

Chinn, M. D. and H. Ito (2008). A new measure of financial openness. *Journal of Comparative Policy Analysis 10*(3), 309–322.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005). Consumption and portfolio choice over the life cycle. *Review of Financial Studies 18*(2), 491–533.

Cogley, T. and T. J. Sargent (2008, April). The market price of risk and the equity premium: A legacy of the Great Depression? *Journal of Monetary Economics 55*(3), 454–476.

Collins, S. M. (2007). Comments on financial globalization, growth, and volatility in developing countries by Eswar Prasad, Kenneth Rogoff, Shang-Jin Wei, and M. Ayhan Kose. In A. Harrison (Ed.), *Globalization and poverty. NBER Conference Report*, pp. 510–516. University of Chicago Press.

Constantinides, G. M., J. B. Donaldson, and R. Mehra (2002). Junior can't borrow: A new perspective on the equity premium puzzle. *Quarterly Journal of Economics 117*(1), 269–296.

Curcuru, S. E., T. Dvorak, and F. E. Warnock (2008). Cross-border returns differentials. *Quarterly Journal of Economics 123*(4), 1495.

Damodaran, A. (2016). Equity risk premiums (ERP): Determinants, estimation and implications–the 2016 edition. SSRN Working Paper 2742186.

Dauth, W., S. Findeisen, J. Südekum, and N. Woessner (2017). German robots - the impact of industrial robots on workers. CEPR Discussion Paper 12306.

De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. *Journal of Political Economy 118*(1), 39–75.

De Nardi, M., E. French, J. B. Jones, and J. McCauley (2016). Medical Spending of the U.S. Elderly. *Fiscal Studies 37*(3–4), 717–747.

De Santis, R. A. and M. Lührmann (2009). On the determinants of net international portfolio flows: A global perspective. *Journal of International Money and Finance 28*(5), 880–901.

Deaton, A. (1997). *The analysis of household surveys: a microeconometric approach to development policy*. World Bank Publications.

Deaton, A. S. and C. H. Paxson (1994). Saving, growth, and aging in Taiwan. *Studies in the Economics of Aging*, 331.

Devereux, M. B. and A. Sutherland (2011). Country portfolios in open economy macromodels. *Journal of the European Economic Association 9*(2), 337–369.

Dohmen, T., A. Falk, B. H. H. Golsteyn, D. Huffman, and U. Sunde (2017). Risk attitudes across the life course. *Economic Journal 127*, F95 – F116.

Domingos, P. and M. Pazzani (1997). On the optimality of the simple Bayesian classifier under zero-one loss. *Machine Learning 29*(2), 103–130.

Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics 80*(4), 549–560.

Edison, H. J., R. Levine, L. Ricci, and T. Sløk (2002). International financial integration and economic growth. *Journal of International Money and Finance 21*(6), 749–776.

Encyclopædia Britannica (2015). *Automaton*.

Fagereng, A., C. Gottlieb, and L. Guiso (2017). Asset market participation and portfolio choice over the life-cycle. *Journal of Finance 72*(2), 705–750.

Fama, E. F. and K. R. French (2002). The equity premium. *Journal of Finance 57*(2), 637–659.

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the Penn World Table. *American Economic Review 105*(10), 3150–82.

Ferrero, A. (2010, May). A structural decomposition of the U.S. trade balance: Productivity, demographics and fiscal policy. *Journal of Monetary Economics 57*(4), 478–490.

Frey, C. B. and M. A. Osborne (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change 114*, 254–280.

Gagnon, E., B. K. Johannsen, and J. D. Lopez-Salido (2016). Understanding the new normal: the role of demographics. Finance and Economics Discussion Series 2016-080, Board of Governors of the Federal Reserve System.

Galindo, A., F. Schiantarelli, and A. Weiss (2007). Does financial liberalization improve the allocation of investment?: Micro-evidence from developing countries. *Journal of Development Economics 83*(2), 562–587.

Galor, O. (1992). A two-sector overlapping-generations model: A global characterization of the dynamical system. *Econometrica 60*(6), 1351–1386.

Geanakoplos, J., M. Magill, and M. Quinzii (2004). Demography and the long-run predictability of the stock market. *Brookings Papers on Economic Activity* (1), 241–307.

Gentzkow, M., B. T. Kelly, and M. Taddy (2017). Text as data. *Journal of Economic Literature*. (Forthcoming).

Gentzkow, M. and J. M. Shapiro (2010). What drives media slant? Evidence from US daily newspapers. *Econometrica 78*(1), 35–71.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2017). Bartik instruments: What, when, why and how. Unpublished manuscript.

