Firm Heterogeneity and the Macroeconomy

Immo Schott

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Florence, June 2014

European University Institute
Department of Economics

Firm Heterogeneity and the Macroeconomy

Immo Schott

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Examining Board

Prof. Russell Cooper, Penn State University, Supervisor
Prof. Arpad Abraham, EUI
Prof. E.J. Bartelsman, VU University Amsterdam
Prof. Christian Bayer, University of Bonn

Abstract

The three chapters of this thesis contribute to a literature which emphasizes the importance of microeconomic heterogeneity for macroeconomic outcomes. In my work I focus on firm heterogeneity. I investigate the US labor market implications of a drop in the number of new firms, study the cyclical effects on productivity due to limits in the reallocation of capital across firms, and quantify the effectiveness of a policy which attempted to save jobs in Germany by altering firm incentives for lay-offs. The first chapter of this thesis investigates the role of new firms ('start-ups') in the US labor market. Start-ups and young firms grow faster and create more net jobs than older, incumbent firms. Yet since 2007 the number of start-ups in the US has declined by over 20%, accounting for a large part of the persistently high unemployment rate. I claim that this fact is related to the unprecedented fall in the value of real estate. Based on the empirical evidence I construct a model that captures the idea that start-ups require external financing, for which real estate is used as collateral. I calibrate and compute a quantitative competitive industry model with endogenous entry and exit, firm heterogeneity, labor adjustment costs, and aggregate shocks. It generates a 'jobless recovery' similar to what we observed in the US in the aftermath of the 2007-09 recession and is able to explain over 80% of the increase and persistence in unemployment since the recession. The second chapter, joint work with Russell Cooper, studies the productivity implications of frictions in the reallocation of factors. Recent empirical work has shown that misallocation of factors can have sizeable effects on the levels of aggregate output and productivity. We are interested in the question whether these frictions can also produce important cyclical movements. We find that the effects are quantitatively important in the presence of fluctuations in adjustment frictions and/or the cross sectional variation of profitability shocks. These fluctuations depend on higher order moments of the joint distribution of capital and plant-level productivity rather than mean values alone. Even without aggregate productivity shocks, the model has quantitative properties that resemble those of a standard stochastic growth model and match important facts about the cyclicality of reallocation and firm productivity dispersion. The last chapter, joint work with Russell Cooper and Moritz Meyer, studies the employment and productivity implications of short-time work ('Kurzarbeit') in Germany. During the years 2009-10 this policy was intended to provide incentives for firms to adjust labor input by reducing hours per worker instead of firing workers. Using confidential German firm micro data we estimate a model of costly labor adjustment. We use the estimated model to simulate the effects of the policy during the recent recession, trying to quantify in how far the German short-time work scheme reduced the allocative efficiency of the German labor market.

Contents

Chapter 1

Start-ups, House Prices, and the Jobless Recovery

1.1 Introduction

In this paper I argue that the 'jobless recovery' can be explained through lower job creation by start-ups (firms of age zero). Figure 1-1 shows the result of a simple counterfactual exercise. Had employment by start-ups and young firms been equal to its pre-crisis trend, the unemployment rate at the end of 2011 would have been as low as 6.5% instead of 8.5%. The figure also shows that changes in job *destruction* are not driving the jobless recovery. Even with pre-crisis levels of job destruction the unemployment rate would have been almost as high as we observed.

There has been a renewed interest in jobless recoveries due to the slow recovery of the US labor market following the Great Recession: Although GDP growth rates have been positive since the third quarter of 2009, employment has been slow to follow. Only in the first quarter of 2011 did the unemployment rate fall below its end-of-recession level.[1] In the first quarter of 2013, the unemployment rate stood at 7.7%, compared to the 4.8% unemployment rate in the last quarter prior to the recession (Q4 2007). Employment relative to the working age population in mid-2013 was lower than at the height of the financial crisis.

Relatively little is known about who creates - and who destroys - jobs.[2] Every year several hundred thousand new firms are created, providing millions of new jobs.

---

[1] Throughout this paper I use the NBER recession dates for my business cycle classifications.

[2] In an important empirical contributionHaltiwanger *et al.* (2010) show that by controlling for firm age there remains no systematic relationship between firm size and growth.

While not all of those firms succeed, those that do remain strong engines of job growth over the coming years. This highlights the importance of studying the labor market's extraordinary dynamics, resulting from persistent and large heterogeneity across firms: While some firms expand, others contract, firms are born and firms die.[3] At the heart of these dynamics lie start-ups and young firms. Successful start-ups become vibrant young firms which make up the lion's share of net job creation. A consequence of the prominent role of start-ups is that whenever the *inflow* of new firms into the economy is interrupted this has adverse effects on job creation. The result can then be a jobless recovery. I will argue that the 'credit crunch' and particularly the fall in house prices associated to the recent economic crisis has created such an event. Figure 1-2 shows the strong correlation between the number of start-ups and the house prices index for all US-recession episodes since 1980. The figure plots the evolution of the HPI and the number of new firms throughout the recession periods. The goal of this paper is to quantitatively assess the importance of the decline in the value of real estate - a major funding vehicle for business formation - as a key reason behind the low number new firms and persistently high unemployment.

To this end I develop a quantitative model of heterogeneous firms that operate in a frictional labor market. Firms must post vacancies that are filled with endogenous probability. Wages are determined through bargaining between workers and firms. Unproductive firms may exit the economy, while new firms can enter. During recessions firms shed workers and post fewer vacancies, generating a Beveridge-curve relationship between unemployment and vacancies. All agents own a fixed stock of real estate. Entering firms require a one-period loan to finance start-up costs. They obtain this loan from a bank and use their real estate as collateral. Because new entrepreneurs may strategically default, the risk neutral bank efficiently prices interest rates by charging a default premium to compensate for expected losses. In this way changes in the value of collateral feed back to the entry costs of new firms. Adverse conditions on the housing market can constrain the number of start-ups that enter during a recovery. This link between house prices and firm entry can explain why job creation by start-ups decreased *before* the beginning of the recession in 2007 - a fact that previous models were unable to address. My model generates jobless recoveries if low collateral values prevent some start-ups from entering. Since start-ups have hiring rates over-proportional to their share of output, the link between entry and real estate breaks the strong co-movement in output and unemployment observed in

---

[3]Over the last 35 years the average number of gross jobs created was around 16 million per year, while 14.4 million jobs per year were destroyed. This respectively corresponds to 17% and 15% of the entire labor force. In other words, over 30% of the labor force is reallocated in a given year.
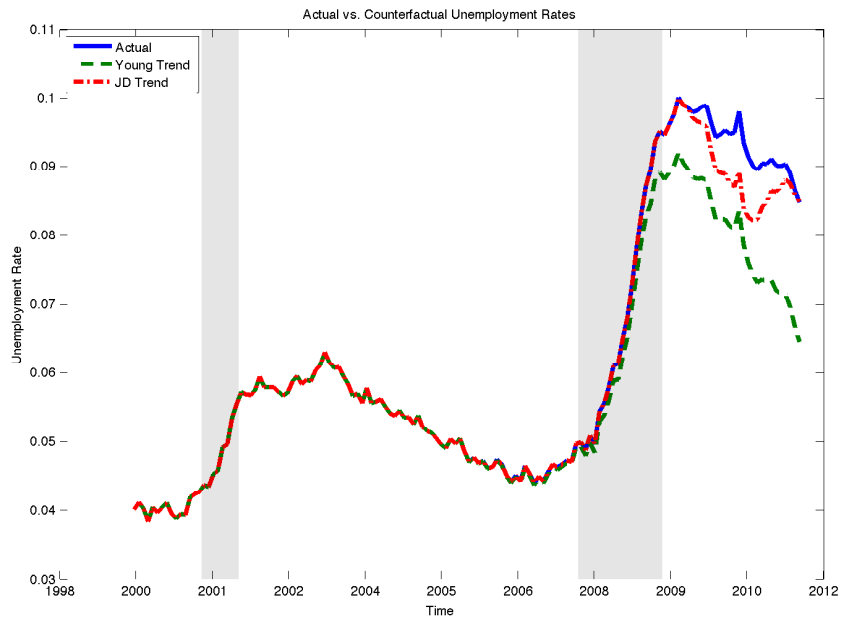
Figure 1-1: The actual unemployment rate is plotted in as the blue solid line. The remaining lines show the counterfactual unemployment rates for the following experiments: The green dashed line labeled 'Young Trend' shows unemployment if gross job creation by young firms (age 5 or below) had been equal to its pre-2006 HP-trend. For the red dashdotted line 'Trend JD' I set gross job destruction (JD) after 2009 equal to its pre-2006 HP-trend. Source: Census, BLS, own computations

otherwise similar models. Additional propagation comes through labor adjustment costs which are chosen to match key moments of the employment change distribution.

1.2 Literature Review

Standard models of the labor market are unable to generate jobless recoveries and sufficient volatility in unemployment and vacancies. The RBC model cannot generate jobless recoveries because shocks are only to aggregate TFP. After a negative shock the reversion to the unconditional mean of TFP increases the marginal benefit of all factor inputs. The Mortensen and Pissarides (1994) search model suffers from the same shortcomings. Furthermore, as pointed out by Shimer (2005) it is unable to generate the volatility in unemployment and vacancies we observe in the data. The competitive industry model (Hopenhayn (1992) and Hopenhayn and Rogerson (1993), henceforth (HR)) introduces entry of new firms and therefore appears as a natural starting point for studying start-ups. The HR setup is a frictionless model in which a market-clearing wage is found via the free-entry condition. The general equilibrium effects induced by this condition are quite powerful in this environment, virtually eliminating any selection effects that could result from the composition of entering and exiting firms (see e.g. Lee and Mukoyama (2012)). I therefore depart from the basic HR model in the following respects. First, I add aggregate shocks to the model since I am interested in the business cycle implications of the model. Second, I add a search-and-matching framework where firms fill vacancies with endogenous probability. This allows me to study the implications of the model for unemployment and vacancies and creates a link between new and incumbent firms through labor market tightness. Third, labor adjustment costs are added to the model in order to match the employment change distribution. Finally, I assume that start-ups must borrow the costs of entry. Potential entrepreneurs use real estate to collateralize a fraction of this loan. As the value of housing falls, the interest rate new entrepreneurs pay on the loan increases. This raises their costs of entry and deters some entrepreneurs from entering. Making entry a function of house prices has several advantages. First, there is empirical evidence on the sensitivity of young firms' hiring behavior with respect to conditions on the credit market. Second, a model with a standard free-entry condition which is not a function of the house price generates entry rates exhibiting excess volatility with respect to the data. The additional dependence on a slow-moving process such as the value of collateral is succesful in generating a realistic volatility of entry. Since the focus of this paper lies on entry, achieving realistic entry rates is crucial. An important assumption of my model is that only new firms need to borrow their overhead costs.
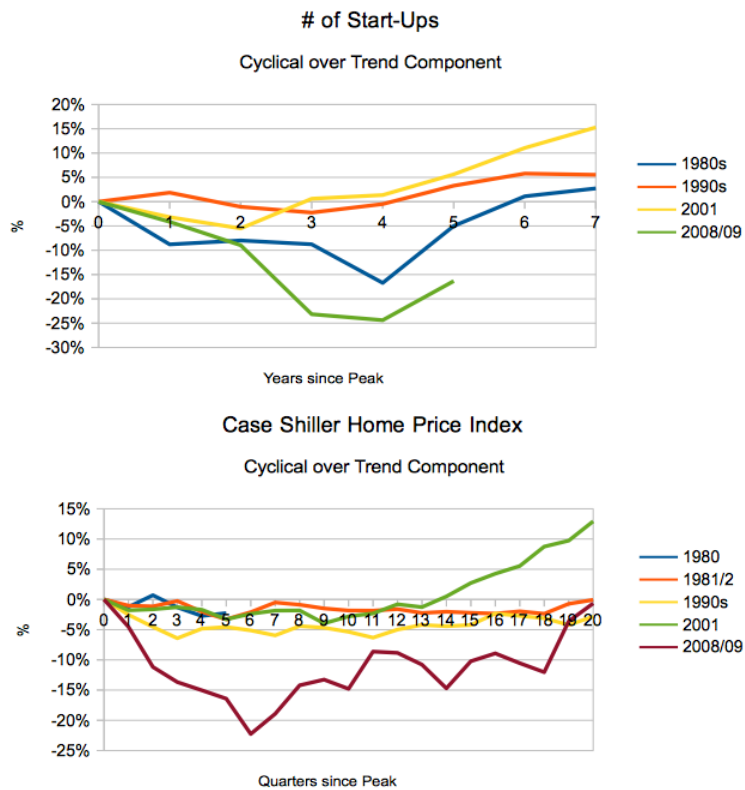
Figure 1-2: Source: BDS and Cash Shiller Home Price Index. HP-filtered. The x-axis shows years/quarters since the respective pre-recession quarter (based on NBER classification).

This is motivated by results of the Survey of Consumer Finances (SCF) which shows that real estate and other personal resources are an important source of business *formation,* but play a much smaller role for expansions of *existing* businesses.

My model is then calibrated to match certain cross-sectional data moments, such as the unemployment-vacancy ratio and the firm age- and employment change distributions. I estimate firm-level labor adjustment costs via a simulated method of moments (SMM) approach. The calibrated model can replicate the average life cycle of firms, the positive correlation between productivity and age, and the negative correlation between employment growth and size observed in the data. I find that the model with house prices affecting credit conditions significantly outperforms alternative specifications, particularly because of its ability to generate jobless recoveries. I perform various policy experiments showing that around 80% of the increase and persistence in unemployment since the end of 2006 can be explained by a model with aggregate TFP shocks and changes in the house price index.

This paper contributes to the literature on the role of start-ups over the business cycle, the impact of financial conditions on the real economy, and jobless recoveries. At the basis of the model lies a heterogeneous-firm framework as in HR, to which I propose the extensions discussed above. An important one is the combination of heterogeneous firms with a standard Mortensen and Pissarides (1994) search-and-matching structure. Other papers that have extended the search framework to multi-worker firms include Cooper *et al.* (2007), Kaas and Kircher (2011), Elsby and Michaels (2013), and Acemoglu and Hawkins (2013). In Cooper *et al.* (2007) labor adjustment costs are estimated in a heterogeneous firm model with search frictions but their framework does not allow for entry and exit. Kaas and Kircher (2011) augment a simplified HR framework with competitive search. Their model can generate sluggish movements of unemployment following a boom but they rely crucially on a time-varying entry cost and the convexity of the recruiting cost function. Furthermore, firms in Kaas and Kircher (2011) are ex-ante heterogeneous, while in my paper they are ex-ante homogeneous and productivity evolves over time. Elsby and Michaels (2013) introduce the Stole and Zwiebel (1996) bargaining framework to the multi-worker firm environment but do not study entry.

A second important extension to the HR model is the financing friction for new businesses. The paper which is most closely related in this respect is Siemer (2013). Siemer develops a heterogeneous firm model with entry and exit based on Khan and Thomas (2013) in which all firms borrow through optimal lending contracts with financial intermediaries. A financial shock overproportionally increases borrowing

costs for small and young firms and reduces entry. The main difference of my model is that it generates jobless recoveries, i.e. underproportional employment growth during recoveries. While in Siemer's model the financial shock produces a 'missing generation' of entrants, I model a link between the hiring conditions of incumbents and entrants through the endogenous labor market tightness θ. This implies that during a recovery firms benefit from an initially low θ, which increases hiring and lets entry rates return relatively quickly to their pre-recession value. In Siemer's model the mass of entrants is a 1:1 mapping of the financial shock, implying that entry levels jump back to their unconditional mean once the financial shock has passed. In the data, however, we observe that entry rates continued to be at historically low levels even *after* financial conditions in the US had returned to "normal", as measured e.g. by various financial stress indeces. In my setup I model entry costs as a function of the value of house prices (collateral). This helps me to explain why entry rates decreased *prior* to the recent recession, why they continue to be low relative to their pre-recession trend, and why job creation by incumbent firms recovered *before* job creation by start-ups.[4] Other related work focusing on start-ups includes Coles and Kelishomi (2011), Clementi and Palazzo (2010), and Lee and Mukoyama (2012). Coles and Kelishomi (2011) study single-worker firms with a two-stage entry process. They show that thus replacing the free entry condition in the standard matching framework significantly enhances the aggregate properties of the model. Lee and Mukoyama (2012) study the cyclical properties of entrants vs. exiters but rely on an entry cost parameter which is exogenously pro-cyclical. Clementi and Palazzo (2010) replace the free entry condition of a standard competitive industry model with a fixed mass of potential entrants and show that entry and exit can propagate the effects of aggregate shocks.[5] Using a standard free-entry condition Hawkins (2011) finds the opposite result. However, he overstates the cyclicality of entry. To the best of my knowledge the previous literature on heterogeneous firms has not been succesful in finding an entry specification that allows for cyclicality in start-up job creation without misspecifying its cyclicality (see e.g. Clementi and Palazzo (2010), Hawkins (2011), Lee and Mukoyama (2012), Berger (2012)). The connection between entry costs and the value of real estate helps to smooth entry rates considerably over the business cycle and generates a realistic degree of fluctuations. Papers which study the link between entrepreneurship and housing collateral empirically are Fort *et al.*

---

[4]House Price Indeces such as the All-Transactions House Price Index for the United States by the FHFA clearly shows a decline *prior* to the end of 2007. The HPI in Q1 2013 stood at 86% of its Q4 2007 value.

[5]Sedlacek (2011) uses a reduced form specification of the free-entry condition to obtain realistic entry dynamics and reproduce key facts of US firm dynamics.

(2013) and Schmalz *et al.* (2013). Fort *et al.* (2013) estimate a VAR and conclude that the collapse in housing prices accounts for a significant part of the large decline in job creation by young firms during the recent recession. Liu *et al.* (2013*a*) also find a signifcant effect of house prices on unemployment. Schmalz *et al.* (2013) empirically link house price shocks to entrepreneurial activity and employment in new firms.

Following the seminal publications by Kiyotaki and Moore (1997) and Bernanke *et al.* (1999) there now exists a vast theoretical literature on the linkages between the financial sector and the real economy. The impact of credit constraints on macroeconomic outcomes has been studied both in the context of search-and-matching and heterogeneous firm models.[6] A large number of theoretical and empirical papers has found a sizeable effect of credit conditions on the real economy during the recent recession (see e.g. Jermann and Quadrini (2012), Gilchrist and Zakrajšek (2012), and Chodorow-Reich (2013)), but these models do not study entry and exit. In my model start-ups need to borrow in order to pay the entry costs, making firm entry a function of credit conditions. A similar mechanism is modeled in Liu *et al.* (2013*b*), where land prices enter a firm's collateral constraint. As in Chaney *et al.* (2012) they find that variations in the collateral value have significant effects on investment.

The jobless recovery has been the topic of Galí *et al.* (2012), Drautzburg (2013), Bachmann (2011), and Berger (2012). In Berger's model firms lay off unproductive workers during recessions. Differently from my paper, the focus of Berger (2012) is on the intensive margin of job creation. While my mechanism is otherwise complementary to Berger's, I show that introducing financing costs for entrants can not only generate jobless recoveries, it also significantly contributes to limiting the volatility of the entry rate. Drautzburg (2013) models an occupational choice problem and estimates that approximately one third of the change in start-up job creation following the recent recession can be attributed to higher risk. Bachmann (2011) explains the jobless recovery through a combination of adjustment costs which generate a jobless recovery after a short and shallow recession. For more severe recession episodes such as the 2008/09 case the model cannot reproduce the observed dynamics, however. Galí *et al.* (2012) argue that the 2008/09 downturn only produced a *quantitative* change in the relation between GDP and employment. However, by comparing the trajectories of GDP, unemployment, job destruction, the house price index (HPI), and start-up job creation for different recession episodes it becomes clear that series differ substantially compared to the other post-1980 recessions. I show those series in

---

[6]Credit constraints in a standard search-and-matching framework were studied by Dromel *et al.* (2010) and Petrosky-Nadeau (2013), who find that the presence of constraints can impact both the level and the persistence of unemployment. Financial constraints have first been introduced into heterogeneous firm models by Midrigan and Xu (2014), Khan and Thomas (2013) and Siemer (2013).

Figures -9 and -10 in Appendix A.1. In particular the link between the HPI and start-up activity (Figure 1-2) stands out as a particular feature of the 2009/09 recession, as the next section shows.

1.3 Facts

This section presents some stylized facts about job destruction and job creation, enterprise dynamics, firm survival, and the link between credit conditions and start-ups. Throughout this paper I will refer to firms of age zero as start-ups or entrants, while firms aged one to five years will be referred to as young firms. A start-up is defined as a new *firm*, not as a new establishment. Unless otherwise noted the data comes from the US Census' Business Dynamics Statistics (BDS) database. Details regarding all the data used in this paper can be found in the Data Appendix. Robustness checks and additional information about the stylized facts can also be found in the appendix.

1.3.1 Firm Dynamics

The 2008/09 recession produced the largest drop in employment since the beginning of the Census' BDS database in 1977. This was the result of both an increase in gross job destruction and a decrease in gross job creation. Persistently low job creation rates are what made the recovery 'jobless'. In 2008/09 fewer jobs were destroyed than during the 2001 recession.[7] Most of it took place on the intensive margin, that is through downsizings of existing firms. Firm deaths only contributed to around 18% of all gross job destruction since 2008.[8] On the other hand, the years 2008 and 2009 marked the largest decreases in gross job creation in the entire Census data. This is summarized in the first *Stylized Fact*. It is robust to employing alternative data sources as I show in the appendix.

Stylized Fact 1: ***The Great Recession was mainly a crisis of low job creation.***

Start-ups play a crucial role for the US economy. The main reason for this is their contribution to job creation. While start-ups create around three million new jobs each year the net contribution of incumbent firms is typically negative. The

---

[7]This holds both in absolute numbers and for the HP-filtered cyclical component. See Appendix.

[8]The average since 1977 was 17.66%. A similar point can be made for establishment deaths. The fraction of gross JD from establishment deaths since 2008 was 30.53%, the average since 1977 was 35.38%.

recent recession has left its mark: While net job creation by incumbent firms quickly recovered since the end of the recession, job creation by start-ups in 2011 was at its lowest point since the beginning of the Census BDS series in 1977.[9] At the same time the average size of a start-up has virtually remained unchanged at around 6 employees. This suggests an important extensive margin effect: fewer entrepreneurs start a business. The drop in start-up hiring stands out as a main factor for low job creation since 2008. While gross job creation remained low for all firm ages after the recession trough, the largest decreases in gross job creation occured among start-ups, followed by the youngest firms. Specifically, in 2011 start-ups created about 700,000 fewer new jobs than in 2007. This is a feature of the 'Great Recession' we do not observe to this extent for the other recessions covered by the BDS data. Figure 1-3 compares changes in absolute gross job creation by firm age relative to the respective pre-recession year across different recession episodes.[10] During the 1980 recessions start-up employment initially increased. During the 2001 recession it decreased but quickly rebounded. The 2008/09 recession was different: Not only was there an unprecedented fall in job creation by start-ups and young firms, this decline persisted even after the official end of the recession. Only the recession in the early 1990s bears similarity to the 'Great Recession' in the sense that hiring by young firms decreased and remained low until several years after the recession trough. The magnitude of this effect is smaller and the relative effect on start-ups is weaker than in 2008/09, however. Interestingly, apart from the recent downturn the 1990/91 epsiode was the only recession where house prices were below trend for a prolonged period of time as Figure -10 in Appendix A.1 shows. These stylized facts summarize the above results:

Stylized Fact 2: *The decrease in job creation was largely due to lower job creation by start-ups and young firms.*

Start-ups have employed around 3 million jobs per year since 1977. Coles and Kelishomi (2011) have pointed out, that this number has been relatively inelastic over the cycle.[11] Using the most recent data the correlation between (the cyclical components of) GDP and job creation by start-ups is 0.356. Job creation by incumbent

---

[9]This result holds across regions and sectors, as is also discussed in Haltiwanger *et al.* (2012). They also note, however, that states that were hit hardest by the financial crisis suffered larger decreases in startup employment, a point that I will take up further below.

[10]Data for all available age groups is shown. Choosing different base years leaves results virtually unchanged. Furthermore, qualitatively identical results can be obtained by plotting job creation *rates* or the cohort's fraction of total job creation (available upon request).

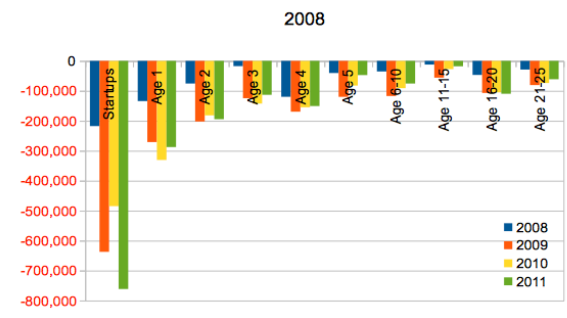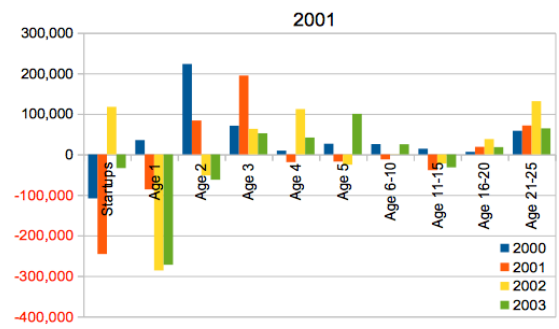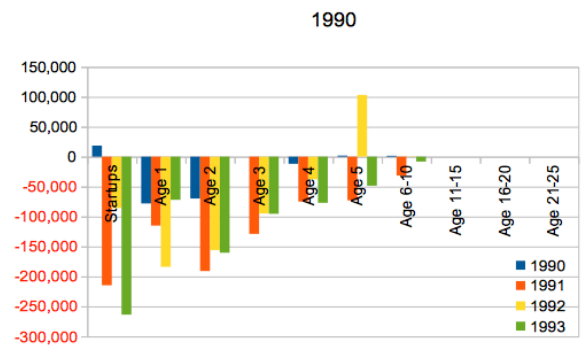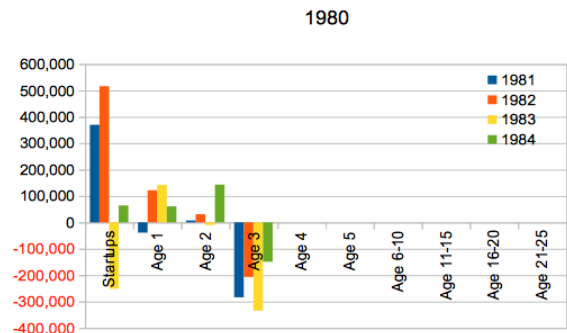[11]See Figure -6 in the Appendix for an updated version of a graph used in Coles and Kelishomi (2011).

Figure 1-3: The y-axis shows changes in gross job creation relative to base years 1979, 1989, 1999, and 2007. For age group bins averages are shown. Source: BDS.

firms, on the other hand, has a higher procyclicality (0.756).[12]

Stylized Fact 3: *Employment in incumbent firms is more strongly procyclical than employment in start-ups.*

However, while start-up job creation is less correlated to fluctuations in GDP, it nevertheless shows more volatility over time than gross job creation by established firms. The standard deviation of the cyclical component of job creation over its trend is about 40% larger for entrants than for incumbent firms (0.10 vs. 0.07).

Stylized Fact 4: *Job creation by start-ups is more volatile than job creation by incumbents.*

I divide firms into four size categories, 1-19, 20-99, 100-499, and 500+ employees. The size and age distribution of firms and establishments can be seen in Tables 3 and 4 in the Appendix. It is noteworthy that while over 95% of firms have less than 100 employees, it is large firms that employ almost half of the workforce. The average firm size is 21.43 workers. The distribution of start-ups shows that the vast majority of start-ups (98.1%) are small firms with less than 20 employees.[13] The age distribution of firms shows that start-ups make up about 11% of all firms. This is an important statistic that my model is going to match. Start-ups and young firms show overproportional employment growth: Start-ups employ around 3% of the labor force, yet contribute 18.7% of total job creation. On the other hand, young firms show higher-than-average rates of job destruction. A significant fraction of which is the result of firm exit. These 'up or out dynamics' were first described by Haltiwanger *et al.* (2010). Conditional on survival, young firms grow considerably faster than their mature counterparts. As Figure 1-4 shows, in the BDS data around 50% of gross job destruction is the result of firm exit during the first three years of a firm's life. On the other hand, for firms older than 20 years less than 15% of all gross job destruction is the result of firm exit. The total firm exit rate is 8.8% per year.

---

[12]These numbers change only marginally by considering alternative subsets of the data. For example the correlation between GDP and job creation by start-ups between 1982-2007 is 0.33 and for gross job creation by incumbents it is 0.72. The correlation between GDP and employment in incumbent firms is 0.5002.

[13]Very large start-ups are rare and should be treated with caution, as practise shows they are often temporary entities that get folded into other firms later on.
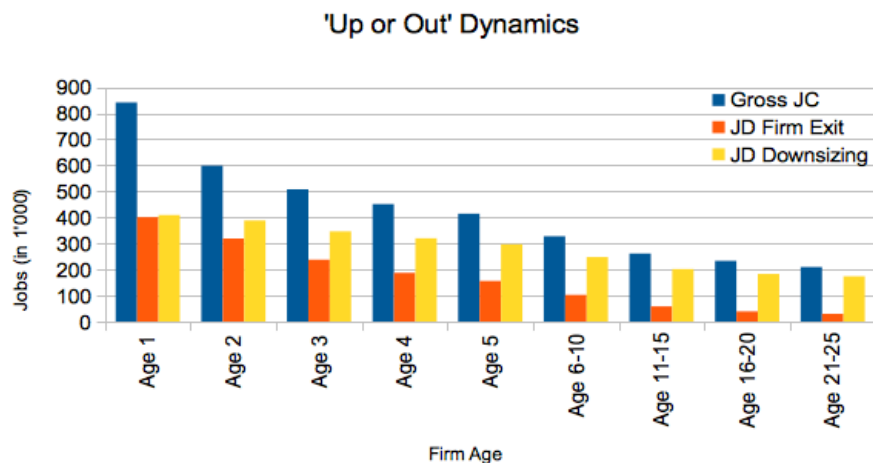
'Up or Out' Dynamics

Figure 1-4: Gross Job Creation and Job Destruction from Exit and Downsizing by Age. Source: Census 1977-2011. Own computations.

1.3.2 Housing and Credit supply

In the wake of the financial crisis there have been numerous initiatives to monitor credit conditions for small businesses.[14] This section will show that after 2007 start-ups have been facing higher costs of obtaining credit. This is important because besides personal wealth, banks are the most important source of funding for start-ups. I present evidence that the sharp drop in the value of real estate, which is a predominant source of collateral for business formation, is connected to the ongoing difficulty for start-ups to obtain financing. State-level regression results indicate that changes in the value of real estate have a negative effect on the number of start-ups, even when controlling for macroeconomic conditions and demand effects.

Start-ups and young firms rely heavily on external liquidity but they face a different initial lending environment and more challenges than mature firms in obtaining credit.[15] Start-ups do not have an established credit record and typically face restric-

---

[14]Those include the reports issued by the Congressional Oversight Panel for the 'Troubled Asset Relief Program' (TARP) ({United States Congressional Oversight Panel} (2010) and {United States Congressional Oversight Panel} (2011)), surveys by the National Federation of Independent Business (NFIB) summarized in Dennis Jr. (2010), Dennis Jr. (2011), and Dennis Jr. (2012), and the proceedings of the annual conference of the US Securities and Exchange Commission (SEC) on 'Small Business Capital Formation' (2009), which can be found at http://www.sec.gov/info/smallbus/sbforum.shtml. Furthermore the Federal Deposit Insurance Corporation (FDIC) has increased the periodicity of its 'Quarterly Banking Profile'. See also Bassett *et al.* (2011) for a discussion of the data and http://www.federalreserve.gov/events/conferences/2010/sbc/agenda.htm for a forum on this topic organized by the Federal Reserve in July 2010. See also evidence in Foundation (2013), the Survey of Small Business Finance, and the Statistics of Business Owners.

[15]See Board (2011), Siemer (2013), Robb and Robinson (2012), the Survey results by the Federal

tions in their access to commercial bonds or other means of financing available to older firms. Why would banks not lend (enough) to start-ups? One reason is the general deterioration in the lending environment of firms.[16] More important, however, was the decline of the value of real estate and household net worth, which acts as collateral for initial loans. According to Avery *et al.* (1998) loans having a personal guarantee account for 55.5% of small business credit dollars. Results from the 2009 and 2010 Surveys of Consumer Finances indicate that personal savings or assets were used as collateral to initiate more than 70% of new businesses, making personal resources the most important funding source of entrepreneurs (Board (2011)). Also Mann (1998), Moon (2009), Dennis Jr. (2010), and Robb and Robinson (2012) highlight the importance of collateralized loans for small business finance. This collateral takes the form of personal assets - mostly real estate. The decrease in the value of real estate has made pledging personal commitments more difficult. Figure 1-5 shows that net mortgage equity extraction dropped from around 8% of disposable personal income in 2006 to around -6% at the end of 2010, the lowest value on record.[17] Although not all home equity is used for start-up financing, this 'deleveraging' by households which accompanied the dramatic decline in household net worth implies that the amount of equity available for start-up equity has been severely curtailed.[18]

State-level regressions: The impact of HPI

I test the hypothesis of a positive link between the value of real estate and the labor market by combining state-level data on house prices and establishment birth. Table 1 shows various state-level regressions. I use establishment births from the BLS Business Employment Dynamics (BDM) as the dependent variable (BIRTH). Al-

---

Reserve Banks of New York and Atlanta, as well as Dennis Jr. (2011), who report that young and small firms almost universally use one or more types of credit from a financial institution. Evidence in Siemer (2013) further supports the fact that young and small firms are more sensitive to changes in the credit environment.

[16]There was an increase in the costs of external finance during the last recession. Both the number and the dollar amount of approved C&I loans fell by around 20% over the course of the crisis. The drop in number and dollar amount of small loans (under $1 Million) was particularly severe. At the end of Q1 2013 the volume of small loans was only 84.84% of its pre-recession value, and their share of all C&I loans fell from a pre-crisis average of 32.33% to 22.39% in Q1 2013. This decrease in the number of loans was accompanied by an increase in the interest rate spread between smaller and riskier loans and the federal funds rate. Figures -11 and -12 in Appendix A.1 show the spread by loan size and risk.

[17]Thanks to Bill McBride for providing me with his estimates.

[18]Further evidence comes from FDIC data on used and unused home equity lines which I produce in Figure -14 in Appendix A.1. While unused commitments typically exceed outstanding home equity loans, the 2008/09 recession generated an earlier and steeper decline in unused equity lines. While part of this decline reflects drawdowns of existing lines a large portion represents a reduction of the credit supply by banks, as Bassett *et al.* (2011) argue in a similar context.

Mortgage Equity Withdrawal
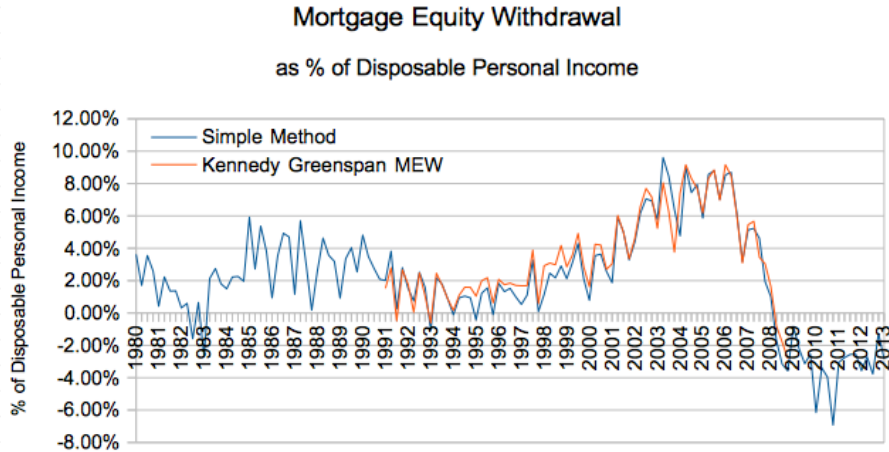
as % of Disposable Personal Income

Figure 1-5: Net equity extraction, or mortgage equity withdrawal (MEW) using the Flow of Funds and BEA data. For comparison the original Kennedy-Greenspan method is plotted as well. Not seasonally adjusted. Sources: Federal Reserve and BEA. This graph was created using the methodology proposed in Kennedy (2011) for updating the estimates presented in Greenspan and Kennedy (2005).

though this is establishment-level and not firm-level data I use this dataset because of its quarterly frequency. The data is available from Q2 1993-Q2 2013. The main explanatory variable is the state-level HPI, which comes from the Federal Housing Finance Agency (FHFA). As additional controls I use two alternative cyclical indicators: the state-level unemployment rate (UE) and state-level personal income (PI). I use personal income as a cyclical indicator because state-level GDP is only available on an annual basis. All variables have been HP-filtered. I am controlling for year- and state-effects and use cluster-robust standard errors in all regressions.[19] Summary statistics for all variables can be found in Appendix A.0. The first column in Table 1 shows a simple regression of BIRTH on HPI. The HPI is positively correlated with the number of new establishments at the state-level. This relationship is robust to controlling for cyclical indicators: Columns (2), (3), and (4) control for personal income and unemployment, which are both significant at the 5%-level and have the expected sign. Column (4) controls for both UE and PI jointly. Here, the state-level unemployment rate is no longer significant at the 10%-level. The last stylized fact summarizes the above results:

---

[19]I removed the states AK, DC, DE, HI, ND, SD, VT, WV, and WY from the analysis because of an FHFA warning. The HPI from those states have been derived from fewer than 15,000 transactions over the last ten years. Using a fixed-effect estimator leaves the results virtually unchanged. The same is true for using the variables in levels or logs instead of the cyclical component of the HP-filtered data. Results are available upon request.

	(1)	(2)	(3)	(4)
HPI	11.9366*	9.4346*	10.2039*	8.7394*
	(2.32)	(2.36)	(2.04)	(2.14)
PI		0.0153***		0.0149***
		(13.98)		(14.67)
UE			-87.2835*	-38.4972
			(-2.58)	(-1.13)
_cons	-50.4743	96.9491***	-48.6150	-50.1817
	(-1.87)	(5.27)	(-0.62)	(-0.69)
N	3276	3276	3276	3276
r2	0.0567	0.0775	0.0590	0.0779

Dependent variable: Establishment Birth. t statistics in parentheses.

All regression include year- and state dummies.

Source: BLS, FHFA, BEA.

All series are quarterly and have been HP-filtered with $\lambda = 1600$.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.1: Descriptive Regressions at the state level

Stylized Fact 5: *During the 2008/09 recession the financing environment for start-ups deteriorated. Credit supply by commerical banks decreased, and the value of real estate - widely used as collateral - fell.*

This section has produced five stylized facts about job creation and destruction. One, high unemployment is mainly driven by low job creation figures. Two, a large part of the decrease in job creation was due to the behavior of the youngest firms. Start-ups constitute the single largest contributor to net job creation. It is more volatile but less cyclical than job creationb by incumbents. Job creation by start-ups has taken a prolonged dive since the onset of the recent crisis. Five, there was a decrease in the availability of external finance for start-ups. Credit supply by commercial banks dropped during the 2008/09 recession, partly because declining property values diminished the value of collateral. I now present my model which figures a collateral channel and uses exogenous variation in the value of collateral to replicate the *Stylized Facts* presented above.

16

1.4 The Model

The economy consists of a fixed mass of workers and entrepreneurs. There is a competitive bank which provides start-up financing and is jointly owned by all agents. Each worker and each entrepreneur owns one unit of housing h, the price of which is q^h. Housing provides utility to all agents and can serve as collateral when entrepreneurs finance start-up loans. Workers can supply labor and consume their income, either from wages or home production. Entrepreneurs own the production process which utilizes labor to generate a single consumption good. Output is a function of labor and two types of profitability: idiosyncratic and aggregate. Shocks to profitability can be interpreted as changes in productivity or demand. Both types of profitability evolve persistently over time. Time is discrete and a period refers to one quarter.

The labor market is frictional. To hire unemployed workers firms must post vacancies v which are filled with endogenous probability. Labor is perfectly divisible. Following the standard search and matching literature a matching function captures those frictions. It is denoted as $m(U,V) = \mu U^\gamma V^{1-\gamma}$. Its inputs are the unemployment rate U and the vacancy rate V. Vacancies posted by firms are filled with probability $H(\theta) = m/V$. An unemployed worker finds a job with probability $\phi(U,V) = m/U$. The ratio $\theta \equiv V/U$ is a sufficient labor market statistic to compute the vacancy-filling and job-finding rates in this economy. Employed workers may lose their job if the entrepreneur they are matched with exits or decides to reduce employment in his production site. The worker takes both the job-finding rate and the job-destruction rate as exogenous. The workers' compensation for their labor input is specified through a bargaining process between the entrepreneur and the worker, where the entrepreneur has all the bargaining power.[20]

A fixed cost to production guarantees that firms exit when they receive a sufficiently low profitability draw. New firms that enter the economy must pay start-up cost c_e. To finance c_e, new firms obtain an intra-period loan from the bank. A fraction of the loan can be secured by collateral, for which agents use their real estate h. Changes in the value of collateral q^h lead to variations in the effective cost of entry $\tilde{c}_e$ and hence in the number of firms that enter the economy. Shocks to q^h are exogenous. I estimate a process of q^h and its cross-correlation coefficient with aggregate TFP from the US data.

The timing of events in my model is based on the setup in Hopenhayn and Rogerson (1993): Prior to the realizations of new aggregate and idiosyncratic shocks, firms decide whether to continue operating or exit. For new entrants exiting implies that

---

[20]This is following Cooper *et al.* (2007).

intra-period loans are defaulted on. Then the aggregate state realizes and new firms enter the economy without knowing their idiosyncratic productivity draw. The idiosyncratic shocks realize and all firms decide on their desired employment level. Bargaining takes place between workers and entrepreneurs, after which production occurs, and compensations are paid. The model is now explained in more detail.

1.4.1 Workers

Workers can either be employed or unemployed. They derive utility $\varphi(h)$ from housing independently of their employment status. When they are unemployed they receive an outside option $b(a)$, which can vary with the aggregate state a. This ouside option reflects the returns to home production. With probability $\phi(U, V)$ an unemployed worker is able to find a job, thus becoming employed next period. We can write the value of being unemployed as

$$W^u(a, h) = Z(b(a) + \pi^b) + \varphi(h) + \beta E_{a'|a}[\phi(U, V)W^e(a', h) + (1 - \phi(U, V))W^u(a', h)],$$
$$(1.1)$$

where $Z(\cdot)$ describes the worker's utility from consumption and redistributed profits made by the bank π^b. The term $\varphi(h)$ describes utility from housing. The discount factor is β, and $\phi(\cdot)$ is the job finding rate which depends on the current unemployment rate U as well as the number of vacancies V. The utility function $Z(\cdot)$ is assumed to be strictly increasing and concave. For simplicity I assume that there is no disutility from labor. The expectations operator in (1.1) is taken over the future values of unemployment and unemployment.