Goos, M. and A. Manning (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *Review of Economics and Statistics 89*(1), 118–133.

Gorton, G., S. Lewellen, and A. Metrick (2012). The safe-asset share. *The American Economic Review 102*(3), 101–106.

Gourinchas, P.-O. and O. Jeanne (2013). Capital flows to developing countries: The allocation puzzle. *Review of Economic Studies 80*(4), 1484–1515.

Gourinchas, P.-O. and H. Rey (2007). From world banker to world venture capitalist: Us external adjustment and the exorbitant privilege. In *G7 Current Account Imbalances: Sustainability and Adjustment*, pp. 11–66. University of Chicago Press.

Gourinchas, P.-O. and H. Rey (2013). External Adjustment, Global Imbalances and Valuation Effects. NBER Working Papers 19240, National Bureau of Economic Research, Inc.

Gourinchas, P.-O., H. Rey, and N. Govillot (2010). Exorbitant privilege and exorbitant duty. Technical report, Institute for Monetary and Economic Studies, Bank of Japan.

Graetz, G. and G. Michaels (2015). Robots at work. CEPR Discussion Paper 1335.

Gregory, T., A. Salomons, and U. Zierahn (2016). Racing with or against the machine? Evidence from Europe. ZEW Discussion Paper 16-053.

Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature 28*, 1661–1707.

Guiso, L., M. Haliassos, and T. Jappelli (2002). *Household portfolios*. MIT press.

Guiso, L. and P. Sodini (2012). Household finance: An emerging field. In M. H. Constandinides, G. and R. Stulz (Eds.), *Handbook of the Economics of Finance*. Elsevier Science.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The NBER patent citation data file: Lessons, insights and methodologial tools. NBER Working Paper No. 8498.

Haltenhof, S., S. J. Lee, and V. Stebunovs (2014). The credit crunch and fall in employment during the Great recession. *Journal of Economic Dynamics and Control 43*, 31–57.

Heaton, J. and D. Lucas (1997). Market frictions, savings behavior, and portfolio choice. *Macroeconomic Dynamics 1*(1), 76–101.

Heaton, J. and D. Lucas (2000). Portfolio choice in the presence of background risk. *Economic Journal 110*(460), 1–26.

Hintermaier, T. and W. Koeniger (2010). The method of endogenous gridpoints with occasionally binding constraints among endogenous variables. *Journal of Economic Dynamics and Control 34*(10), 2074–2088.

Holmstrom, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of Economics 112*(3), 663–691.

Hsieh, C.-T. and P. J. Klenow (2007). Relative prices and relative prosperity. *American Economic Review 97*(3), 562–585.

Huggett, M. and G. Kaplan (2016). How large is the stock component of human capital? *Review of Economic Dynamics 22*, 21 – 51.

IMF (2000). Offshore financial centers. IMF Background Paper.

IRF (2014). World robotics: Industrial robots. Technical report, International Federation of Robotics.

Ju, J. and S.-J. Wei (2010). Domestic institutions and the bypass effect of financial globalization. *American Economic Journal: Economic Policy 2*(4), 173–204.

Ju, J. and S.-J. Wei (2014). A solution to two paradoxes of international capital flows. *Economic and Political Studies 2*(1), 3–43.

Judd, K. (1998). *Numerical Methods in Economics* (1 ed.), Volume 1. The MIT Press.

Kiyotaki, N. and J. Moore (1997). Credit cycles. *Journal of Political Economy 105*(2), 211–248.

Kortum, S. and J. Lerner (1999). What is behind the recent surge in patenting? *Research Policy 28*(1), 1–22.

Kortum, S. and J. Putnam (1997). Assigning patents to industries: Tests of the Yale technology concordance. *Economic Systems Research 9*(2), 161–176.

Kose, M. A., E. Prasad, K. Rogoff, and S.-J. Wei (2010). Financial globalization and economic policies. In D. Rodrik and M. R. Rosenzweig (Eds.), *Handbook of Development Economics*, Volume 5, pp. 4283–4359. Elsevier.

Krueger, D. and A. Ludwig (2007). On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare. *Journal of Monetary Economics 54*(1), 49–87.

Kuhle, W. (2017). Demographic change and the rates of return to risky capital and safe debt. In B. Bökemeier and A. Greiner (Eds.), *Inequality and Finance in Macrodynamics*, pp. 177–190. Cham: Springer International Publishing.