By constrast, when a worker is currently employed he receives a compensation ω as defined by the state-contingent contract. With (endogenous) probability δ the worker loses his job and receives the value of unemployment $W^u(a', h)$ next period. With the remaining probability he continues to be employed.

$$W^e(a, h) = Z(\omega(a) + \pi^b) + \varphi(h) + \beta E_{a'|a}[(1 - \delta)W^e(a', h) + \delta W^u(a', h)] \qquad (1.2)$$

1.4.2   Entrepreneurs

Entrepreneurs own the production process. Income from firms constitutes the entrepreneurs' only source of income and they consume all profits within the period.[21] Entrepreneurs are assumed to be risk-neutral. They produce using a production technology $F(e)$, where e represents the number of workers. The production function has the properites $F_e(e) > 0$ and $F_{ee}(e) < 0$, meaning it exhibits decreasing returns to labor, which might arise from fixed factors such as capital or materials, from imperfect substitutability for consumers of the goods produced by different firms or from managerial span-of-control as in Lucas (1978). At the end of a period entrepreneurs decide whether to continue operation or exit, based on their expectation of future shocks.[22] At the same time new entrepreneurs enter the economy. After the realization of uncertainty, entrepreneurs make hiring and firing decisions. A fraction χ of the workforce is separated exogenously (quits) each period. Given the state vector the entrepreneurs and the workers bargain over a compensation $\omega(a, \varepsilon, e)$. The firm's state vector at time t is $(a, \varepsilon, e, \theta)$, where θ reflects labor market tightness, as explained in more detail below. The profit function is given by

$$\pi(a, \varepsilon, e) = a\varepsilon F(e) - e\omega(a, \varepsilon, e) - \Gamma - \mathbb{C}. \tag{1.3}$$

Output is affected by two multiplicative profitability shocks a, and ε. While the former is an aggregate shock, the latter affects only idiosyncratic profitability. The term Γ is a fixed cost of operation to induce exit in low profitability states.[23] $\mathbb{C}$ defines a cost function given by

$$\mathbb{C} \equiv -F_v \mathbf{1}^+ - c_v v^2 \mathbf{1}^+ - F_f \mathbf{1}^- - c_f f^2 \mathbf{1}^-.$$

The indicator function $\mathbf{1}^+$ is equal to one if the firm is hiring and equal to $\mathbf{1}^-$ if the firm is firing. The number of vacancies posted is v and the amount of fired workers is f. There are two types of costs connected to hiring. One is a fixed cost F_v. The other is a quadratic cost c_v. The respective cost associated to firing are given by F_f

---

[21]See e.g. evidence in Moskowitz and Vissing-Jorgensen (2002) who show that entrepreneurial risk is not diversified and that dividends from the firm are the only source of income for owners.

[22]As in Hopenhayn and Rogerson (1993), since there is no additional information gained between periods, the exit decision is taken at the end of a period. This is mainly a computational convenience. Since I have one-period loans in my model the end-of-period exit decision is necessary to obtain default in the same period the loan was issued.

[23]The entrepreneur's problem is stated net of housing utility and net of redistributed bank profits because in the baseline model these values do not affect incumbent entrepreneurs's optimal decision. In Appendix A.4 I outline a model where the price of collateral also enters the incumbent entrepreneur's problem.

and c_f.

The Employment decision   A firm that is operation when idiosyncratic profitabilities are realized is called an incumbent, or 'continuing' firm. This firm employed e_{-1} workers last period and faces a shock x, where $x = (a, \varepsilon)$ consists of the aggregate and idiosyncratic productivity realization. Also part of the firm's state vector is the aggregate labor market tightness θ. This determines how effective the firm can hire new workers. In order to compute the expected value of θ' firms require knowledge about Γ, the joint distribution over (n, ε) and its law of motion. This is described in detail below. The state vector is summarized by $s = (x, e_{-1}; \theta)$. The value function for a continuing firm is denoted $Q^c(s)$. Because there are fixed costs to variations in employment, the entrepreneur faces a discrete choice problem within the period. He can decide between posting vacancies, remaining inactive, and firing workers. Vacancies must be reposted each period. The value $Q^c(s)$ is thus given by the maximum of the values of posting vacancies, firing, and inaction.

$$Q^c(s) = \max\{Q^v(s), Q^n(s), Q^f(s)\} \tag{1.4}$$

The three Bellman equations will now be described in turn. Because the entrepreneur can choose to exit at the end of the period, the continuation value in each case is given by the maximum of the expected value of continuing and exiting. The value of exit is $Q^x(e)$ and will be described below. The value of posting vacancies Q^v is given by

$$Q^v(s) = \max_v \pi(a, \varepsilon, e) + \beta E_x \max\{Q^c(x', e'; \theta'), Q^x(e)\}, \tag{1.5}$$

and the evolution of employment is given by the number of quits and the vacancy filling rate $H(\theta)$

$$e = e_{-1}(1 - \chi) + H(\theta)v,$$

The value of firing workers is

$$Q^f(s) = \max_f \pi(a, \varepsilon, e_{-1}(1 - \chi) - f) + \beta E_x \max\{Q^c(x', e'; \theta'), Q^x(e)\}. \tag{1.6}$$

Lastly, the value of inaction is given by

$$Q^n(s) = \pi(a, \varepsilon, e_{-1}(1 - \chi)) + \beta E_x \max\{Q^c(x', e'; \theta'), Q^x(e)\}. \tag{1.7}$$

Here E_x denotes the expectation conditional on the current value of x. The maximum operator nested on the right-hand side of all three Bellman equations reflects the fact that a firm can make a decision about exiting before the next period. Since this is decided before the realization of new information the choice can be made in the current period. Conditional on this period's employment choice the entrepreneur must evaluate the expected value of being active next period, given by $E_x[Q^c(x', e'; \theta')]$, and compare this to the present discounted value of exiting, given by $Q^x(e)$. This value is defined below. The policy function for employment will be denoted $\phi^e(s)$. The employment policy function will be characterized by different cutoff values in the (x, e_{-1}) space. For a given e_{-1} there exists a region of inaction over the values of the idiosyncratic shock due to the presence of fixed costs. An example is given in Figure -16 in Appendix A.2. For values higher than a cutoff profitability, the firm hires new workers, while for values below a lower cutoff profitability workers are shed. Note that changes in employment do not take 'time-to-build' because I want to rule this out as a driver of jobless recoveries.

The Wage Contract   We can now define the optimal wage contract between workers and entrepreneurs. The contract specifies $w(S)$, the compensation for a worker in a firm with state S, where $S = (a, \varepsilon, e, \theta)$ is the firm's state vector. A simplifying assumption is that entrepreneurs are able to make take-it-or-leave-it offers, i.e. the workers have zero bargaining power.[24] The firm thus chooses the wage subject to the worker's participation constraint. This is given by $W^e(a) \geq W^u(a)$. It says that the employed workers' outside option must be at least as large as the remuneration offered by the contract. In equilibrium the participation constraint will hold with equality, implying $Z(w(S)) + \varphi(h) = Z(b(a)) + \varphi(h)$, or $w(a) = b(a)$.[25] This is a simple way in which the model generates movements in the wage without the complexity of adding aggregate labor demand as an additional state variable. I assume the folowing functional form for the outside option $b(a) = b_0 a^{b_1}$. The parameter b_0 is part of the model calibration, while b_1 is estimated from the data. Importantly, $b_1 < 1$.

---

[24]As in Cooper *et al.* (2007) and many other papers this assumption is employed to facilitate the computation of the optimal contract. See Elsby and Michaels (2013) and Acemoglu and Hawkins (2013)for a different approach based on Stole and Zwiebel (1996). Kaas and Kircher (2011) introduce a competitive search procedure. This simplification does not change my results qualitatively as long as the elasticity of the bargained wage with respect to aggregate profitability is not larger than 1, for which to the best of my knowledge no evidence exists. In Appendix A.4 I show some intuition for a model with an alternative bargaining rule based on Stole and Zwiebel (1996).

[25]Formally, the profit maximizing contract results from the following optimization problem: $\hat{\pi}(a, \varepsilon, e) = \max_{w(S)} a\varepsilon F(e) - ew(S)$ subject to $W^e(a) \geq W^u(a)$.

The Exit Decision   At the end of a period, before any new information about the exogenous shocks arrives, an incumbent entrepreneur has to decide whether he wants to continue operating or exit next period.[26] The exit decision is thus based on the expected future value of the firm, which ensures that a firm will never post vacancies and exit in the same period. If the entrepreneur decides to exit, he will reduce the amount of workers to zero (paying the firing costs for the e remaining workers) and generate zero revenue. However, he avoids paying the fixed cost of operation. Any outstanding debt obligations are defaulted on. The value of exiting is given by

$$Q^x(e) = 0 - F_f - C_f e \leq 0. \tag{1.8}$$

This formulation implies that once a firm has decided to exit, it can not re-enter the market. All future profits are zero. The firm decides to exit whenever the expected value of continuing its operation is below the expected value of exiting with the current stock of employment carried over from the last period, e.

$$E_{a',\epsilon'|a,\epsilon}\left[Q^c(a',\varepsilon',e,\theta')\right] - Q^x(e) < 0. \tag{1.9}$$

Here F, the fixed cost of operation, induces exit for low realizations of ε since $Q^x(e)$ is always non-positive. The associated exit policy function will be denoted $\phi^x(s)$ and takes a value of one if the firm exits, and zero otherwise. Because $Q^c(x)$ is increasing in ε, for a given e, a, and θ the exit policy function is characterized by a threshold productivity level $\bar{\varepsilon}^x$ below which a firm exits. This threshold is defined as the lowest realization of ε such that the expected value of continuing exceed the value of exiting.

Definition. The threshold productivity level $\bar{\varepsilon}^x$ below which a firm exits is defined as

$$\begin{cases} \bar{\varepsilon}_t^x = \inf\left\{\varepsilon \in S : E_{a',\epsilon'|a,\epsilon}Q^c(a',\varepsilon',e_{-1},\theta') \geq Q^x(a',e_{-1})\right\} & \text{or} \\ \bar{\varepsilon}_t^x = 0 & \text{if this set is empty} \end{cases}. \tag{1.10}$$

Each period a fraction $F(\bar{\varepsilon}^x)$ of new entrants exits, while the remaining fraction continues operating. The cutoff $\bar{\varepsilon}^x$ is (weakly) decreasing in a, and (weakly) increasing in θ and e_{-1}. The intuition for this is straightforward: Because a is persistent, an

---

[26]Note that no additional information is revealed between the end of the current period and the time of the exit decision. Therefore the firm can determine in period t whether it will choose to exit in period $t + 1$. This insight makes the computation of the problem easier and brings the timing of the exit decision in line with the default decision by entrants.

increase in a raises the expected value of the continuing firm. At the same time the increase in a has no effect on the value of exit. Increases in θ decrease the firm's value and hence increase $\bar{\varepsilon}^x$. An increase in θ lowers the number of workers a firm that posts vacancies is able to hire, but has no effect on $Q^x(\cdot)$. The cutoff $\bar{\varepsilon}^x$ can never decrease in θ because the effect of θ on the firm's value function $Q^c(s)$ is less or equal to zero. Because the adjustment costs are increasing in e_{-1}, everything else equal a higher employment stock has a positive effect on $\bar{\varepsilon}^x$.

The Entry Process   At the beginning of each period there is a continuum of ex-ante identical potential entrants. The entry decision is made before the idiosyncratic profitability is known. Entrants do not pay a fixed cost of operation Γ in the first period. Instead, to enter, potential entrants must pay a start-up cost $\tilde{c}_e$, which they compare to the expected value of entry Q^e. The cost $\tilde{c}_e \equiv c_e \cdot \tilde{R}$ consists of a positive physical entry cost c_e times the interest rate charged by the bank, $\tilde{R}$ (defined below).[27] If the value function Q^c is known, the value of entry gross of entry costs is given by the value of an incumbent firm evaluated at zero employment and the expected initial productivity draw

$$Q^e(a, \theta) \equiv \int_\epsilon Q^c(a, \varepsilon_0, 0, \theta) d\nu.$$

Once an entrepreneur has decided to enter he receives an initial profitability draw ε_0 from a distribution ν, which may differ from the distribution of incumbents productivity draws. After the initial period, profitability evolves identically to that of all other incumbent firms. Employment in start-ups is given by the amount of succesful hires, $e = H(\theta)v$. The value of entry is increasing in a and decreasing in θ. Total start-up job creation is $\int_\epsilon \phi^e(a, \varepsilon_0, 0, \theta) d\nu$. Firms entering in period t have mass M_t, which is pinned down via a free-entry condition. Free entry requires that the cost of entry be equal to the value of entry.

$$\tilde{c}_e = Q^e(a, \theta). \tag{1.11}$$

Proposition 1. *There is a unique M_t which solves (1.11).*

The proof can be found in Appendix A.2. The logic is that as M_t increases, labor market tightness θ goes up since more firms are hiring. This negatively affects Q^e since a firm needs to post more (costly) vacancies to fill the same number of jobs. On

---

[27]I restrict attention to the case where $c_e < c^*$, where $c^* > 0$ is a number such that if $c_e \geq c^*$ no positive entry rates exist and the equilibrium is one of no firms. In the numerical solution of the model it will furthermore be the case that $c_e \geq \Gamma$, meaning that entrants have to pay a cost higher than the fixed cost of operation.

the other hand, $\tilde{c}_e$ is increasing in θ as the next section will show. This is a result of the exit threshold $\bar{\varepsilon}^x$ which is increasing in the labor market tightness. With Q^e monotonically decreasing and $\tilde{c}_e$ monotonically increasing in θ the intersection where (1.11) holds is unique.

1.4.3 The Bank

The bank is owned by all agents in the economy and behaves competitively, i.e. makes zero profits. To pay the entry cost c_e new firms must obtain a loan from the bank. Firms that are still in operation at the end of the period pay back the loan plus any interest payments that accrued. Entrepreneurs can use their real estate as collateral to secure part of the loan. This can be thought of as a shortcut for the idea that in reality some loans are completely secured by real estate while others are not. Putting down collateral for a loan is desirable because uncollateralized loans are risky for the bank, while collateralized loans are not. A start-up entrepreneur may strategically choose to exit and hence walk away from his obligations before the loan has to be repaid. Therefore, the bank efficiently prices interest rates by charging a default premium on the *uncollateralized* fraction of the loan in order to compensate itself for expected losses.[28] The collateralized fraction of the loan is riskless for the bank, hence the intra-period interest rate for it is 1. The fraction of the loan that can be collateralized depends on the value of real estate, q^h. The diagram in Figure 1-6 illustrates the structure of the loan.

Interest Rates and the Value of Collateral   Default occurs with positive probability, i.e. whenever a borrowing firm chooses to exit. In that case the bank claims the collateral which was used to secure the loan. No repayment is received for the uncollateralized fraction of the loan because - as can be seen from (1.8) - profits are non-positive if a firm exits. Payment of collateralized loans can always be enforced by the bank in case of default, hence the intra-period interest rate for this part of a loan is equal to the risk-free rate 1. This corresponds to the bottom area in Figure 1-6. The remaining fraction of the loan is *not* secured by collateral and the bank charges a loan rate $\hat{R} \geq 1$. Since the bank is perfectly competitive the loan rate is determined by a zero-profit condition $\hat{R}(1 - F(\bar{\varepsilon}^x)) = 1$. This implies that the risk-neutral bank receives the same expected return as the risk-free rate, which is 1. The total loan

---

[28]This is similar to the mechanism in Townsend (1979) and Bernanke *et al.* (1999) where the bank faces a costly state-verification problem. In my model state-verification is costless but in case of defalt the bank is unable to recuperate any fraction of the initial loan because wages are paid before the intra-period loan is reimbursed. I choose this timing of events in order to eliminate the default dimension from the worker-firm bargaining problem.
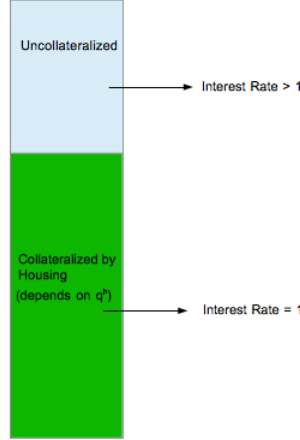
Figure 1-6: The Intra-period Loan. For the collateralized fraction of the loan the intra-period interest rate is 1. The uncollateralized part includes a positive default risk for which the bank charges a no-default interest rate larger than unity.

rate paid on c_e is denoted $\tilde{R}$ and is given by a combination of the risk-free rate and $\hat{R}$. Proposition 2 defines it.

Proposition 2. *The loan rate $\tilde{R}$ is given by*

$$\begin{cases} \tilde{R} = \frac{q^h}{c_e} + \left(\frac{c_e - q^h}{c_e}\right) \cdot \hat{R} & \text{if } q^h < c_e \\ \tilde{R} = 1 & \text{if } q^h \geq c_e \end{cases},$$

where

$$\hat{R} = \left(\int_{\bar{\varepsilon}^x}^{\infty} d\nu\right)^{-1}.$$

The proof can be found in the appendix. The intuition for the result is that if $q^h \geq c_e$ the new entrepreneur can fully collateralize his loan, which implies that he pays the risk-free rate on the intra-period loan. If $q^h < c_e$ he receives the risk-free rate only on a fraction $\frac{q^h}{c_e} < 1$ of the loan. The difference $c_e - q^h$ of the loan c_e has to be borrowed at the risky interest rate $\hat{R}$. This rate is increasing in the probability of receiving an initial profitability draw below $\bar{\varepsilon}^x$. If an entrant would *never* choose to exit, then the integral $\int_{\bar{\varepsilon}^x}^{\infty} d\nu = 1$ and $\hat{R} = \tilde{R} = 1$. Changes in $\tilde{R}$ are a key driver for the dynamics of the model because changes in the cost of entry have important effects on the number of entrants and hence on job creation and unemployment. Since $\bar{\varepsilon}^x$ is (weakly) decreasing in a and (weakly) increasing in θ it follows that both $\hat{R}$ and $\tilde{R}$ are (weakly) decreasing in a and (weakly) increasing in θ. Furthermore, the effective loan rate $\tilde{R}$ is (weakly) decreasing in q^h.

1.4.4 Equilibrium

The distribution over incumbent firms In the absence of aggregate shocks (as in Hopenhayn and Rogerson (1993)) it is possible to solve for a stationary distribution of incumbent firms λ^*. Although my model incorporates aggregate shocks it is useful to spell out the transition of the firm distribution here, since the non-stochastic simulation method is based on it. The distribution over incumbent firms in period t is given by λ_t. The mass of entering firms shall be denoted M_t. I will drop the time subscripts for notational convenience. The transition from any λ to λ' will be written as $\lambda' = T(\lambda, M)$. The operator T is linearly homogeneous in λ and M jointly. This implies that if we doubled the amount of firms in this economy and doubled the amount of entrants the resulting distribution would be unchanged.

Assuming that some initial distribution λ_0 exists and given the policy functions for employment and exit the law of motion of the distribution over incumbent firms is given as follows. For any set $(e\ x)' \in E \times X$, where E and X respectively denote the employment and exogenous shock space the law of motion for λ is

$$\lambda'((e\ x)' \in E \times X) = \int_{x \in x'} \int_{E \times X} (1 - \phi_x(x, e; \theta)) \times 1_{\{\phi_e(x,e;\theta) \in e'\}} \times F(dx'|x)\lambda(dex)$$

$$+ M \times \int_{x \in x'} \int_{0 \times X} \times 1_{\{\phi_e(x,0;\theta) \in e'\}} \times F(dx'|x)\nu(dx). \quad (1.12)$$

This defines the operator T. For the case without aggregate shocks $x = \varepsilon$ and a stationary distribution λ^* exists.[29]

Endogenous and Exogenous processes The law of motion for the labor market tightness θ follows the law of motion

$$\theta' = \mathbb{H}(a, a', \lambda).$$

---

[29]Equation (1.12) can be most easily read by fixing an exogenous state x', then integrating over the space of incumbents $(E \times X)$ and selecting those for whom the policy function $\phi_e(\cdot)$ prescribes e'. The term $F(dx'|x)$ defines the probability that a firm with current productivity x has productivity x' next period. This is multiplied with λ to obtain the mass of these firms. The second term refers to entrants, who have mass M. Their initial employment is equal to zero and they cannot exit in the same period as they enter, otherwise the structure is identical. The stationary equilibrium with entry and exit is given by $\lambda^* = (I - \pi')^{-1}(\pi' * E)$, where λ is the distribution over incumbents, π is the transition matrix and E is the distribution over entrants.

The knowledge of $\mathbb{H}$ requires the joint distribution over employment and idiosyncratic profitability, which is (theoretically) infinitely-dimensional. I follow the approach developed by Krusell and Smith (1998) described in the following section. Briefly, the approach consists of postulating a functional form for $\mathbb{H}$ which entrepreneurs use to make their optimal decisions. From a subsequent simulation of the model one can check the consistency between the actual law of motion of θ and the one predicted by the guess of $\mathbb{H}$. The resulting equilibrium must be such that $\mathbb{H}$ must track the evolution of θ very accurately. This is explained in more detail below.

Unemployment in the model follows $U' = (1-U)\delta(U,V) + (1-\phi(U,V))U$, where $\delta(U,V)$ is the separation rate and $\phi(U,V)$ describes the job-finding rate. I assume that the logarithms of both a, ε, and q^h follow autoregressive processes.

$$\ln a_t = \rho_a \ln a_{t-1} + v_{a,t} \ , \ v_a \sim N(0, \sigma_a) \tag{1.13}$$

$$\ln \varepsilon_t = \rho_\varepsilon \ln \varepsilon_{t-1} + v_{\varepsilon,t} \ , \ v_\varepsilon \sim N(0, \sigma_\varepsilon) \tag{1.14}$$

$$q_t^h = \rho_q q_{t-1}^h + v_{q,t} \ , \ v_q \sim N(0, \sigma_q) \tag{1.15}$$

The initial productivity of entrants is determined by a drawn from $v_\nu \sim N(0, \sigma_\nu)$ and then evolves according to (1.14). In the simulation I enforce a correlation coefficient between q^h and a obtained from the data.

Equilibrium For a given λ_0 a recursive competitive equilibrium consists of (i) value functions $Q^c(a, \epsilon, e_{-1}; \theta)$ and $Q^e(a, \epsilon_{i,0}, \theta)$, (ii) policy functions $\phi^e(a, \epsilon, e_{-1}; \theta)$ and $\phi^x(a, \epsilon, e_{-1}; \theta)$, (iii) bounded sequences of non-negative negotiated wages $\{w_t\}_{t=0}^\infty$ and interest rates $\{\hat{R}_t\}_{t=0}^\infty$, unemployment $\{U_t\}_{t=0}^\infty$, vacancies $\{V_t\}_{t=0}^\infty$, incumbent measures $\{\lambda_t\}_{t=0}^\infty$ and entrant measures $\{M_t\}_{t=0}^\infty$ such that (1) $Q^c(a, \epsilon, e_{-1}; \theta)$, $\phi^e(a, \epsilon, e_{-1}; \theta)$, and $\phi^x(a, \epsilon, e_{-1}; \theta)$ solve the incumbent's problem, (2) $\{w_t\}_{t=0}^\infty$ satisfies the worker's participation constraint, and $\{\hat{R}_t\}_{t=0}^\infty$ is given by the bank's zero-profit condition, (3) labor market tightness $\{\theta_t\}_{t=0}^\infty$ is determined by the ratio of vacancies $\{V_t\}_{t=0}^\infty$ over unemployment $\{U_t\}_{t=0}^\infty$, (4) the measure of entrants is given by the free-entry condition (1.11), (5) λ_t evolves according to (1.12).[30]

---

[30] In the economy with aggregate shocks this equilibrium is boundedly rational because the law of motion for θ is approximated.

27

1.4.5 Calibration

I calibrate the model at quarterly frequency. The steady state equilibrium without aggregate shocks matches US non-farm establishment level data. All paramter values together with their calibration targets are listed in Table 1.2. The parameters can be divided into two groups. The first group consists of parameters that are either taken from the existing literature or backed out given static calibration targets. The second group of parameters is estimated with a simulated method of moments (SMM) procedure. The first group includes the discount factor β, the curvature of the profit function α, the value of leisure parameters, the parameters governing the evolution of the aggregate states, as well as the parameters of the matching function. β and α are taken from the literature. I fit AR(1)-processes to the data to back out the persistence and innovation parameters for a and q^h. For a I use US output from 1977-2011, while for q^h I use the purchase-only HPI from 1977-2011. Both series are HP-filtered. The correlation between output and HPI is 0.628. I enforce this correlation coefficient onto the simulated processes. Recall that workers' value of leisure is $b(a) = b_0 a^{b_1}$. To estimate b_1 I use (HP-filtered, seasonally adjusted) average weekly wages from the Quarterly Census of Employment and Wages (QCEW) between 2001 and 2011. The correlation between the cyclical component of this series and GDP is 0.49, which is almost identical to the value used in Cooper *et al.* (2007). I calibrate b_0 to match an average firm size of 21.43 from the BDS data.Regarding the parameters of the matching function, I assume a constant returns to scale function which takes the form

$$m = \mu U^\gamma V^{1-\gamma} = \mu V \theta^{-\gamma},$$

where $\theta \equiv \frac{V}{U}$ measures labor market tightness. The job-finding rate of a worker is defined as $\phi = m/U$, which given the functional form for the matching function takes the form $\phi = \mu \theta^{1-\gamma}$. Similarly the vacancy-filling rate for firms, $H = m/V$ takes the form $H = \mu \theta^{-\gamma}$. Based on BLS data the average unemployment rate over the time of my sample (1977-2010) was 6.3%, which serves as my target for the steady state. I target a vacancy-filling probability of 0.71, in line with empirical evidence in Den Haan *et al.* (2000), Pissarides (2009), Shimer (2012), and Elsby and Michaels (2013). The same studies suggest a steady-state value of $\theta = 0.70$. The matching elasticity γ is set to 0.60.[31] My target for the vacancy-filling rate together with a choice of γ implies a matching efficiency parameter of $\mu = 0.5732$.

The cost parameters in $\mathbb{C}$ and the parameters governing the idiosyncratic prof-

---

[31]This is based on a survey by Pissarides and Petrongolo (2001). Cooper *et al.* (2007) estimate this parameter to be .36, Hall (2005*b*) finds 0.72.

Calibrated Parameters	Symbol	Value	Source
Discount Factor	β	0.99	implies $r^{ann} = 4\%$
Curvature of profit function	α	0.65	Cooper *et al.* (2007)
Autocorrelation of a	ρ_a	0.95	Output 1977-2011
Standard deviation of ν_a	σ_a	0.05	Output 1977-2011
Autocorrelation of q^h	ρ_q	0.9565	HPI 1977-2011
Standard deviation of ν_q	σ_q	0.08	HPI 1977-2011
Correlation q^h and a	$\rho_{a,q}$	0.628	same as above
Base wage	b_0	0.9	Average firms size 21.43
Sensitivity of outside option to a	b_1	0.49	BLS QCEW
Matching elasticity	γ	0.6	Pissarides and Petrongolo (2001)
Match efficiency	μ	0.5732	$H = 0.71, \theta = 0.7$
Estimated Parameters	**Symbol**	**Value**	**Calibration Target / Source**
Fixed costs vacancies	F_v	0.01	Inaction in Δe
Quadratic costs vacancies	c_v	0.005	Small changes in Δe
Fixed costs firing	F_f	0.01	Inaction in Δe
Quadratic costs firing	c_f	0.005	Small changes in Δe
Fixed costs of operation	Γ	3.3	Firm Exit Rate 8.8% (BDS)
Autocorrelation of ε	ρ_ε	0.97	Firm size distribution
Standard deviation of ε	σ_ε	0.02	Distribution of Δe, JC
Std dev of initial productivity	σ_ν	0.02	Start-up Fraction of JC = 18.7%

Table 1.2: Parameter Values. The first block consists of calibrated parameters, the parameters in the second block consists were estimated via SMM.

itability process are consistently estimated via SMM. This entails finding the vector of structural parameters Θ which minimizes the (weighted) distance $L(\Theta)$ between data moments and moments of the model. The distance is defined as

$$L(\Theta) = \left(\Gamma^D - \Gamma^M(\Theta) \right) \Xi \left(\Gamma^D - \Gamma^M(\Theta) \right)',$$

where Γ^D are data moments and $\Gamma^M(\Theta)$ are moments from a simulation of the model, given parameters Θ. The weighting matrix is Ξ. I solve the dynamic programming problem and generate policy functions given a parameter vector Θ. From the simulation of the model I then obtain $\Gamma^M(\Theta)$.[32] The algorithm finds the parameter vector Θ which minimizes $L(\Theta)$. The parameter vector is $\Theta = (F_f, c_f, F_v, c_v, \Gamma, \rho_\varepsilon, \sigma_\varepsilon, \sigma_\nu)$. I restrict the model such that $F_f = F_v$ and $c_f = c_v$, i.e. hiring and firing costs are symmetric. The moments Γ^D chosen to estimate Θ are motivated by Cooper *et al.* (2012) and Berger (2012) and are reported in the column 'Data' in Table 1.3. The first four moments are derived from the distribution of employment changes for continuing establishments using Census BDS data between 1985-1999. The first row reports the inaction rate, i.e. the fraction of establishments that did not undergo any employment change over the course of one year. The high value suggests that fixed costs of labor adjustment are important. The second column $|\Delta e| \leq .1$ reports the fraction of 'small' employment changes of under 10% in absolute value. Rows 3 and 4 report large positive and negative employment changes of over 30%. These large changes are very prevalent in the data, indicating large changes in firm-level productivity over time. Row 5 is the firm exit rate from the BDS data between 1977 and 2011. From the same data comes the fraction of gross job creation through firm birth, which is around 19%. Both Θ and Γ^D consist of six (unique) elemts, but there exists no direct mapping between them. The following can be said about the identification, however: The fixed costs F_f and F_v play a crucial role for generating inaction, while the quadratic costs are identified through small employment changes, $|\Delta e| \leq .1$. The quadratic costs also play an important role for generating exit among large plants. The operational overhead cost Γ is used to pin down the exit rate. Start-up job creation largely depends on the initial productivity draw, whose variance is governed by σ_ν. The persistence of the idiosyncratic shock ρ_ε is crucial for determining the shape of the size and age distributions and affects the frequency of employment adjustments. The variance of ε is important for large adjustments and the size distribution of firms. It indirectly affects all moments in $\Gamma^M(\Theta)$. The stationary model is further discussed below.

---

[32]Equilibrium is enforced during all of these estimations, meaning that the entrepreneur's beliefs about θ are consistent.

	Data Moments	Model Moments
$\Delta e = 0$	0.38	0.36
$|\Delta e| \leq .1$	0.09	0.07
$\Delta e > .3$	0.15	0.15
$\Delta e < -.3$	0.14	0.16
Exit Rate	8.8%	6.3%
Start-up JC	18.7%	17.0%

Table 1.3: Data Moments and SMM estimates. Column 3 estimates the benchmark model using symmetric adjustment costs (AC) for hiring and firing. The employment change numbers are taken from Berger (2012) who uses LBD averages between 1985-1999 . The exit rate and start-up JC rate are computed using BDS data.

1.5   Computational Strategy

Firms need to forecast θ' in order to compute the expected vacancy-filling rate $H(U, V)$. The variable θ is determined in equilibrium. While firms take this function as given, it must be consistent with the relationship generated by the model. In the stationary model without aggregate shocks there is a steady state value θ^* which can easily be determined. Including aggregate shocks creates a non-trivial computational problem, which I solve similarly to Krusell and Smith (1998). The free-entry condition is given by (1.11). The labor-market tightness θ is now a slow-moving state variable about which firms must generate consistent forecasts. The solution of this model is non-trivial since firms need to forecast the entire cross-sectional joint distribution of employment and productivity in order to forecast labor market tightness in the following period. In the presence of aggregate shocks, this distribution moves over time and the state-space becomes (theoretically) infinite-dimensional. Following the seminal work of Krusell and Smith (1998) an approximate solution can be found by postulating that firms track only several moments of this joint distribution. The first moments usually turns out to be a sufficient statistic. The word *sufficient* typically means that the forecast generates a high R^2. However, as Den Haan (2010) has shown, it should also be verified that the maximum forecast errors that result from the approximated law of motion are small. In the present framework firms are ultimately interested in forecasting θ', the labor market tightness next period. The perceived law of motion of θ is denoted $\theta' = \mathbb{H}(\theta, A', A)$, where $\mathbb{H}(\cdot)$ is to be determined as part of the solution of the model. Firms make their forecasts of θ' conditional on the current realizations of θ and A, as well as on possible future realizations A'. The solution algorithm first postulates an initial guess for for $\mathbb{H}(\cdot)$. Next, policy functions are computed given the guess. Following a simulation, the parameters of $\mathbb{H}(\cdot)$ are

updated. This procedure is repeated until the current guess and the updated version of $\mathbb{H}(\cdot)$ are sufficiently close (consistency) and until $\mathbb{H}$ tracks the evolution of θ with high accuracy. I guess a log-linear prediction rule for θ'.

$$\log \theta_t = b_0 + b_1 \log \theta_{t-1} + b_2 \log A_t + b_3 \log A_{t-1} + b_4 \cdot I_{A_t \neq A_{t-1}}$$

The last term, $I_{A_t \neq A_{t-1}}$, is an indicator function which takes the value of one if $A_t \neq A_{t-1}$. The coefficients that minimize the stopping criterion are given by

$$\log \theta_t = -0.0087 + 0.9939 \cdot \log \theta_{t-1} + 20.996 \cdot \log A_t - 21.095 \cdot b_3 \log A_{t-1} + 0.2327 \cdot I_{A_t \neq A_{t-1}}.$$

This functional form for $\mathbb{H}(\cdot)$ generates an $R^2 = 0.9994$ and a maximum forecast error of 0.005%. Accuracy plots can be found in the Appendix A.2. Note that without financing friction (i.e. no variation in q^h) the computational problem is much easier to solve. When the only shocks are to a the model behaves very similarly to the standard HR model. In particular, the free entry condition reduces the computational burden because the future value of θ can be computed without a Krusell-Smith type algorithm for the cross-sectional distribution. The reason is that with free entry aggregate labor demand becomes perfectly elastic and for each a there exists one value of θ which is consistent with equilibrium. Free-entry of new firms makes the tightness parameter θ respond 1:1 to changes in the aggregate state a.[33] However, such a model generates unrealistically volatile entry rates and basically reduces the model to a function of the aggregate state a, with some propagation through the adjustment costs.

The simulation of the model is carried out using a non-stochastic simulation technique. The algorithm does *not* draw a random sequence of idiosyncratic shocks for each firm and play out the policy function for a large number of periods. Instead, my algorithm computes the exact mass of firms at each grid point jointly representing idiosyncratic productivity and employment. This solution method is applicable for both the stationary and non-stationary version of the economy. The main advantages of this approach are its speed and the fact that it eliminates sampling error. Den Haan (2010) showed that this latter source of error can become important in Krusell-Smith type solution algorithms. The details of this algorithm are laid out in

---

[33]The labor market tightness 'jumps' with the aggregate state when the only shocks are to a. The true and the approximated law of motion are almost indistinguishable. A regression which ignores past realizations of θ produces an $R^2 > 0.99$ and a maximum forecast error of 0.0052%. The R^2 is not equal to 1 because θ influences the interest rate $\hat{R}$ which effects the number of entrants and hence the labor market tightness. Including past realizations of θ into the regression increases the R^2 to over 0.99999999.

Firm Age	Firms		Employment	
	Data	Model	Data	Model
Start-Ups	11%	12%	3%	5%
Age 1-2	16%	19%	6%	8%
Age 3-5	17%	22%	8%	13%
Age 6-20	41%	40%	27%	23%
Age 21+	15%	8%	56%	51%

Table 1.4: Age-distribution of firms. Census BDS data and results from the stationary model.

Appendix A.3.

1.6 Quantitative Results

This section describes the numerical results. I evaluate the performance of the stationary model with respect to non-targeted moments and then discuss the results of the model with aggregate shocks.

1.6.1 Results of the stationary model

Table 1.3 showed the match of targeted moments. The employment change distribution as well as the exit rate and job creation by start-ups generated by the model are very close to their counterparts in the US data. The fit of the firm-age distributions of the calibrated model is shown in Table (1.4). The model matches the age distribution of firms well but slightly underpredicts the amount of old firms.

1.6.2 Results with Aggregate Shocks

I now add aggregate shocks to the model in order to assess the business cycle properties of the model and evaluate its quantitative performance. To demonstrate the effect of shocks to aggregate productivity and the HPI, impulse repsonse functions are generated. I also test alternative model specifications without financial frictions and without adjustment costs in order to build some intuition about the respective effects those features on the results. Finally, I show a policy experiment which allows me to back out the effects of the decrease in the HPI on the increase and persistence of unemployment during and after the Great Recession. The main results are summarized in Tables 1.5 and 1.6.

	σ_U	ρ_U	σ_V	ρ_V	$\rho_{U,V}$	σ_θ	ρ_θ
US Data	0.13	0.948	0.16	0.93	-0.896	0.316	0.94
Benchmark Model	0.13	0.996	0.17	0.91	-0.86	0.303	0.943
constant q^h	0.17	0.995	0.198	0.95	-0.94	0.359	0.984
constant a	0.02	0.99	0.02	0.90	-0.89	0.03	0.97

Table 1.5: Business Cycle Statistics of the Model. Source: FRED, FHFA, and BLS. Data (1995Q1-2010Q4) and model moments have been computed as log deviations from mean/trend. Vacancy data starts in 2001Q1. σ denotes the standard deviation and ρ the autocorrelation of unemployment (U), vacancies (V), and labor market tightness (θ). The term $\rho_{U,V}$ is the correlation between unemployment and vacancies.

Results of the Benchmark Model

This section describes the results of the benchmark model which includes shocks to a and q^h. The model is able to match the key statistics of the US labor market regarding unemployment, vacancies, and their joint movement. Those statistics are reported in the first row of Table 1.5. The resulst of the benchmark model are in the second row. Both the volatility and the autocorrelation of unemployment, vancancies, and labor market tightness are close to their counterparts in the data. However, the persistence of unemployment is overstated, while the persistence of vacancies is understated with respect to the data. The correlation between unemployment and vacancies is strongly negative, as in the data. Given that the model was not calibrated to generate these moments the close fit can be considered a success of the calibration strategy.

The second set of results focuses on the cyclicality and volatility of employment in start-ups vis-a-vis incumbent firms. Two of the stylized facts presented in Section 2 were that job creation by incumbents is more strongly pro-cyclical, while job creation by start-ups is more volatile around its trend. Those facts are summarized in the first row of Table 1.6. The first two columns show the correlation between GDP and job creation by entrants (E) and incumbents (I). The last two rows report the standard deviation of the cyclical component of job creation over the trend component for the two groups.[34] The model generates the lower pro-cyclicality of job creation by entrants with respect to incumbents. The good fit in the correlation between GDP and job creation by new firms is achieved through the effect of house prices q^h on the entry process as will be explained below. Furthermore, the model replicates the higher correlation between GDP and job creation by incumbent

---

[34]I divide the series by their respective trend in order to control for the fact that otherwise the large number of jobs created by incumbents blows up the standard deviation of the series. An alternative measure that delivers similar results is the coefficient of variaion.

firms. This has been an important feature of the recovery after the Great Recession.[35]

The benchmark model can generate 'jobless recoveries' through the effect of house prices q^h on the start-up process. Imagine a situation where both aggregate profitability and the HPI are below their unconditional means. Now both shocks start reverting back but - as we will see below - the effects on unemployment and total output of the two shocks differ significantly. Other than the shock to aggregate profitability the shock to q^h exerts only very mild influence on total output. By directly impacting entry, the decrease in q^h has a large effect on hiring by start-ups, and thus on unemployment. The fraction of total hiring by start-ups is overproportional to their share of total output. Therefore, if the number of entrants decreases, the effect on unemployment is larger than the effect on GDP. Incumbent firms are only indirectly affected by the HPI through an effect on θ. On the other hand, shocks to a have the effect that hiring - and most importantly - output by incumbent firms changes. Since the lion's share of total output is produced by incumbent firms, an increase in a after an initial negative shock has an immediate effect on output *and* employment. This is why a shock to a alone cannot generate a jobless recovery. It requires the effect on entry - exerted by shocks to q^h - to make the unemployment rate react sluggishly and uncouple it from the strong co-movement with GDP. The impulse response functions will show this in more detail.

Results of the Alternative Model Specifications

We can now compare the benchmark results to those of the model without financial frictions or without shocks to aggregate productivity. The results are summarized in the last two rows of Tables 1.5 and 1.6. Table 1.5 shows that the business cycle statistics of the model without the financial friction are similar to the benchmark model. The volatility of unemployment and vacancies, as well as the correlation between the two is slightly overstated. Furthermore, θ is more volatile than in the data. The fact that the model produces similar moments as the benchmark model is not very surprising given the similarity of the model without the financial friction to Cooper *et al.* (2007), who find similar results. The model without shocks to a, on the other hand, is unable to capture some of the key US business cycle statistics. In particular, the model does not generate enough variation in unemployment and vacancies. The reason is that variations in q^h have a strong effect on start-ups but only an indirect effect (through labor market tightness) on incumbent firms. The movements in θ generated by changes in q^h are by themselves not sufficient to generate

---

[35]See the additional material, e.g. Figure -6 in Appendix A.1.

	$\rho(Y, N^E)$	$\rho(Y, N^I)$	$\sigma(c/t)^E$	$\sigma(c/t)^I$
US Data	0.35	0.76	0.10	0.07
Benchmark Model	0.34	0.65	0.20	0.07
constant q^h	0.60	0.79	0.30	0.07
constant a	0.06	0.13	0.08	0.06

Table 1.6: Data and Model Moments. Source: BDS 1977-2011. The resulting model moments have been computed using time aggregation. Data and model moments have been computed as log deviations from mean/trend. $\rho(Y, N^E)$ and $\rho(Y, N^I)$ show the correlation between GDP and gross job creation by entrants and incumbents. The standard deviation of the cyclical over the trend component of job creation by start-ups are $(\sigma(c/t)^E)$ and $\sigma(c/t)$ for incumbent firms.

the observed time-series volatility. Table 1.6 shows the model performance regarding job creation by entrants and incumbents.

Impulse Response Functions

In order to disentangle the respective effects of θ and a I show several impulse response functions in Figures 1-7-1-9. Figure 1-7 studies a negative shock to aggregate profitability, Figure 1-8 shows results for a negative shock to q^h, and in Figure 1-9 both shocks occur simultaneously. For comparability between the IRFs the size of the (negative) shocks to a and q^h were chosen to generate the same contemporaneous increase in unemployment. The figures are all constructed in the same way: The first panel shows the effect of the shock to the exogenous state. The second panel (clockwise) shows the effects on unemployment and GDP. The third panel plots the labor market tightness θ, while the last panel shows the effect on start-up activity.

I start with Figure 1-7 where the effects of a drop in a are analyzed. The first panel shows that in period $t = 10$ aggregate profitability falls by 1.22%. This results in a contemporaneous increase of the unemployment rate by 5.8%, and a fall in GDP by 1.35%. Labor market tightness falls, both because incumbent firms post fewer vacancies and because there are fewer entrants. The last panel also shows that the mass of entrants quickly rebounds after the initial shock. The reason is that the entrants are facing a trade-off between the lower aggregate profitability and the decreased labor market tightness. The latter has the effect of making it more profitable for potential entrants to start operating. Starting in period $t = 14$ the mass of entrants is above its unconditional mean, beginning to restore the total mass of firms to its pre-recession value.