Kuhle, W., A. Ludwig, and A. Boersch-Supan (2007, December). Demographic Change, and the Equity Premium . NBER Retirement Research Center Papers NB 07-07, National Bureau of Economic Research, Inc.

Laeven, L. and F. Valencia (2008). Systemic banking crises: A new database. IMF Working Paper 224.

Laeven, L. and F. Valencia (2012). Systemic banking crises database: An update. IMF Working Paper 163.

Lai, R., A. D'Amour, A. Yu, Y. Sun, D. M. Doolin, and L. Fleming (2011). Disambiguation and co-authorship networks of the U.S. patent inventor database (1975 -2010). Fung Institute.

Lane, P. R. and G. M. Milesi-Ferretti (2007). The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of International Economics 73*(2), 223–250.

Levchenko, A. A., R. Rancière, and M. Thoenig (2009). Growth and risk at the industry level: The real effects of financial liberalization. *Journal of Development Economics 89*(2), 210–222.

Levine, R. (2005). Finance and growth: theory and evidence. In P. Aghion and S. Durlauf (Eds.), *Handbook of Economic Growth*, Volume 1, pp. 865–934. Elsevier.

Lewis, E. (2011). Immigration, skill mix, and capital skill complementarity. *Quarterly Journal of Economics 126*(2), 1029–1069.

Lewis, K. K. (1999). Trying to explain home bias in equities and consumption. *Journal of Economic Literature 37*(2), 571–608.

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica 46*(6), 1429–1445.

Lucas, R. E. (1990). Why doesn't capital flow from rich to poor countries? *The American Economic Review 80*(2), 92–96.

Lybbert, T. J. and N. J. Zolas (2014). Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity. *Research Policy 43*(3), 530–542.

Magerman, T., B. V. Looy, and X. Song (2010). Exploring the feasibility and accuracy of latent semantic analysis based on text mining techniques to detect similarity between patent documents and scientific publication. *Scientometrics 82*(2), 289–306.

Malmendier, U. and S. Nagel (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics 126*(1), 373–416.

Manning, C. D., P. Raghavan, and H. Schütze (2009). *An Introduction to Information Retrieval*. Cambridge University Press.

Matsuyama, K. (2004). Financial market globalization, symmetry-breaking, and endogenous inequality of nations. *Econometrica 72*(3), 853–884.

Matsuyama, K. (2007). Credit traps and credit cycles. *American Economic Review 97*(1), 503–516.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*(2), 145–161.

Mendoza, E. G., V. Quadrini, and J.-V. Ríos-Rull (2009). Financial integration, financial development, and global imbalances. *Journal of Political Economy 117*(3), 371–416.

Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*(4), 373–413.

Michaels, G., A. Natraj, and J. van Reenen (2014). Has ICT polarized skill demand? Evidence from eleven countries over twenty-five years. *Review of Economics and Statistics 96*(1), 60–77.

Milesi-Ferretti, G. M. and P. R. Lane (2010). Cross-border investment in small international financial centers. *IMF Working Papers* (10/38), 1–32.

Modigliani, F. and R. H. Brumberg (1954). Utility analysis and the consumption function: An interpretation of cross-section Data. In Kenneth K. Kurihara (Ed.), *Post-Keynesian Economics*, pp. 388–436. Rutgers University Press.

Modigliani, F. and R. H. Brumberg (1990). Utility analysis and aggregate consumption functions: An attempt at integration. In Andrew Abel (Ed.), *The Collected Papers of Franco Modigliani: Volume 2, The Life Cycle Hypothesis of Saving*, pp. 128–197. The MIT Press.

Mody, A. and A. P. Murshid (2005). Growing up with capital flows. *Journal of International Economics 65*(1), 249–266.

Moretti, E. (2011). Local labor markets. *Handbook of Labor Economics 4*, 1237–1313.

Murphy, K. P. (2012). *Machine learning: a probabilistic perspective*. MIT press.

Nagaoka, S., K. Motohashi, and A. Goto (2010). Patent statistics as an innovation indicator. In B. H. Hall and N. Rosenberg (Eds.), *Handbook of Economics of Innovation*, pp. 1083–1127. Amsterdam: North Holland.

Notowidigdo, M. J. (2011). The incidence of local labor demand shocks. NBER Working Paper No. 17167.

OECD (2003). *Turning Science Into Business: Patenting and Licensing at Public Research Organisations*. OECD Publishing.

OECD (2004). Understanding business patenting and licensing: Results of a survey. In *Patents, Innovation and Economic Performance. OECD Conference Proceedings*, pp. 89–110.