Now I turn to analyzing the implications of a negative shock to q^h. The first panel of Figure 1-8 shows that in period $t = 10$ q^h decreases by 4.12%.[36] The shock
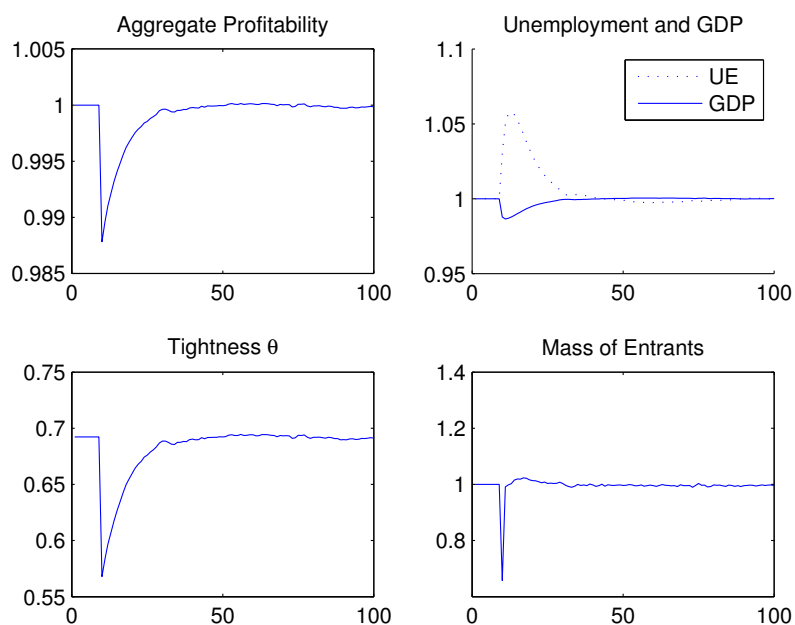
Figure 1-7: Impulse Response Functions for a shock to a. Simulation results from 10,000 repititions of 200 periods.
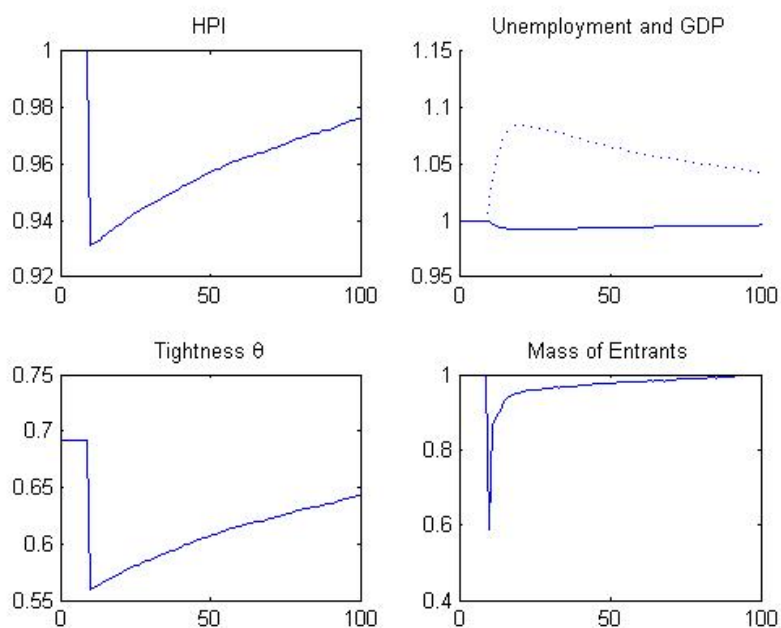

Figure 1-8: Impulse Response Functions for a shock to q^h. Simulation results from 10,000 repititions of 200 periods.

generates an increase in unemployment of 5.8%. This can be seen in the second panel. The shock to q^h produces a smaller decrease in GDP (0.48%) than the shock to a. This is because incumbent firms are only indirectly affected by the HPI shock, namely through the effect on θ which is displayed in the third panel. Labor market tightness decreases when the shock occurs and then slowly recovers. For incumbents firms and hiring entrants this implies that following the shock to q^h the vacancy-filling probability $H(\theta)$ increases. This has the effect that job creation by incumbent firms increases. The last panel shows the effect on the number of start-ups. The most important difference with respect to the effects of a shock to a is that the mass of entrants is affected both more severely and for a longer period of time. After a rebound to around 92% of its steady-state value in $t = 11$ the entry rate is only gradually moving back towards its unconitional mean.The shock to a generated a tradeoff between lower profitability and lower θ, which induced high entry rates after aggregate productivity had been beginning to recover. The outcome generated by the drop in q^h is different in the sense that the higher entry costs outweigh the effects of the drop in θ for new entrants. This is the main takeaway from Figures 1-7 and 1-8: In the context of the model, a jobless recovery must be the result of a simultaneous shock to both a and q^h. While the mean reversion of aggregate profitability brings GDP back to its pre-recession value, the slow recovery of the HPI has almost no output effect, but a large positive effect on the unemployment rate. Therefore, although GDP is above its recession trough, the decline in the unemployment rate is strongly underproportional to this decrease.

Figure 1-9 shows results for a simultaneous shock to a and q^h. The first panel plots the two shock processes. The second panel shows that the average increase in unemployment is 10.2%, while GDP drops by 1.59%, both of which is lower than the sum of the effects of the individual shocks. Both shocks are mean reverting but the persistent q^h shock keeps the unemployment rate high although GDP has practically recovered its pre-shock value (after $t = 20$ average GDP stands at 0.9978 of the pre-shock value). The effect on the number of entrants is strong. There is a sharp rebounce in the periods after the initial shock but no overshooting, as the dampening effect of the low q^h prevails over the mean reversion in the shock to a.

---

[36]This is a failry large shock compared to the decrease in the HPI during the Great Recession. The average HPI growth between 2007Q1 and 2011Q1 was -1.46% per quarter, the minimum was -2.88%.
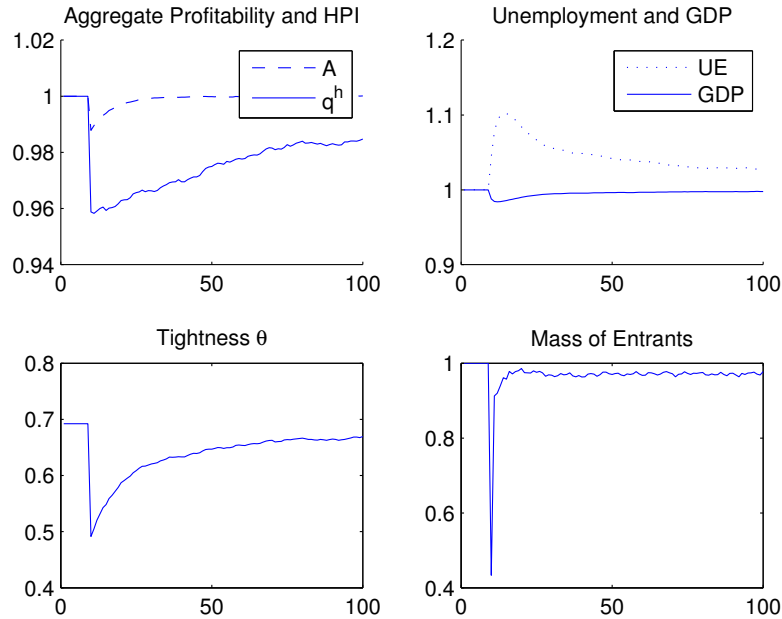
Figure 1-9: Impulse Response Functions for a shock to a and q^h. Simulation results from 1,000 repititions of 200 periods.

Policy Experiment

Tables 1.5 and 1.6 showed that the model is able to match they key properties of the US labor market as well the cyclicality and volatility of job creation by entrants and incumbent firms. The impulse response functions were meant to create some intuition about the effect of the two shocks. I now test in how far the model can replicate the relationship between the cyclical components of GDP growth and unemployment during the 'Great Recession'. To evaluate the model's performance in this respect I feed in the observed house price index between 1990Q1 and 2013Q1 (see Figure -10). Furthermore, I pick the sequence of aggregate productivity shocks to match the cyclical component of GDP over the same period. I simulate the model for 93 periods after some initial periods for the model to reach the stationary distribution. I choose 93 periods because this corresponds to the number of quarterly observations. The results are presented in Figure 1-10. The co-movement of the two time series is extremely strong, particularly during the 'Great Recession', indicated by the third shaded area. The simulated data is able to explain 72.23% of the variation of the unemployment rate observed in the data. For the period starting in 2006 the simulated data can even explain 84.66% of the movement in the unemployment rate. The recovery is 'jobless' because of the ongoing negative influence of the low HPI on start-
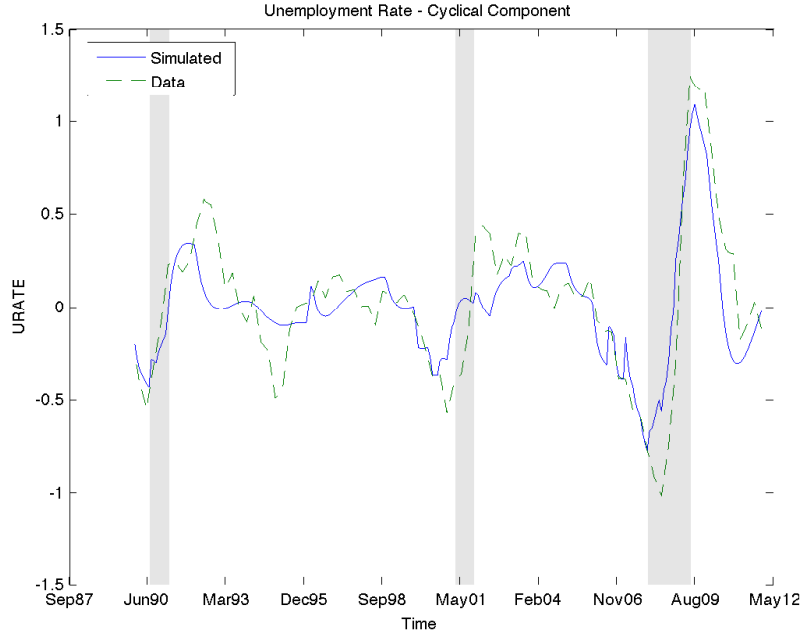
Figure 1-10: Cyclical component of the unemployment rate. Data vs. simulation using estimated processes for a and q^h between 1990 and 2011. Shaded areas correspond to NBER recession dates.

up job creation. Like in the data this leads to high levels of unemployment even after the official recession end. We see that job creation by start-ups decreased prior to the beginning of the recession. The model has this feature simply because the drop in the HPI precedes the decline in aggregate productivity.[37] Net job creation by incumbents begins to recover before job creation by start-ups. This is the case because at the end of the recession incumbent firms take advantage of the high vacancy filling probability due to the low θ, while hiring for start-ups remains costly because of the ongoing low q^h which increases the cost for setting up shop. What the model is unable to match is the time lag in the respective troughs of the HPI and job creation by start-ups. In the simulation job-creation by start-ups coincides with the trough in the HPI series, while in the data job creation by start-ups was lower in 2011 than in 2009.

In Appendix A.2 I repeat this experiment when there are only shocks to a or q^h. Figures -17 and -18 show that although the variation in q^h generates a lot of movement in the unemployment rate it is not enough to reproduce the large increase in unemployment which accompanied the recent recession.[38]

---

[37]The HPI showed negative growth rates as early as Q12006, while the NBER dates the beginning of the recession in Q42007.

[38]The q^h shock alone explains about 59.25% and the a shock alone about 56.93% of the variation in unemployment.
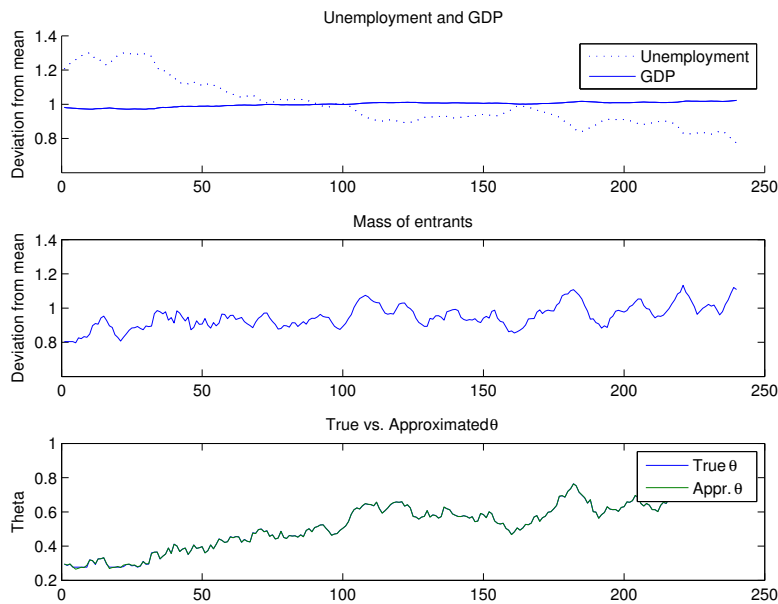
Figure 1-11: Sample simulation when the only shocks are to aggregate profitability. The first panel shows unemployment and GDP. The second panel shows the mass of entrants, and the last panel shows the true and approximated values of θ.

1.6.3 Evaluation of Results

This is a rich model in which the mapping from parameters to moments is not immediately clear. I therefore show several additional Figures here to help build some intuition for the results. Figure 10 shows results for a sample simulation of the benchmark model. ... We see that the model produces bursts of entry, particularly in reaction to changes in the aggregate shock a, which are larger than those observed in the data. Part of this is smoothed out by time aggregation, however.

The results of two sample simulations of the model without the financial friction and without shocks to aggregate profitability are shown in Figures 1-11 and 1-12 . In Figure 1-11 the only exogenous variation comes from changes in a. The first panel shows unemployment and GDP. The comovement between the two series is strong (the correlation between the two series in -0.995). For this reason the model is unable to generate jobless recoveries. An increase in unemployment can only result from a low realization of the aggregate shock a. However, once a returns to its unconditional mean the unemployment rate reverts back to its pre-recession value almost immediately. The second panel shows the mass of entrants, which reacts strongly to changes in a. In fact, the procyclicality of entry is around 60% larger than in the data. The last panel shows the true and the approximated laws of motion

41

Figure 1-12: Sample simulation when the only shocks are to HPI. The first panel shows unemployment and GDP. The second panel shows the mass of entrants, and the last panel shows the true and approximated values of θ.

of θ. The series are almost indistinguishable as labor market tightness θ moves virually 1:1 with the aggregate state. Figure 1-12 shows a simulation of the model when a is fixed at its unconditional mean. The only exogenous variation comes from movements in q^h. The first panel highlights that those exogenous shocks cause variations in the unemployment rate, while having almost no effect on GDP. As was discussed above, this is the main feature of the model which generates jobless recoveries. The mass of entrants, plotted in the second panel is much less volatile compared to the previous case where all exogenous shocks occured via a. Finally, the true and approximated law of motion of θ are shown in the third panel. Again, the fit is very good (see Appendix A.3 for details).

1.7 Conclusion

The recent recession which lasted from the end of 2007 until mid-2009 was severe in many respects. Because the unemployment rate remains far above its pre-crisis level the recovery has been described as *jobless*. Second, the recession was accompanied by an unprecedented fall in the value of real estate. In this paper I claim that these two facts are related. As the main channel through which house prices can exert this

influence on the unemployment rates I propose the process of lending to new firms. The model captures the idea that start-ups require external financing, for which real estate is used as collateral. As the value of this collateral falls, start-up costs increase and the number of new firms declines.

The number of start-ups in the US has fallen by over 20% since 2007. Never since the beginning of the data series in 1977 have there been as few openings of new firms or as few jobs created through firm birth than in 2010 and 2011. Young firms' below-trend job creation can account for almost all of the persistently high unemployment rate after the end of the recession.

I calibrate and compute a quantitative competitive industry model with endogenous entry and exit, firm heterogeneity, labor adjustment costs, and aggregate shocks. This model is able to match key moments of the firm distribution and employment at the micro- and macro-level. It captures the importance of new firms for employment and generates a jobless recovery. The model is able to explain over 80% of the increase and persistence in unemployment since 2007. I find that the effects of a 'technology shock' alone on the unemployment rate are neither strong nor persistent enough to fit the US data. I estimate that absent the deterioration of value of real estate, the increase in the unemployment rate would have been at around 40% of the actual increase. Furthermore, my mechanism generates a realistic procyclicality and time series variation in entry rates, something that previous studies have had difficulties with. Entry emerges as an important factor for the propagation of aggregate shocks. In contrast to previous studies my framework establishes a structural link between house prices, entrepreneurial activity, and the jobless recovery. This setup is suited to explain why start-up job creation began to decrease *prior* to the recent recession, and why - contrary to older, incumbent firms - it remains at low levels.

Chapter 2

Capital Reallocation and Aggregate Productivity

2.1 Motivation

Frictions in the reallocation of capital and labor are important for understanding aggregate productivity. With heterogenous plants, the assignment of capital, labor and other inputs across production sites impacts directly on aggregate productivity. Frictions in the reallocation process thus lead to the misallocation of factors of production (relative to a frictionless benchmark). This point lies at the heart of the analysis of productivity across countries in Hsieh and Klenow (2009), Bartelsman *et al.* (2013) and Restuccia and Rogerson (2008).[1]

In this paper we consider the **cyclical dimension of reallocation** in the presence of capital adjustment costs. In important empirical contributions, Eisfeldt and Rampini (2006) and Kehrig (2011) show that capital reallocation is pro-cyclical and that the cross-sectional productivity dispersion behaves counter-cyclically.[2] This not only underlines the importance of heterogeneity in the production sector but also suggests that frictions in the adjustment to capital may produce cyclical effects on output over the business cycle. One contribution of this paper is to specify a dynamic equilibrium model to further understand these findings about cyclical reallocation and dispersion in productivities.

Not properly taking cross-sectional heterogeneity into account will also lead to a mis-measurement of total factor productivity (TFP). We are interested in the cyclical

---

[1]More specific differences with these and other studies are discussed below.

[2]Eisfeldt and Rampini (2006) use dispersion in firm level Tobin's Q, dispersion in firm level investment rates, dispersion in total factor productivity growth rates, and dispersion in capacity utilization. Kehrig (2011) constructs dispersion measures based on TFP estimates using the estimation strategy in Olley and Pakes (1996).

component of the output loss resulting from frictions in the adjustment process which will be reflected in a mis-measured TFP. This relates to the question how micro-frictions like physical adjustment costs translate into aggregate outcomes. **We find that if the only shocks in the economy are to aggregate TFP, then the productivity loss from costly reallocation has no cyclical element.** This is consistent with results on the aggregate implications of lumpy investment, as in Thomas (2002), Khan and Thomas (2003) and Gourio and Kashyap (2007). If an aggregate model behaves as if there were no non-convexities at the plant-level, then the distortions in the allocation of capital across plants with different productivities will matter only for aggregate *levels*. As a result, the distribution over plants' capital stock and idiosyncratic productivity can be extremely well approximated by its first moment.

In addition to shocks to TFP, we also study shocks to plants' investment opportunities as in Eisfeldt and Rampini (2006), together with shocks to the distribution of idiosyncratic productivity as in Bloom (2009), Bloom *et al.* (2012), Gilchrist *et al.* (2014), or Bachmann and Bayer (2013). **Those shocks create cyclical movements in reallocation and productivity as well as time-varying productivity dispersion. Cross-sectional heterogeneity now plays an important role for shaping aggregate dynamics.** In the presence of those shocks, reallocation is correlated with measured aggregate productivity. The cross-sectional joint distribution over plants' capital stock and idiosyncratic productivity is a slow-moving object in this environment and tracking its evolution only by its first moment is insufficient: higher order moments are needed to characterize the outcome of the planner's problem, in particular the covariance of the cross-sectional distribution between plants' capital stocks and profitability.

Importantly these features of our model are interrelated. The fact that the covariance matters as a moment for determining the optimal allocation is indicative of the significance of reallocation effects. If this covariance did not matter for describing optimal allocations, for example because it is constant over time or perfectly correlated with the mean, then it could not have a cyclical effect on aggregate output. Thus the covariance that matters from the perspective of the Krusell and Smith (1998) approach is precisely the moment that reflects gains to capital reallocation.

This last point is worth stressing. Studies following Krusell and Smith (1998) routinely find that only first moments of distributions are needed to summarize cross sectional distributions. In our economy, the covariance of the cross sectional distribution between a plant's capital and its profitability is needed in the state space of the problem. When there are shocks either to the capital adjustment process or to the

cross sectional distribution, this covariance evolves in response to these shocks. In the presence of such shocks the approximate solution to the planner's problem using only average capital fails: the solution requires higher order moments.

As a final exercise, we study the business cycle properties of an economy driven by shocks to adjustment rates and to the cross sectional distribution of idiosyncratic shocks assuming constant aggregate total factor productivity.[3] This exercise provides a basis for "adverse" aggregate productivity shocks and the serial correlation of the Solow residual. **The aggregate moments produced by this economy are very similar to the moments of the standard stochastic growth model.** In particular: (i) the Solow residual is pro-cyclical and positively serially correlated, (ii) consumption, investment and output are positively correlated, (iii) consumption is smoothed, (iv) reallocation is pro-cyclical and (v) the standard deviation of productivity across plants is counter-cyclical. The first three properties match those of the standard RBC model. The last two properties match those stressed by Eisfeldt and Rampini (2006) and Kehrig (2011). In our setting a reduction in the Solow residual comes from variations in the distribution of shocks, not an adverse shock to total factor productivity.

2.2 Frictionless Economy

To fix basic ideas and notation, start with an economy with heterogeneity and no frictions. The planner maximizes

$$V(A, K) = max_{K', k(\varepsilon)} \ u(c) + \beta E_{A'|A} V(A', K') \tag{2.1}$$

for all (A, K). The constraints are

$$c + K' = y + (1 - \delta)K, \tag{2.2}$$

$$\int_{\varepsilon} k(\varepsilon) f(\varepsilon) d\varepsilon = K, \tag{2.3}$$

$$y = A \int_{\varepsilon} \varepsilon k(\varepsilon)^{\alpha} f(\varepsilon) d(\varepsilon). \tag{2.4}$$

The objective function is the lifetime utility of the representative household. The state vector has two elements: A is aggregate TFP and K is the aggregate stock of

---

[3]This analysis shares some features with Bloom *et al.* (2012) and Bachmann and Bayer (2013). Differences and similarities are made clear in the next sections.

capital. There is a distribution of plant specific productivity shocks, $f(\varepsilon)$ which is fixed and hence omitted from the state vector.

There are two controls in (2.1). The first is the choice of aggregate capital for the next period. The second is the assignment function, $k(\varepsilon)$, which allocates the given stock of capital across the production sites, indexed by their current productivity.

At the beginning of the period, A as well as the idiosyncratic productivity shocks ε realize. While aggregate capital K requires one period time-to-build, the reallocation of existing capital takes place instantaneously and is given by $k(\varepsilon)$.

The resource constraint for the accumulation of aggregate capital is given in (2.2). The constraint for the allocation of capital across production sites in given in (2.3), where $f(\cdot)$ is the distribution function for ε.

From (2.4), total output, y, is the sum of the output across production sites. The production function at any site is

$$y(k, A, \varepsilon) = A\varepsilon k^{\alpha} \tag{2.5}$$

where k is the capital used at the site with productivity ε.[4] The idiosyncratic productivity ε is persistent, parameterized by $\rho_{\varepsilon} \in [0, 1]$. We assume $\alpha < 1$ as in Lucas (1978).[5] In this frictionless environment, a plants' optimal capital stock is entirely determined by ε.

The assumption of diminishing returns to scale, $\alpha < 1$, implies that the allocation of capital across production sites is non-trivial. There are gains to allocating capital to high productivity sites but there are also gains, due to $\alpha < 1$, from spreading capital across production sites.

Let $k(\varepsilon) = \xi(\varepsilon)K$, so that $\xi(\varepsilon)$ is the fraction of the capital stock going to a plant with productivity ε. Then (2.4) becomes:

$$y = AK^{\alpha} \int_{\varepsilon} \varepsilon \xi(\varepsilon)^{\alpha} f(\varepsilon) d(\varepsilon) = AK^{\alpha}(\mu + \phi) \tag{2.6}$$

---

[4]Labor and other inputs are not made explicit. One interpretation is that these inputs have no adjustment costs and are optimally chosen each period, given the state. In this case, the marginal product of labor (and other inputs) will be equal across production sites. This does not imply equality of the marginal products of capital. Adding labor adjustment, perhaps interactive with capital adjustment, would be a natural extension of our model. Presumably, adding labor frictions would enhance our results. Bloom *et al.* (2012) include labor adjustment costs while Bachmann and Bayer (2013) assume flexible labor.

[5]As in Cooper and Haltiwanger (2006*a*), estimates of α are routinely below unity. This is interpreted as reflecting both diminishing returns to scale in production and market power due to product differentiation. For simplicity, our model ignores product differentiation and treats the curvature as reflecting diminishing returns. The analysis in Kehrig (2011) includes product differentiation at the level of intermediate goods.

where $\mu = \bar{\varepsilon} \int_\varepsilon \xi(\varepsilon)^\alpha f(\varepsilon) d(\varepsilon)$ and $\phi = Cov(\varepsilon, \xi(\varepsilon)^\alpha)$.[6] As is well understood from the Olley and Pakes (1996) analysis of productivity, aggregate output will depend on the covariance between the plant-level productivity and the factor allocation.

In the frictionless economy with time invariant distribution $f(\varepsilon)$ and costless reallocation of capital, this covariance is constant so that the joint distribution of plant-specific capital and ε is not part of the state vector. As this analysis progresses, this will not always be the case.

2.2.1 Optimal Choices

Within a period, the condition for the optimal allocation of capital across production sites is given by $\alpha A \varepsilon k(\varepsilon)^{\alpha-1} = \eta$ for all ε, where η is the multiplier on (2.3). This condition is intuitive: absent frictions, the optimal allocation equates the marginal product of capital across production sites.

Working with this condition,

$$k(\varepsilon) = \frac{\eta}{\alpha A \varepsilon}^{\frac{1}{\alpha-1}}. \tag{2.7}$$

Using (2.3),

$$\eta = A\alpha K^{\alpha-1} \left(\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon \right)^{1-\alpha}. \tag{2.8}$$

The multiplier is the standard marginal product on an additional unit of capital times the effect of the ε distribution on productivity.

Putting these two conditions together,

$$k(\varepsilon) = K \frac{\varepsilon^{\frac{1}{1-\alpha}}}{\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon}. \tag{2.9}$$

Substituting into (2.4) yields

$$y = AK^\alpha \left(\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon \right)^{1-\alpha}. \tag{2.10}$$

This is a standard aggregate production function, AK^α, augmented by a term that captures a "love of variety" effect from the optimal allocation of capital across plants. With a given distribution $f(\cdot)$ the idiosyncratic shocks magnify average aggregate productivity as the planner can reallocate inputs to the more productive sites.

---

[6]This uses $E(XY) = EX \times EY + cov(X,Y)$, where $\bar{\varepsilon}$ is the mean of the plant-specific shock.

The condition for **intertemporal optimality** is $u'(c) = \beta E V_K(A', K')$ so that the marginal cost and expected marginal gains of additional capital are equated. Using (2.1), this condition becomes

$$u'(c) = \beta E u'(c') \left[(1-\delta) + A'\alpha K'^{\alpha-1} \left(\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon \right)^{1-\alpha} \right]. \qquad (2.11)$$

The left side is the marginal cost of accumulating an additional unit of capital. The right side is the discounted marginal gain of capital accumulation. Part of this gain comes from having an extra unit of capital to allocate across production sites in the following period. The productivity from these production sites depend ons two factors, the **future** values of: aggregate productivity, A' and the cross sectional distribution of idiosyncratic shocks, $f(\varepsilon)$.

The choice of k for each plant within a period is independent of the choice between consumption and saving. The planner optimally allocates capital to maximize the level of output and then allocates output between consumption and capital accumulation. Clearly, once we allow for limits to reallocation, the capital accumulation decision will depend upon the future allocation of capital across production sites. In this way, variations in the distribution of $f(\cdot)$ can impact on the capital accumulation choice.

2.2.2 Aggregate Output and Productivity

For this economy, there is an interesting way to represent total output. This is seen from defining

$$\tilde{A} \equiv A \int_\varepsilon \varepsilon k(\varepsilon)^\alpha f(\varepsilon) d\varepsilon \qquad (2.12)$$

so that

$$y = \tilde{A} K^\alpha. \qquad (2.13)$$

from (2.4).

Researchers interested in measuring TFP from the aggregate data will typically uncover $\tilde{A}$ rather than A. This is the mis-measurement referred to earlier. As the discussion progresses, we will refer to $\tilde{A}$ as the Solow residual, as distinct from aggregate TFP.[7] There are three factors which influence $\tilde{A}$. The first one is A. The influence of A, aggregate TFP, on $\tilde{A}$, measured TFP, the Solow residual is direct and has been central to many studies of aggregate fluctuations. Second, the distribution $f(\varepsilon)$. Variations in $f(\varepsilon)$ influence $\tilde{A}$ because variations in the cross sectional distri-

---

[7]Thanks for Susanto Basu for urging us to make these terms clear.

bution of the idiosyncratic shocks lead to different marginal productivities of plants and thus changes in the Solow residual. Finally, there is the allocation of factors, k. If factors are optimally allocated, then the distribution of capital over plants does not have an independent effect on $\tilde{A}$. However, the existence of frictions may imply that, in a static sense, capital is not efficiently allocated. In that case, even with $f(\varepsilon)$ fixed, the reallocation process will lead to variations in $\tilde{A}$.[8]

Starting with Olley and Pakes (1996), many researchers have recognized the dependence of aggregate productivity on factor allocation. In many studies the underlying frictions are due to policies which influence steady state productivity across countries.[9] Our analysis differs from these studies in a couple of important ways. We next focus on (i) frictions through adjustment costs to capital, (ii) dynamic inefficiency brought about through the adjustment process so that the magnitude of the inefficiency and thus aggregate productivity are endogenous and (iii) the behavior of aggregate productivity over business cycles.

2.3 Capital Adjustment Costs

The allocation of capital over sites with heterogeneous idiosyncratic productivity has important effects on measured total factor productivity. In a frictionless economy there are no cyclical effects of reallocation on productivity. However, there is ample evidence in the literature for both non-convex and convex adjustment costs. Introducing these adjustment costs will enrich the analysis of productivity and reallocation.[10]

There are two distinct frictions to study, corresponding to the two dimensions of capital adjustment. The first is "costly reallocation" in which the friction is associated with the allocation of capital across the production sites. The second is "costly accumulation" in which the adjustment cost refers to the cost of accumulating rather than allocating capital.

Our focus here lies on studying the presence of costs to the reallocation (assignment) process. We introduce a special type of adjustment costs that is very tractable, although not very informative about the source of the friction. Following Calvo (1983)

---

[8]This decomposition of productivity taken from Olley and Pakes (1996) highlights the interaction between the distribution of productivity and factors of production across firms. Gourio and Miao (2010) use a version of this argument, see their equation (45), to study the effects of dividend taxes on productivity. Khan and Thomas (2008) study individual choice problems and aggregation in the frictionless model with plant specific shocks. Basu and Fernald (1997) also discuss the role of reallocation for productivity in an aggregate model.

[9]Bartelsman *et al.* (2013) discuss these other studies in their analysis of productivity differences over 24 economies.

[10]In contrast to **?**, there are no borrowing frictions. They argue that these frictions do not create large losses in aggregate productivity.

and more recently adopted to study investment decisions by Sveen and Weinke (2005), assume that each period a Bernoulli draw determines the fraction $\pi \in [0, 1]$ of plants the planner can costlessly reallocate capital between. This represents a stochastic investment opportunity. The remaining fraction of plants $1 - \pi$ produces with its beginning-of-period capital stock. This structure of adjustment costs captures the fact that plants adjust their capital stock infrequently. Applying a law of large numbers, the plant-specific shocks ε are assumed to be equally distributed over the fractions π and $1 - \pi$ of adjustable and non-adjustable plants. The two distributions of plants will be referred to as F^a and F^n. This also implies that $E(\varepsilon)$ is time-invariant and the same across adjustable and non-adjustable plants.

By assumption, π is not dependent on the state of the plant. This simplification makes our analysis tractable. At the same time it does not preclude a role for the cross sectional distribution in the state space of the problem. Besides tractability, there are other arguments for this specification.

First, a model with just non-convex adjustment costs, or a mixture of non-convex and quadratic adjustment costs, as in Cooper and Haltiwanger (2006a), captures inaction and bursts of investment but misses small adjustments. While not as elegant as the state dependent adjustment model, the constant hazard structure does generate inaction, bursts of investment as well as smaller adjustment rates. A similar point about price adjustment is used in Midrigan (2011) to justify a constant adjustment rate specification[11]

Second, the focus of our analysis is on (2.12): the impact of the cross sectional distribution of profitability shocks on the Solow residual and thus output. The constant hazard assumption allows us to isolate the effects of the cross sectional distribution through its effects on the allocation of capital and hence output rather than through adjustment costs alone. This does not deny the significance of adjustment costs but rather focuses solely on the output effects of the cross sectional distribution. There is an important cost to this specification: there is no option value of waiting. In a model with non-stochastic fixed costs, if adjustment is not made in the current period, it is available for sure in the next one. Once adjustment costs are stochastic, the option value of waiting is reduced.

2.3.1 The Planner's Problem

For the dynamic program of the planner in the presence of adjustment costs, the state vector contains aggregate productivity A, the aggregate capital stock K, and

---

[11]See also Costain and Nakov (2013).

Γ. The high-dimensional object Γ describes the joint distribution over capital (at the start of the period) and productivity shocks across plants. Γ is needed in the state vector because the presence of adjustment costs implies that a plant's capital stock may not reflect the current draw of ε. As noted above, there is time variation in the probability of adjustment π. Furthermore, there are shocks to the variance of idiosyncratic productivity shocks, parameterized by λ. Changes in the variance of the cross-sectional idiosyncratic productivity, as recently highlighted in Bloom (2009) and Gilchrist *et al.* (2014), have an effect on output. Such changes can be interpreted as variations in uncertainty. Consider a mean-preserving spread (MPS) in the distribution of ε. In a frictionless economy such a spread would incentivize the planner to carry out more reallocation of capital between plants because capital can be employed in highly productive sites. Let $s = (A, K, \Gamma, \lambda, \pi)$ denote the vector of aggregate state variables. Note the assumed timing: changes in the distribution of idiosyncratic shocks are known in the period they occur, not in advance.[12] The adjustment status of a plant is given by $j = a, n$, where a stands for 'adjustment', while n stands for 'non-adjustment'.

Given the state, the planner makes an investment decision K' and chooses how much capital to reallocate across those plants whose capital stock can be costlessly reallocated, $(k, \varepsilon) \in a$. Let $\tilde{k}_j(k, \varepsilon, s)$ for $j = a, n$ denote the capital allocation to a plant that enters the period with capital k and profitability shock ε in group j after reallocation. The capital of a plant in group $j = a$ is adjusted and is optimally set by the planner to the level $\tilde{k}_a(k, \varepsilon, s)$. The capital of a plant in group $j = n$ is not adjusted so that $\tilde{k}_n(k, \varepsilon, s) = k$.

The choice problem of the planner is:

$$V(A, K, \Gamma, \lambda, \pi) = max_{\tilde{k}_a(k, \varepsilon, s), K'} \ u(c) + \beta E_{[A', \Gamma', \lambda', \pi'|A, \Gamma, \lambda, \pi]} V(A', K', \Gamma', \lambda', \pi') \quad (2.14)$$

subject to the resource constraint (2.2) and

$$y = \int_{(k, \varepsilon) \in F^a} A\varepsilon \tilde{k}_a(k, \varepsilon, s)^\alpha d\Gamma(k, \varepsilon) + \int_{(k, \varepsilon) \in F^n} A\varepsilon \tilde{k}_n(k, \varepsilon, s)^\alpha d\Gamma(k, \varepsilon), \quad (2.15)$$

which is simply (2.4) split into adjustable and non-adjustable plants. Here F^j is the set of plants in group $j = a, n$. The fraction of plants whose capital stock can be

---

[12]Other models, such as Bloom *et al.* (2012), include future values of λ in the current state as a way to generate a reduction in activity in the face of greater uncertainty about the future. We include the implications of this alternative timing as part of the results below.

adjusted is equal to π

$$\int_{(k,\varepsilon)\in F^a} f(\varepsilon)d\varepsilon = \pi \qquad (2.16)$$

and the amount of capital over all plants must sum to total capital K:

$$\pi \int_{(k,\varepsilon)\in F^a} \tilde{k}_a(k,\varepsilon,s)d\Gamma(k,\varepsilon) + (1-\pi) \int_{(k,\varepsilon)\in F^n} \tilde{k}_n(k,\varepsilon,s)d\Gamma(k,\varepsilon) = K. \qquad (2.17)$$

As the capital is plant specific, it is necessary to specify transition equations at the plant level. Let $i = \frac{K'-K}{K}$ denote the gross investment rate so that $K' = (1-\delta+i)K$ is the aggregate capital accumulation equation. To distinguish reallocation from aggregate capital accumulation, assume that the capital at **all** plants, regardless of their reallocation status, have the same capital accumulation. The transition for the capital (after reallocation) this period and the initial plant-specific capital next period is given by

$$k'_j(k,\varepsilon,s) = (1-\delta+i)\tilde{k}_j(k,\varepsilon,s), \qquad (2.18)$$

for $j = a, n$. Due to the presence of frictions $\tilde{k}_a(k,\varepsilon,s)$ is not given by (2.9). Notice that A affects unadjustable and adjustable plants in the same way. This implies that the optimal reallocation decision will occur independently of A. The shock to A will have an effect on the mis-measured part of TFP only in the presence of a capital accumulation problem, since the total amount of capital in adjustable and non-adjustable plants may differ.

The quantitative analysis will focus on reallocation of capital, defined as the fraction of total capital that is moved between adjustable plants within a period. Following a new realization of idiosyncratic productivity shocks, the planner will reallocate capital from less productive to more productive sites. Aggregate output is thus increasing in the amount of capital reallocation.

As $\tilde{k}_a(k,\varepsilon,s)$ denotes the post-reallocation capital stock of a plant with initial capital k, the plant-level reallocation rate would be $r(k,\varepsilon,s) = |\frac{\tilde{k}_a(k,\varepsilon,s)-k}{k}|$. Aggregating over all the plants who adjust, the aggregate reallocation rate is

$$R(s) \equiv 0.5 \int_{(k,\varepsilon)\in F^a} r(k,\varepsilon,s)d\Gamma(k,\varepsilon). \qquad (2.19)$$

The multiplication by 0.5 is simply to avoid double counting flows between adjusting plants.

2.3.2 Joint Distribution of Capital and Productivity

In the presence of reallocation frictions, the state space of the problem includes the cross sectional distribution, Γ. Consequently, when making investment and reallocation decisions the planner needs to forecast Γ'. It is computationally not feasible to follow the joint distribution of capital and profitability shocks over plants, we represent the joint distribution by several of its moments. These forecast the marginal benefit of investment.

The right set of moments is suggested by the following expression for aggregate output, taken from (2.15)

$$y = \pi(\bar{\varepsilon}\mu_a + \phi_a) + (1 - \pi)(\bar{\varepsilon}\mu_n + \phi_n), \tag{2.20}$$

where $\mu_j \equiv E(\tilde{k}_j(k, \varepsilon, s)^\alpha)$ and $\phi_j \equiv Cov(\varepsilon, \tilde{k}_j(k, \varepsilon, s)^\alpha)$, for $j = a, n$. Instead of Γ we retain μ_n and ϕ_n in the state vector of (2.14).

These two moments contain *all* the necessary information about the joint distribution of capital and profitability among non-adjustable plants. The information about capital in plants in F^A is not needed since capital in those plants can be freely adjusted, independently of their current capital stock. Together, μ_n and ϕ_n are sufficient to compute the output of those plants whose capital cannot be reallocated and thus to solve the planner's optimization problem. Note that by keeping μ_n and ϕ_n in the state space, we are not *approximating* the joint distribution over capital and productivity since the two moments can account for all the variation of the joint distribution. This feature of our choice of moments allows us to compare it with common approximation techniques in the spirit of Krusell and Smith (1998).

The covariance term ϕ_n is crucial for understanding the impact of reallocation on measures of aggregate productivity. If the covariance is indispensable in the state vector of the planner, then the model is not isomorphic to the stochastic growth model. That is, if the covariance is part of the state vector, then the existence of heterogeneous plants along with capital adjustment costs matters for aggregate variables like investment over the business cycle.

When either A or π is stochastic, it is possible to follow the evolution of these moments analytically.[13] The choice of $\tilde{k}_a$ for adjustable plants, along with the respective ε shocks at these plants, maps into values of the moments μ_a and ϕ_a. Together with the new realization of exogenous shocks at the beginning of the next period these map into the next period moments μ'_n and ϕ'_n. The laws of motion for the two states

---

[13]The analytics hold for the evolution of the mean, (2.21), but not the covariance, (2.22), when λ is stochastic.

μ_n and ϕ_n are given by

$$\mu'_n = \pi'\mu_a + (1 - \pi')\mu_n \qquad (2.21)$$

and

$$\phi'_n = \pi'\rho_\varepsilon\phi_a + (1 - \pi')\rho_\varepsilon\phi_n. \qquad (2.22)$$

Together these laws of motion define the law of motion of the joint distribution Γ, allowing us to follow the evolution of this component of the aggregate state.[14] Equations (2.20)-(2.22) permit us to study the trade-off regarding the optimal allocation of capital across sites. The planner can increase contemporaneous output by reallocating capital from low- to high-productivity sites in F^a. This will increase the covariance between profitability and capital, ϕ_a, while at the same time decreasing μ_a because $\alpha < 1$. A fraction $1 - \rho_\pi$ of currently adjustable plants will not be able to adjust its capital stock tomorrow. The planner therefore has to trade off the higher instantaneous output from reallocation with the higher probability of a mismatch between $\tilde{k}_n(k, \varepsilon, s) = k$ and the realization of ε' for plants in F^n tomorrow. This is captured in the laws of motion (2.21) and (2.22).

2.3.3 Stationary Equilibria

To fix ideas we can analyze the stationary economy where π and λ are not varying over time. In this environment a stationary distribution Γ^* exists. Using (2.21) it follows that $\mu_n = \mu_a = \mu^*$. Furthermore, stationary values ϕ_a^* and ϕ_n^* exist. Using (2.22) one can show that ϕ_n converges to

$$\phi_n^* = \phi_a^* \frac{\pi\rho_\varepsilon}{1 - (1 - \pi)\rho_\varepsilon}. \qquad (2.23)$$

Hence (2.20) becomes

$$y = \bar{\varepsilon}\mu^* + \Lambda\phi_a^*, \qquad (2.24)$$

where $\Lambda \equiv \frac{\pi}{1-(1-\pi)\rho_\varepsilon}$ is a function of parameters. Λ is (weakly) increasing in both π and ρ_ε.[15] Intuitively, an increase in π increases total output because more plants' capital stock can be costlessly adjusted. An increase in ρ_ε, the persistence of idiosyncratic productivity shocks, implies that the probability of a plant switching status and being

---

[14]Note that $\phi' = Cov(k(\varepsilon)^\alpha, \varepsilon')$ is an expectation. The term ε' is made up of two components, one is the persistent part, and one is an i.i.d. part, denoted η. Rewrite $\varepsilon' = \rho_\varepsilon\varepsilon + (1 - \rho_\varepsilon)\eta$ to obtain $\phi' = Cov(k(\varepsilon)^\alpha, \rho_\varepsilon\varepsilon + (1 - \rho_\varepsilon)\eta) = \rho_\varepsilon\phi$.

[15]Formally, $\frac{\partial\Lambda}{\partial\pi} = \frac{1-\rho_\varepsilon}{[1-(1-\pi)\rho_\varepsilon]^2} \geq 0$, $\frac{\partial\Lambda}{\partial\rho_\varepsilon} = \frac{\pi(1-\pi)}{[1-(1-\pi)\rho_\varepsilon]^2} \geq 0$. The cross-derivatives are given by $\frac{\partial^2\Lambda}{\partial\rho_\varepsilon\partial\pi} = \frac{\partial^2\Lambda}{\partial\pi\partial\rho_\varepsilon} = \frac{1}{[1-(1-\pi)\rho_\varepsilon]^2} - \frac{2\pi}{[1-(1-\pi)\rho_\varepsilon]^3}$.
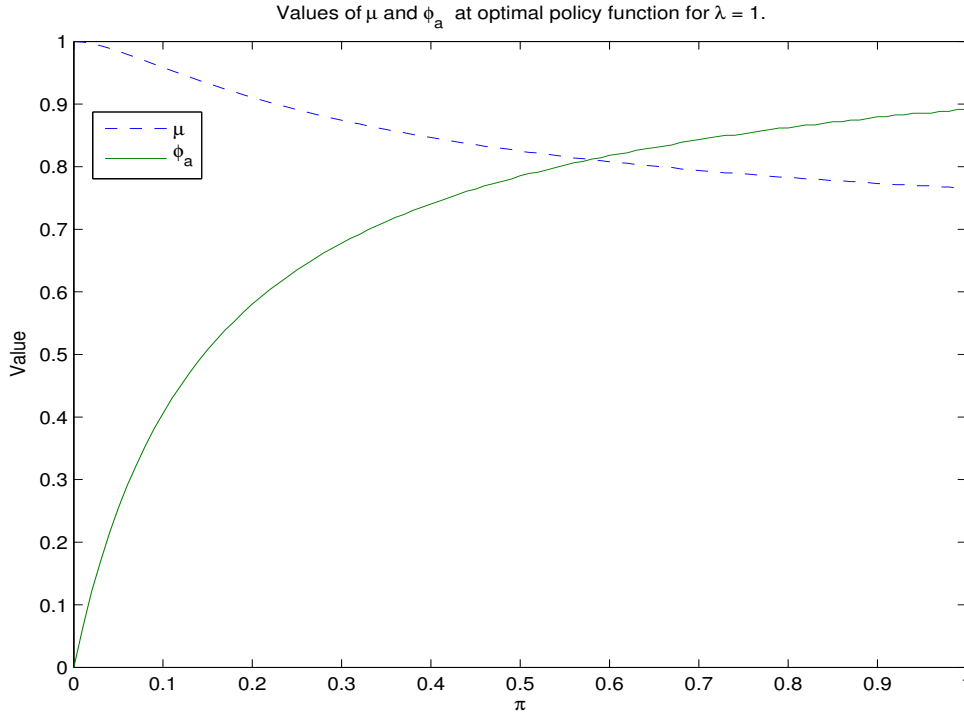
Figure 2-1: Values of μ and ϕ_a in stationary equilibrium for various π. Economy with $\lambda = 1$ and $\rho_\varepsilon = .9$

non-adjustable with a mismatch between ε and k is decreased.[16]

Figure 2-1 shows equilibrium values of μ^* and ϕ_a^* in stationary economies for different values of π. As $\pi \to 0$ the planner reallocates less capital between plants. A value of $\mu^* = 1$ implies $\phi_a^* = 0$, because $k(\varepsilon) = 1$ for all sites, meaning that the capital level is independent of ε. On the other hand, as the fraction of adjustable plants increases, ϕ_a^* increases.