OECD (2011). *OECD Guide to Measuring the Information Society 2011*. OECD Publishing.

OECD (2015). Pensions at a glance 2015: OECD and G20 indicators.

Pang, J. and H. Wu (2009). Financial markets, financial dependence, and the allocation of capital. *Journal of Banking & Finance 33*(5), 810–818.

Parente, S. L. and E. C. Prescott (2000). *Barriers to riches*. MIT Press.

Pierce, J. R. and P. K. Schott (2016). The surprisingly swift decline of US manufacturing employment. *American Economic Review 106*(7), 1632–1662.

Poterba, J. M. (2001, April). Taxation and Portfolio Structure: Issues and Implications. NBER Working Papers 8223, National Bureau of Economic Research, Inc.

Poterba, J. M. and A. Samwick (2001). Household portfolio allocation over the life cycle. In *Aging Issues in the United States and Japan*, pp. 65–104. University of Chicago Press.

Prasad, E., R. G. Rajan, and A. Subramanian (2006). Patterns of international capital flows and their implications for economic development. In *Proceedings of the 2006 Jackson Hole Symposium*, pp. 119–158. Federal Reserve Bank of Kansas.

Prasad, E., R. G. Rajan, and A. Subramanian (2007). Foreign capital and economic growth. *Brookings Papers on Economic Activity 1*, 153–230.

Quinn, D. P. and A. M. Toyoda (2008). Does capital account liberalization lead to growth? *Review of Financial Studies 21*(3), 1403–1449.

Rajan, R. G. and L. Zingales (1998). Financial dependence and growth. *American Economic Review 88*(3), 559–586.

Reis, R. (2013). The Portuguese slump and crash and the euro crisis. *Brookings Papers on Economic Activity*, 143–210.

Rodrik, D. (1998). Who needs capital-account convertibility? In S. Fischer, R. N. Cooper, R. Dornbusch, P. M. Garber, C. Massad, J. J. Pollak, D. Rodrik, and S. S. Tarapore (Eds.), *Essays in International Finance*, Number 207, pp. 55–65. Princeton University.

Samuelson, P. A. (1964). Theoretical notes on trade problems. *Review of Economics and Statistics 46*(2), 145–154.

Scherer, F. M. (1984). Using linked patent and R&D data to measure interindustry technology flows. In Z. Griliches (Ed.), *R&D, Patents, and Productivity*. University of Chicago Press.

Scherer, F. M. and D. Harhoff (2000). Technology policy for a world of skew-distributed outcomes. *Research Policy 29*(4-5), 559–566.

Schmookler, J. (1966). *Invention and Economic Growth*. Harvard University Press.

Silverman, B. S. (2002). *Technological Resources and the Logic of Corporate Diversification*. Routledge.

Sørensen, B. E., Y.-T. Wu, O. Yosha, and Y. Zhu (2007). Home bias and international risk sharing: Twin puzzles separated at birth. *Journal of International Money and Finance 26*(4), 587–605.

Te Kaat, D. M. (2016). International capital flows and the allocation of credit across firms. Unpublished Manuscript.

Tille, C. and E. Van Wincoop (2010). International capital flows. *Journal of International Economics 80*(2), 157–175.

Tolbert, C. M. and M. Sizer (1996). U.S. commuting zones and local labor areas: A 1990 update. Economic Research Service Staff Paper No. 9614.

United Nations (2015). World population prospects 2015.

USPTO (2014). Extended year set. All technologies (utility patents) report. Technical report, United States Patent and Trademark Office.

USPTO (2015). Types of patents. Technical report, United States Patent and Trademark Office.

Von Hagen, J. and H. Zhang (2011). International capital flows and aggregate output. CEPR Discussion Papers 8400.

Von Hagen, J. and H. Zhang (2014). Financial development, international capital flows, and aggregate output. *Journal of Development Economics 106*, 66–77.

WIPO (2016). WIPO IP Facts and Figures 2016. Technical Report No. 943E/16.

Wurgler, J. (2000). Financial markets and the allocation of capital. *Journal of Financial Economics 58*(1-2), 187–214.

Zeldes, S. P. (1989). Optimal consumption with stochastic income: Deviations from certainty equivalence. *Quarterly Journal of Economics 104*(2), 275–298.

Zhang, H. (2014). Trade and financial integration, extensive margin, and income divergence. Research Collection School of Economics Working Paper 7.

Zucman, G. (2013). The missing wealth of nations: Are Europe and the US net debtors or net creditors? *Quarterly Journal of Economics 128*(3), 1321–1364.