2.4 Quantitative Results

With exogenous movements in π and λ no stationary distribution of Γ exists and the two moments μ_n and ϕ_n become part of the state vector. This problem can no longer be solved analytically. This section presents quantitative results.

In the stationary economy, reallocation effects only mattered for aggregate **levels**. When are reallocation effects likely to play a role for aggregate **dynamics**? One key prerequisite is that the economy be subject to shocks that cause the distribution Γ to

---

[16]In the extreme case of iid shocks to idiosyncratic productivity shocks the planner would be more reluctant to allocate large amounts of capital to high-productivity sites, decreasing aggregate output.

move over time. Without movements in Γ the benefits from reallocation are constant and the covariance term ϕ is not required to forecast Γ'. The reasons why Γ may vary and the implications of its variability will be clear as the analysis proceeds.

In keeping with the distinction noted earlier between reallocation and accumulation, the initial quantitative analysis, presented in section 2.4.1 is for an economy with a fixed capital stock, thus highlighting reallocation. The economy is then enriched to allow for capital accumulation in section 2.4.2.

For each of these models, this section focuses on the effects of capital reallocation on aggregate productivity. In addition, we present evidence on whether higher order moments are needed in the solution of the planner's optimization problem in Section 2.5. As highlighted in the introduction, these two themes are connected: higher order moments are needed to follow the evolution of Γ precisely when capital reallocation matters for the cyclical movements in productivity.

We solve the model at a quarterly frequency, using these **baseline** parameters. Following the estimates in Cooper and Haltiwanger (2006a), we set $\alpha = 0.6$.[17] We assume log-utility and a depreciation rate $\delta = 0.025$. Assuming an annual interest rate of 4% implies a discount factor $\beta = 0.987$. We set the mean of π to $\bar{\pi} = 0.5$. This implies that plants adjust their capital stock on average every two quarters. Sveen and Weinke (2005) treat changes in the capital stock of under 10% in absolute value as maintenance and hence use $\pi = 0.08$. In our setup, the choice of π mainly affects aggregate *levels*, not transitions. Aggregate profitability takes the form of an AR(1) in logs

$$\ln a_t = \rho_a \ln a_{t-1} + \nu_{a,t}, \qquad \nu_a \sim N(0, \sigma_a), \tag{2.25}$$

where $\rho_a = 0.9$ and $\sigma_a = 0.005$. Idiosyncratic profitability shocks are log-normal and evolve according to a law of motion with time-varying variance

$$\ln \varepsilon_t = \rho_\varepsilon \ln \varepsilon_{t-1} + \lambda_t \nu_{\varepsilon,t}, \qquad \nu_{\varepsilon,t} \sim N(0, \sigma_\varepsilon). \tag{2.26}$$

The parameters of the idiosyncratic shock process are $\rho_\varepsilon = 0.9$ and $\sigma_\varepsilon = 0.2$. The parameter λ governs the mean-preserving spread of the normal distribution from which idiosyncratic profitability ε is drawn. It has a mean of 1 and variance σ_λ

$$\lambda_t = \rho_\lambda \lambda_{t-1} + \nu_{\lambda,t}, \qquad \nu_{\lambda,t} \sim N(1, \sigma_\lambda). \tag{2.27}$$

---

[17]This curvature is 0.44 in Bachmann and Bayer (2013) and 0.4 in Bloom *et al.* (2012).

We set $\rho_\lambda = 0.82$ as in Gilchrist *et al.* (2014). Finally, the process of π follows

$$\pi_t = \rho_\pi \pi_{t-1} + \nu_{\pi,t}, \qquad \nu_{\pi,t} \sim N(\bar{\pi}, \sigma_\pi), \qquad (2.28)$$

with $\rho_\pi = 0.9$. In order to be able to compare the effect of different shocks, the standard deviations of the innovations, $\sigma_\pi = 0.03$ and $\sigma_\lambda = 0.014$ are set to generate the same amount of variation in output as shocks to A. Section 2.4.4 explores the sensitivity of our findings to this parameterization. The number of plants is set at 10,000 for these simulations. The computational strategy is discussed in further detail in the Appendix.

2.4.1 Capital Reallocation

Table 2.1 shows measures of the efficiency of the allocation of capital and the cyclicality of the Solow residual. These two aspects of the economy are inherently linked. Aggregate productivity is endogenous and responds to changes in the amount of capital reallocated.

The column labeled 'R/R^{*}' for 'Reallocation' measures the time series average of the cross-sectional reallocation of capital across plants as defined in (2.19), relative to the frictionless benchmark without adjustment costs. The column labeled $E_t(\sigma_i(arpk_{it}))$ measures the time series average of the cross sectional standard deviation of the average revenue product of capital. The column labeled G shows the output gap, defined as $G(s) = \frac{y^{FL}(s) - y(s)}{y^{FL}(s)}$, output in state s relative to the frictionless benchmark.[18] The column labeled $\sigma(\tilde{A}/A)$ reports the standard deviation of the Solow residual relative to TFP. The columns $C(R, \tilde{A})$ and $C(\sigma_i(arpk_{it}), \tilde{A})$ show the correlation between the Solow residual and respectively capital reallocation and the standard deviation of the average revenue product of capital. These two columns provide a link back to the facts, noted in the introduction, about the cyclical behavior of reallocation and dispersion in productivity.

The first block of Table 2.1 reports results for the frictionless economy. The second block of results introduces capital adjustment costs.

Frictionless Economy

The first row of Table 2.1 shows the results for the frictionless economy, $\pi = 1$, without time series variations in TFP, the volatility of the idiosyncratic shocks λ,

---

[18]The frictionless output $y^{FL}(s)$ is a function of s because changes in λ affect the output achieved in the frictionless case.

Case	R/R^*	$E_t(\sigma_i(arpk_{it}))$	G	$\sigma(\tilde{A}/A)$	$C(R,\tilde{A})$	$C(\sigma_i(arpk_{it}),\tilde{A})$
Frictionless						
nonstochastic	1 (0)	0 (0)	0 (0)	0 (0)	na $(-)$	na $(-)$
stochastic A	1 (0)	0 (0)	0 (0)	0 (0)	na $(-)$	na $(-)$
stochastic λ	1 (0)	0 (0)	0 (0)	0.083 (0.002)	0.950 (0.006)	na $(-)$
Frictions						
nonstochastic	0.491 (0)	1.09 (0)	0.106 (0)	0 (0)	na $(-)$	na $(-)$
stochastic A	0.491 (0)	1.09 (0)	0.106 (0)	0 (0)	na $(-)$	na $(-)$
stochastic π	0.487 (0.008)	1.09 (0.007)	0.110 (0.001)	0.077 (0.004)	0.995 (0.004)	-0.977 (0.002)
stochastic λ	0.491 ($3.68e-06$)	1.09 (0.01)	0.106 ($7e-05$)	0.064 (0.001)	0.936 (0.005)	0.929 (0.006)
stochastic π,λ	0.492 (0.005)	1.09 (0.006)	0.108 (0.0008)	0.10 (0.003)	0.817 (0.012)	-0.194 (0.03)

Table 2.1: Capital Reallocation Model: Productivity Implications

Results from 100 simulations with T=1000, standard deviations in parentheses below. $\frac{R}{R^*}$ measures the time series average of the cross-sectional reallocation of capital across plants, relative to the frictionless benchmark, R^*. $E_t(\sigma_i(arpk_{it}))$ is the mean standard deviation of the average revenue product of capital. G refers to the output gap relative to the frictionless benchmark. The column $\sigma(\tilde{A}/A)$ shows the standard deviation of measured vs. real TFP. The last columns $C(R,\tilde{A})$ and $C(\sigma_i(arpk_{it}),\tilde{A})$ show the correlation between mismeasured TFP and respectively capital reallocation and the standard deviation of the average revenue product of capital. The "na" entry means that the correlation is not meaningful as one of the variables is constant.

or the fraction of adjustable sites π.[19] This case serves as a benchmark. Without frictions, the marginal product of capital is equalized across plants and our measure of the inefficiency of the capital allocation, $E_t(\sigma_i(arpk_{it}))$, is zero. The first-best output is achieved. The mis-measurement of TFP is constant. The amount of capital reallocation is time-invariant and hence plays no role for aggregate productivity.

The second row, 'stochastic A' introduces variation in aggregate profitability. Variations in A have no effect on the reallocation of capital in this economy, because the planner reallocates capital across plants within a period. Consequently the amount of reallocation is the same as without variations in A. The allocation is efficient, $\tilde{A}$ varies only with A. The only difference with respect to the benchmark in the previous row is the variability of output, which is driven by changes in aggregate profitability. Since A enters total output multiplicatively all variation in output stems from variation in A. There is no endogenous propagation. As before, the amount of capital reallocation is time-invariant.

---

[19]In this abbreviated problem, the planner solves $V(\Gamma) = max_{k(\varepsilon)} \ u(c) + \ \beta EV(\Gamma')$ subject to the resource constraint (2.2) and total production given by (2.15).

The third row 'stochastic λ' presents results for the frictionless economy with stochastic variance of idiosyncratic productivity shocks. The parameter λ is chosen to generate the same coefficient of variation of output as the previous case.[20] The resulting allocation has the same rate of reallocation as the benchmark and the cross sectional distribution of the average revenue product of capital is degenerate. Importantly, output and the Solow residual vary with λ, as shown in column $\sigma(\tilde{A}/A)$. This represents a pure reallocation effect through changes in $f(\varepsilon)$ and occurs even under constant A and π. The second to last column shows the high correlation between the amount of capital reallocation and output. The correlation is not equal to one because following a shock to λ, the subsequent change in the planner's chosen allocation of capital produces an overshooting of output. This is a result of the allocation of capital among non-adjustable plants.

This economy presents the simplest case where reallocation is the sole driver of business cycles. To some degree, it looks like an economy driven by exogenous TFP. Here the variations in productivity arise from the endogenous reallocation of capital. The following subsection studies environments where capital adjustment costs amplify this feature.

Costly Capital Reallocation

Setting $\pi < 1$ introduces capital adjustment costs to the frictionless economy, so that only a fraction of all plants' capital stocks can be adjusted within a given period. Costly capital reallocation will have effects on measured productivity and its cyclical properties.

When π is non-stochastic and there are no other aggregate shocks, a stationary joint distribution Γ exists, with the moments (μ_n, ϕ_n) constant, as was shown in Section 2.3.3 above. Table 2.1 shows the results for this case in the row labeled 'nonstochastic'. In this economy the fraction of capital reallocated is far below the frictionless benchmark, as indicated in the second column. With $R < \pi$, the planner's chosen distribution of capital over adjustable plants is different from the distribution in the frictionless case. Although capital in a fraction π of plants could be costlessly reallocated, the reallocation rate is less than π. Instead, reallocation is lower indicating a reduced capital flow beyond the direct influence of $\pi < 1$.

Figure 2-2 plots capital reallocation as a function of π. The dashed line is the

---

[20]For this case, λ takes values between 0.966 and 1.0344. These values are chosen to generate the same amount of output volatility as direct shocks to A. Below we study the implications of larger variability in λ. Note that $\lambda > 1$ can imply that some values of the shock become negative. To avoid this, we apply the MPS to the underlying normal distribution and re-adjust its mean such that mean of the log-normal is preserved.

45° line. The concave solid green line above it shows capital reallocation between adjustable plants (as a fraction of the frictionless benchmark). As $\pi \to 1$ it approaches the allocation derived in (2.9). For total capital reallocation (plotted as the red solid line beneath the 45° line) this implies that it approaches π as $\pi \to 1$.
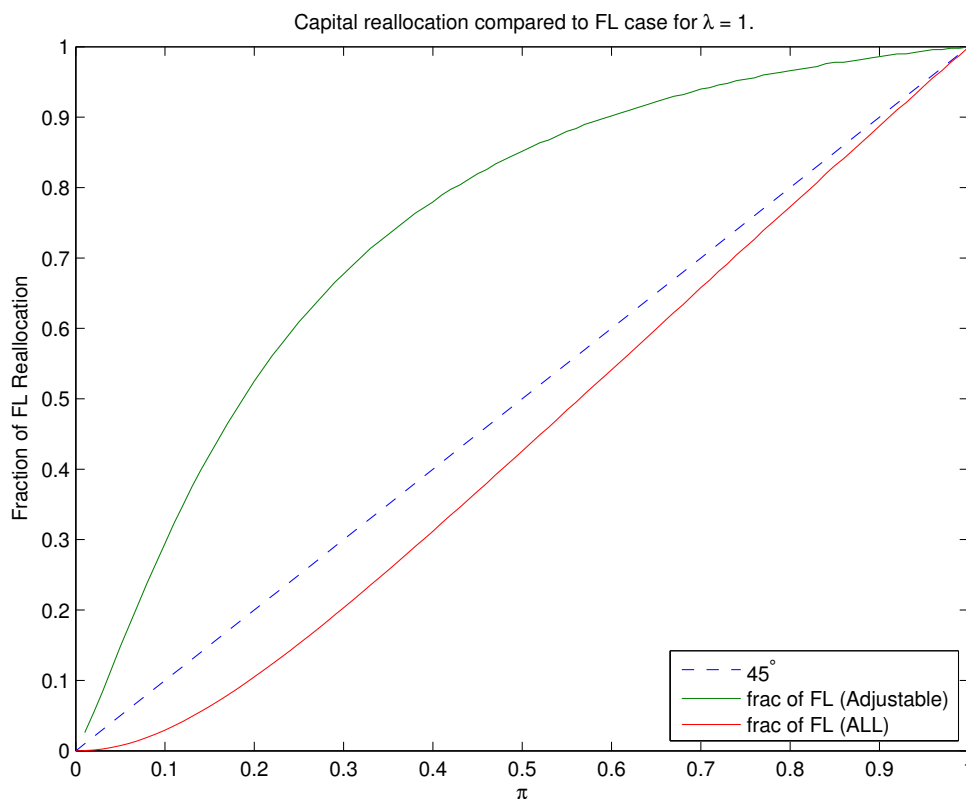
Figure 2-2: Capital Reallocation in adjustable and all plants as fraction of frictionless benchmark in stationary equilibrium for various π. Economy with $\lambda = 1$ and $\rho = .9$.

The inefficiency of the allocation when $\pi < 1$ is highlighted by the column labeled $E_t(\sigma_i(arpk_{it}))$. This measure of the inefficiency of the allocation is larger than zero, reflecting frictions in the reallocation process that stem from two sources. First, the planner chooses not to equalize marginal products between adjustable plants, reflecting the tradeoffs discussed above. Secondly, the marginal products of capital among non-adjustable plants exhibit a high degree of heterogeneity due to the fact that their capital is fixed despite a new realization of idiosyncratic profitability. Because ϕ_n and μ_n converge to their steady-state values output does not vary in this economy. The output gap is positive, directly reflecting the impact of $\pi < 1$. Importantly, the mis-measurement in TFP is constant over time, we only obtain a level-effect.

The row labeled 'stochastic A' allows for randomness in aggregate productivity

with constant π. As explained above, the amount of reallocation is independent of variations in A. Output and $\tilde{A}$ vary only with A. Because $\pi < 1$ the allocation is characterized by a positive standard deviation of average revenue products of capital and a positive output gap.

Variations in π create time series variation in the moments μ_n and ϕ_n, as shown in the row 'stochastic π'. Fluctuations in π lead to pro-cyclical capital reallocation patterns, as shown in column $C(R, \tilde{A})$. But this is not *simply* a correlation. In the presence of adjustment frictions, reallocation **causes** the observed time-variations in output. Variations in π therefore also lead to variations in (mis-measured) total factor productivity. The marginal products of capital are not equalized across plants, neither among the adjustable nor the unadjustable sites. This results in a positive output gap which varies with the evolution of μ_n and ϕ_n. This gap is about 11% of real GDP. Additionally, this economy exhibits counter-cyclical productivity dispersion, as seen in the last column. When π is low, less capital can be reallocated between adjustable plants. This decreases output and increases the standard deviation of marginal products between those plants. Though λ is held fixed, $\sigma_i(arpk_{it})$ nonetheless varies over time.

The row 'stochastic λ' of Table 2.1 studies the effects of time-variation in $f(\varepsilon)$ under costly capital reallocation. Due to the presence of adjustment costs, the marginal products of capital cannot be equalized over time. In addition, the variations in λ lead to changes in the optimal allocation decision by the planner and create considerable time-variation in μ_n and ϕ_n. The resulting fluctuations in output stem from different reallocation choices of the planner that show up in variations of the Solow residual. While variations in π affect output directly through the fraction of plants among which capital can be reallocated, the effect of changes in λ is less direct. Variations in λ induce different reallocation choices but a fraction of the effect on output comes from the fact that the marginal revenue product of capital is changed through productivity draws with larger or smaller tails. As the last two columns show, shocks to λ lead to pro-cyclical reallocation patterns. At the same time they produce a pro-cyclical dispersion in average revenue products of capital. A larger spread in the distribution of shocks leads to more reallocation of capital among adjustable plants by the planner and hence higher output. At the same time the increase in dispersion leads to a larger standard deviation of the marginal products of capital, both among adjustable and non-adjustable plants. This results is driven by the probability of a mismatch between k and ε' for plants in F^n.

The joint effects of changes in π and λ are presented in the last row of Table 2.1. Output varies significantly over time, with variations resulting directly from both

shocks to π and λ. While $\pi < 1$ leads to a positive output gap the presence of a stochastic λ causes additional variation in this gap as was the case before. Notably, mis-measured TFP exhibits significantly more time variation than in the cases of varying λ or varying π alone. This is the result of changes in π and λ jointly affecting the slow-moving joint distribution Γ. Importantly, the correlation between capital reallocation and output is much lower in this environment. This comes about because mis-measured TFP reacts more strongly through changes in λ than π. On the other hand, both exogenous shocks affect the amount of reallocation. The effect of varying π on reallocation, however, is predominantly an extensive margin effect, as a changing fraction of plants can reallocate capital. The effect of λ is on the intensive margin: more capital is reallocated within a given fraction of adjustable plants. Together this explains the observed decrease in the correlation between reallocation and output.

Overall, adjustment frictions reduce reallocation, generating a non-degenerate distribution of average (and marginal) products of capital across plants. The cost is a reduction in output of about 11%, relative to the frictionless benchmark. In all of the experiments, reallocation is pro-cyclical. For these cases, measured variations in TFP are the consequence of reallocation rather than true variations in aggregate productivity. Variations in π lead to counter-cyclical productivity dispersion across firms.

The economy with variations in both π and λ mimic the patterns of pro-cyclical reallocation and counter-cyclical dispersion emphasized by Eisfeldt and Rampini (2006). This will be a leading case as the analysis proceeds.

2.4.2   Endogenous Capital Accumulation

With endogenous capital accumulation, solving (2.14), the capital *reallocation* process has significant interactions with the capital *accumulation* decision. The frictions exert a level effect on the optimal capital stock and induce different dynamics following an exogenous shock. As we saw above, reallocation behaves cyclically in the presence of time-series variation in π and/or λ. Variations in λ and π affect the instantaneous value of existing capital and, because of persistence, the expected future return to capital, too. This affects the planner's incentives to invest. Even absent any frictions to capital accumulation the dynamics of investment and consumption are considerably altered by the presence of exogenous shocks to reallocation or the variance of the idiosyncratic shock.

Adding endogenous capital accumulation does not alter the results on the reallocation process shown in Table 2.1. The reason parallels the argument for the independence of reallocation from A. From (2.10), total output is proportional to

AK^α. Thus just as variations in A scale moments, so will variations in K. Consequently, the analysis focuses on the effects of frictions in reallocation on capital accumulation.

Table 2.2 summarizes results for the endogenous capital accumulation problem, using the baseline parameters, defined earlier. The aggregate capital stock is now endogenous and creates additional variation. The average capital stock (relative to the frictionless benchmark) is shown in the $\bar{K}/\bar{K}^*$ column. The other columns report correlations of reallocation with investment and output, $C(R,i)$ and $C(R,y)$ and the correlation of investment and the Solow residual, $C(\tilde{A},i)$.

Case	$\bar{K}/\bar{K}^*$	$C(R,i)$	$C(R,y)$	$C(\tilde{A},i)$
Frictionless				
nonstochastic	1 (0)	na (−)	na (−)	na (−)
stochastic A	1 (0)	na (−)	na (−)	na (−)
stochastic λ	1 (0)	0.94 (0.01)	0.90 (0.01)	0.99 (0.001)
Frictions				
nonstochastic	0.75 (0)	na (−)	na (−)	na (−)
stochastic A	0.75 (0)	na (−)	na (−)	0.955 (0.09)
stochastic π	0.75 (0.005)	0.97 (0.01)	0.91 (0.003)	0.97 (0.005)
stochastic λ	0.75 (0.0006)	0.93 (0.01)	0.88 (0.01)	0.979 (0.001)
stochastic π, λ	0.75 (0.003)	0.790 (0.01)	0.767 (0.02)	0.964 (0.01)

Table 2.2: Endogenous Capital Accumulation: Aggregate Moments

Results from 100 simulations with T=1000, N=10,000 are reported with standard deviations in parentheses below. Simulations with frictions were computed with a mean of π equal to 0.5, mean of $\lambda = 1$, a ρ of 0.6, N=10,000 plants. $\bar{K}/\bar{K}^*$ reports the average capital stock relative to the frictionless benchmark. $C(R,i)$ is the correlation between reallocation and investment, $C(R,y)$ is the correlation between reallocation and output, and $C(\tilde{A},i)$ is the correlation between mis-measured TFP and investment. The "na" entry means that the correlation is not meaningful as one of the variables is constant.

From Table 2.2, the interaction of costly reallocation and accumulation is evident in a number of forms. First, $\bar{K}$, which is the average capital for a particular treatment, depends on the nature and magnitude of the capital adjustment costs. Even in the absence of any aggregate shocks, the capital stock is around 25% lower when there are adjustment frictions compared to the frictionless case. This comparison of the average capital stocks with and without frictions stands regardless of the source of the shocks.

Second, the addition of the shocks increases the variability of capital. With shocks

to both π and λ the standard deviation of the capital stock is considerably higher than when there are only exogenous productivity shocks.

Third, capital accumulation is positively correlated with both reallocation and the Solow residual. An increase in λ, for example, leads to an increase in investment, reallocation and output. The correlation of reallocation and investment, $C(R, i)$, is informative about the effects of frictions on the incentive to accumulate capital.[21] An increase in π say, will imply that more plants are able to adjust and for this reason alone reallocation will increase. With π correlated, it is likely that more plants will be able to adjust in the future, so investment increases too. The magnitude of this correlation is smaller when only λ is random. Though the same fraction of plants adjusts each period, the gains to adjustment are larger when λ is high. This generates a positive correlation between reallocation and investment.

Finally, reallocation is pro-cyclical in the presence of shocks to either π or λ. This returns to one of the themes of the paper. If variations arise from either changes in the fraction of adjusting plants, through π, or by a change in the spread of the shocks, through λ, output responds. The key to this response is reallocation: the effects on output of getting the right amount of capital into its most productive use. This is captured through $\tilde{A}$.

2.4.3 Impulse Response Functions

Figures 2-3 and 2-4 show impulse response functions for negative shocks to π and λ. The shocks occur in period $t = 5$. The x-axes show time, while the y-axes in panels 2-4 shows the % deviation from the unconditional mean. The drop in the exogenous shock of interest is plotted in the first panel, while all other exogenous shocks are set to their unconditional means.

We first discuss Figure 2-3. The second panel shows the evolution of the two moments μ_n and ϕ_n. The negative correlation between the two series is very high, as changes in π effect the evolution of μ_n and ϕ_n in very similar ways. The third panel illustrates the co-movement between reallocation 'R' and the Solow residual. Following the shock to π less capital can be reallocated between plants, which directly affects $\tilde{A}$. The effects on output and investment are negative, as the last panel shows. Consumption, though, increases in response to the innovation to π, as discussed further below.

Figure 2-4 shows the effects of a negative shock to λ. The second panel shows the evolution of the two moments μ_n and ϕ_n. The sharp drop in ϕ_n is a direct

---

[21]For the nonstochastic and stochastic A models, this correlation is not defined as capital reallocation is constant.
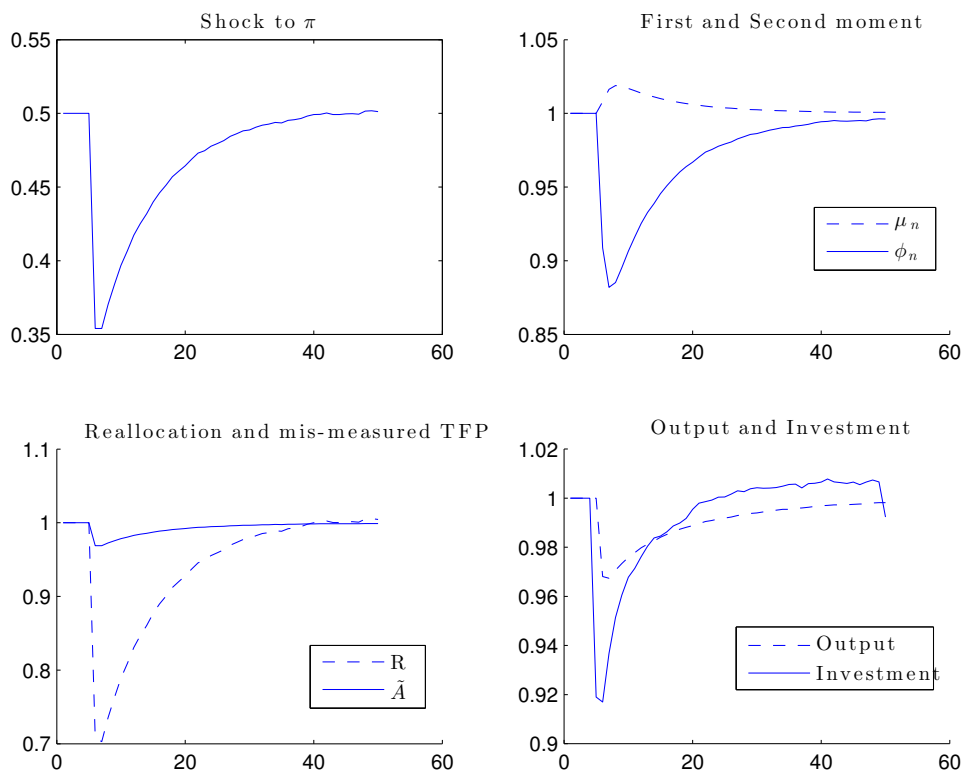
Figure 2-3: Variations in π: Impulse Response Functions. The y-axes show % deviations from unconditional means.

effect of the shock to λ, whereas the increase in μ_n reflects the effects of different reallocation choices. The panel highlights that Γ is a slow moving state variable, implying that μ_n and ϕ_n do not adjust immediately to their new values following a change in λ. Furthermore, the variations in λ have different effects on ϕ_n (direct) and μ_n (indirect), making the two moments imperfectly correlated. Variations in λ produce more cyclicality in ϕ_n than in μ_n.

Panel 3 shows the connection between mis-measured TFP and reallocation, which leads to a cyclical effect on output. In this economy with time-varying idiosyncratic uncertainty in the presence of adjustment costs there is a strong cyclical dimension of capital reallocation. Reallocation is driving time-variations in output.

Output and investment both fall in response to a negative shock to λ. The investment response is quite strong: when λ falls investment opportunities are reduced. Output falls as well due to the reduced dispersion in productivity across plants. These effects are driven by the "love of variety" aspect of the production technology. The large decrease in investment coupled with a smaller reduction in output implies that
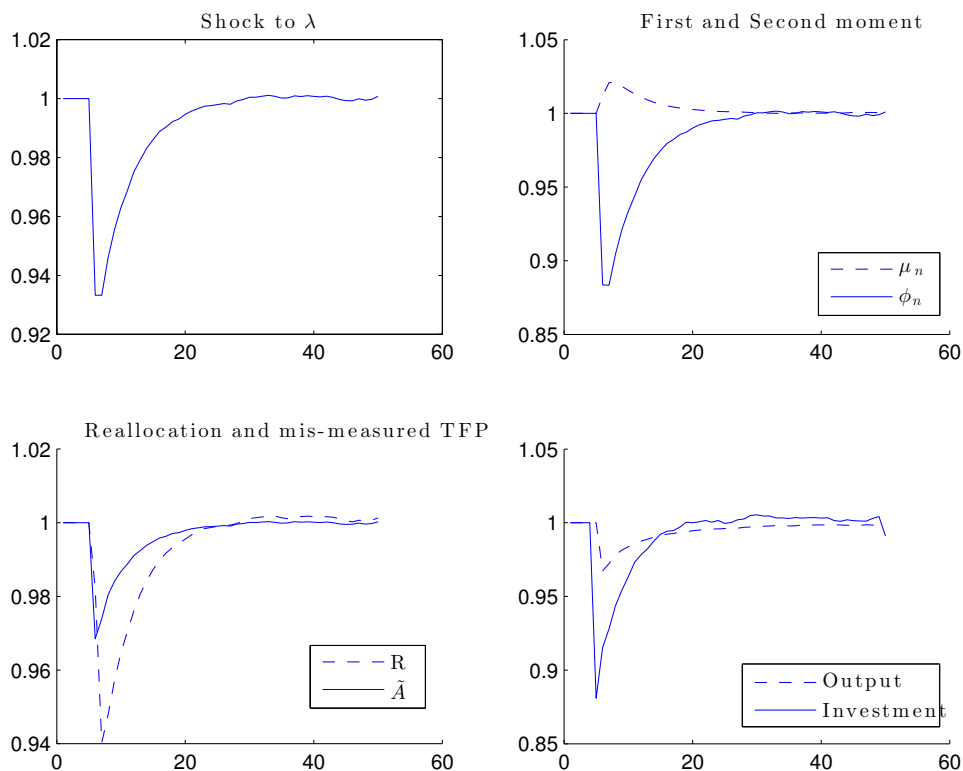
Figure 2-4: Negative shock to λ: Impulse Response Functions. The y-axes show % deviations from unconditional means.

consumption increase at the time of the shock. We return to this point later.

These responses **do not** include the fall in output associated with an *increase* in the dispersion of shocks, as emphasized in Bloom (2009), Bloom *et al.* (2012) and others. As noted above, this reflects a couple of features of our environment: (i) the timing of the shock to λ, (ii) the model of adjustment costs and (iii) the specification of the production function. Nonetheless, as indicated above, the model with both shocks, i.e. the stochastic (π, λ) case, is able to match the two key observations of pro-cyclical reallocation and a counter-cyclical dispersion in capital productivity.

2.4.4 Robustness

The previous results illustrated a couple of themes. First, variations in either π or λ are necessary to generate cyclical movements in reallocation, with resulting effects on mis-measured TFP. Second, evolution of the cross sectional distribution generated dynamics only in the stochastic π and/or λ cases. This is illustrated by the fact that higher order moments are relevant in the planner's optimization problem and the

evolution of these moments are seen in the impulse response functions.

Parameter changes	R/R^*	$E_t(\sigma_i(arpk_{it}))$	G	$\sigma(\tilde{A}/A)$	$C(R,\tilde{A})$	$C(\sigma_i(arpk_{it}),\tilde{A})$
			Frictions			
Baseline	0.492 (0.005)	1.09 (0.006)	0.108 (0.0008)	0.10 (0.003)	0.817 (0.012)	−0.194 (0.03)
$\alpha = 0.8$	0.498 (0.008)	2.33 (0.02)	0.123 (0.0006)	0.25 (0.007)	0.52 (0.04)	0.475 (0.03)
$\bar{\pi} = 0.3$	0.283 (0.005)	1.23 (0.005)	0.207 (0.001)	0.145 (0.004)	0.945 (0.004)	−0.246 (0.04)
$\bar{\pi} = 0.9$	0.899 (0.005)	0.353 (0.01)	0.014 (0.0002)	0.088 (0.003)	0.54 (0.03)	−0.247 (0.05)
$\rho_\pi = 0.5$	0.492 (0.001)	1.09 (0.004)	0.107 (0.0003)	0.07 (0.001)	0.659 (0.01)	0.486 (0.025)
$\rho_\varepsilon = 0.5$	0.429 (0.008)	1.45 (0.01)	0.248 (0.002)	0.105 (0.004)	0.965 (0.002)	−0.487 (0.03)
$\sigma_\lambda = 0.1$	0.491 (0.006)	1.14 (0.02)	0.11 (0.002)	0.49 (0.001)	0.692 (0.01)	0.762 (0.02)
$\sigma_\lambda = 0.1, \rho_\lambda = 0.5$	0.491 (0.006)	1.14 (0.01)	0.112 (0.0006)	0.341 (0.006)	0.554 (0.019)	0.609 (0.02)
timing	0.498 (0.007)	0.930 (0.01)	0.106 (0.002)	0.10 (0.006)	0.84 (0.02)	−0.78 (0.02)

Table 2.3: Capital Reallocation: Robustness

Model with stochastic π and λ. Standard deviations in parentheses.

This section studies the robustness of these findings to alternative values of key parameters. Table 2.3 reports our findings. It has the same structure as Table 2.1. The first column indicates the model. The baseline is the case with adjustment costs and stochastic (π, λ) taken from Table 2.1.

In the second row we show the effects of moving α from 0.60 to 0.80. The increase in the curvature of the revenue function leads to a larger output gap and a higher degree of misallocation. This result is largely driven by the non-adjustable plants: The column R/R^* shows that reallocation among adjustable plants is higher than in the benchmark scenario.

The baseline model assumes $\bar{\pi} = 0.5$. The third and fourth rows of Table 2.3 study the implications of lower and higher adjustment rates. Not surprisingly, the reallocation rate is increasing in π, as frictions are lower. This is consistent with Figure 2-2. The correlation of reallocation and mis-measured TFP is positive, though lower than in the baseline at $\pi = 0.90$.

The standard deviation of actual to mismeasured TFP also varies with $\bar{\pi}$. When $\bar{\pi}$ is high, the response of the planner to a variation in λ is to reallocate capital so that $\sigma(\tilde{A}/A)$ is small compared to the case of low $\bar{\pi}$. This is reflected in the mean standard deviation of the average revenue product of capital.

The table includes two rows in which the serial correlation of shocks is set to 0.5,

lower than their baseline values of $\rho_\pi = 0.9$ and $\rho_\varepsilon = 0.9$. Relative to the baseline, the reduction in the serial correlation of π leads to a reduction in the cyclicality of reallocation. With adjustment opportunities less correlated, the costs of reallocating resources that are subsequently mismatched with productivity is higher. Hence reallocation is less correlated with $\tilde{A}$. This will imply that the correlation of reallocation and investment is lower than in the baseline reflecting the costs of accumulating capital when future adjustment costs are less certain.

When ρ_ε is decreased, the planner has fewer incentives to reallocate capital among adjustable plants. Consequently, the amount of capital reallocation falls and the inefficiency of the solution becomes more pronounced. This can be seen in the larger standard deviation of the marginal products of capital and in a higher output gap.

The row labeled $\sigma_\lambda = 0.1$ increases the variability of λ relative to the baseline where $\sigma_\lambda = 0.014$. This spread is closer to that in Bloom (2009) and Gilchrist *et al.* (2014). Not surprisingly, this extra volatility in the spread of idiosyncratic shocks leads to much more volatility in $\tilde{A}$ relative to the baseline. Reallocation remains pro-cyclical though less compared to the baseline.

The next row shows how a reduction in the serial correlation of λ given the high variance of λ influences these moments. With a lower serial correlation of the shocks to λ, the correlation between reallocation and $\tilde{A}$, though still positive, is considerably lower than the baseline. With less persistent shocks, reallocation is less responsive to variations in λ and π.

The last row is a modification to the model that influences the extent of the "love of variety effect". The row labeled "timing" assumes that the planner knows of a change in the cross sectional distribution of the idiosyncratic shocks one period in advance. That is, the future value of λ is in the current state space. This is the timing used in Bloom (2009) as a way to emphasize the uncertainty effects of a change in the distribution. In our environment, the change in timing has some modest effects relative to the baseline. There is less dispersion in the average product of capital but this dispersion is more negatively correlated with $\tilde{A}$ compared to the baseline. With the alternative timing assumption the planner reallocates more capital when λ is known to remain high, and less capital when λ is known to remain low. This increases the counter-cyclicality of the dispersion and leads to an allocation of capital that is on average closer to the frictionless benchmark.

2.5 Approximation

The previous section showed that the covariance ϕ *matters* for determining the optimal capital allocation. The problem in (2.14) includes Γ, the joint distribution of (k, ε). Using the first two moments of this distribution, μ_n and ϕ_n, the evolution of Γ can be tracked perfectly. This is important for the planner, who has to forecast the expected future output from non-adjustable plants, y'_n. Variations in π and λ generate movements in Γ and hence in y_n. Capital reallocation is tightly linked to changes in the mis-measurement of TFP when stochastic shocks are present.

Movements in Γ may not be captured well by the first moment μ_n alone. In the frictionless case the two moments were perfectly correlated, but this perfect correlation is broken by the existence of time-variation in the adjustment probability π and/or λ. The impulse response functions above showed that both in the case of shocks to π or λ the two moments μ and ϕ were strongly correlated. However, different shocks imply different magnitudes of change in μ, ϕ, and output. A change in λ produces a stronger reaction in ϕ and a smaller reaction in μ compared to a shock in π. Output changes of the *same* magnitude can therefore occur at the same time as *different* changes in μ. This produces the reduced explanatory power of the first moment μ. The significance of reallocation effects is related to the forecasting power of ϕ_n.

Relative to the literature starting with Krusell and Smith (1998), this is an important finding. **In particular, this result is distinguished from preceding papers in that for our environment the approximation of the cross sectional distribution requires higher order moments.**

This section makes two points. First, it emphasizes the importance of including the higher order moments in the state vector. From this we can determine how well the evolution of Γ could be captured by different subsets of its moments under different cases of stochastic π and λ.

Second, we compare the aggregate outcome of the model against a standard stochastic growth model. This allows us to determine to what extent the reallocation effects influence cyclical properties of the model.

2.5.1 Goodness of Fit

Table 2.4 evaluates the importance of the higher order moments.[22] To understand this table, let "DGP" refer to a data set (and moments) created by solving the baseline model (with stochastic π and λ) using (μ, ϕ) in solving the planner's problem. In (2.14), the planner forecasts y'_n, the output from non-adjustable plants next period. The correctly specified regression model including both moments is given by

$$y_{n,t}^{DGP} = \beta_0 + \beta_1 \mu_{n,t} + \beta_2 \phi_{n,t} + \beta_3 s_t + \varepsilon_t, \tag{2.29}$$

where s_t includes π_t and λ_t. Estimation results in $\hat{\beta}_0 = 0$, $\hat{\beta}_1 = 1.6487 = \bar{\varepsilon}$, $\hat{\beta}_2 = 1$, and $\hat{\beta}_3 = 0$ with an $R^2 = 1$. The maximum forecast error (MCFE) is zero. As discussed in Den Haan (2010) a problem of R^2 measures to assess the approximation is that observations generated using the true law of motion are used as the explanatory variable. We construct a series $\hat{\hat{y}}_n$ which is using only the approximate law of motion. The forecast error is defined as $\hat{\hat{\varepsilon}}_{t+1} = |\hat{\hat{y}}_{n,t+1} - y_{n,t+1}|$, and the MCFE is the maximum of this series.

Below we study three cases (experiments). The first takes output of the non-adjusting plants from the DGP and regresses it on an intercept, the exogenous state, and the first moment only. Thus this exercise is about approximating the nonlinear solution with a linear representation. The regression model for the linear approximation is given by (2.29) where we force $\beta_2 = 0$. From Table 2.4, the linear representation is very accurate if only π is stochastic. When λ is random, the resulting movements in the distribution of shocks leads to much greater significance of the cross sectional distribution in forecasting (decisions do not change in this experiment).

The second case actually solves the planner's problem under the (false) assumption that the model is linear. The resulting decision rules and expectations are model consistent by construction, but not data consistent.[23] The goodness of fit measure is computed from a regression of the output of the non-adjusting plants in the DGP using the model consistent estimators from the linearized approximation. As before, the linear beliefs in the stochastic π case are approximately consistent with the outcome. Again this is not the case when λ is random. For this experiment, the linear forecast rule leads to very different allocative decisions by the planner. Consequently, the R^2 is quite low – movements in the cross sectional distribution are very important.

---

[22]Only the stochastic model with frictions is explored. The case of "stochastic A" is not of interest as the higher order moments did not matter. For these experiments, the shocks are held fixed to isolate the effects of the approximation.

[23]The R^2 from the forecast of μ in the linearized version of the model typically exceeds 0.99. In this sense, the solution is internally consistent.

In the third case, the planner uses the DGP to obtain a linear approximation of the law of motion. With this representation, the planner solves the optimization problem. In this case, the expectations about the evolution of the state vector is consistent with the data, but not with the model. Here, none of the experiments generate a good fit. The planner is simply unable to capture the nonlinear movements in the economy with a linear approximation of the law of motion.

Case	R^2	MFCE
Truth, approximated		
Stochastic π	0.9907	0.031%
	(−)	(−)
Stochastic λ	0.966	1.37%
	(−)	(−)
Stochastic π, λ	0.94	2.5%
	(−)	(−)
Linear, consistent		
Stochastic π	0.9908	0.3954%
	(−)	(−)
Stochastic λ	0.6958	0.7289%
	(−)	(−)
Stochastic π, λ	0.7032	1.707%
	(−)	(−)
Linear using DG truth		
Stochastic π	0.94	1.52%
	(−)	(−)
Stochastic λ	0.82	1.339%
	(−)	(−)
Stochastic π, λ	0.948	1.78%
	(−)	(−)

Table 2.4: Different approximation strategies

The first column shows the R^2 of a regression of output from non-adjustable plants on an intercept and the first moment, μ only. The second column reports the maximum forecast error from such a regression.

2.5.2   Comparison to the RBC Model

This section compares the aggregate properties of our model with those of the RBC model. There are two motivations for this exercise.

First, one of the key findings of Thomas (2002) and the literature that followed was the near equivalence between the **aggregate moments** of a model with lumpy investment and the aggregate implications of a real business cycle model with quadratic adjustment costs at the plant-level. This sub-section returns to that theme. Given that higher order moments matter in the planner's optimization problem, it is natural to conjecture that the non-convexities also matter for aggregate moments.

Second, a standard criticism of the RBC model is technological regress: i.e. apparent reductions in total factor productivity. As emphasized in Bloom *et al.* (2012) as well, model economies which induce variations in the Solow residual have the potential to explain technological regress and can potentially match other correlation patterns.

As we shall see, the aggregate moments of the model with stochastic (π, λ) share many of the characteristics of the RBC model. The Solow residual, driven by reallocation, has a serial correlation of nearly 0.92. Consumption, investment and output are positively correlated with the Solow residual and the model exhibits consumption smoothing. In our environment, the puzzle of "What causes a reduction in the Solow residual?" is easily resolved: measured productivity is low when reallocation is low, either due to lower adjustment rates or a contraction in the distribution of profitability shocks.

Our environment is different from Bloom *et al.* (2012) in a couple of important ways. First, our model includes shocks to both the distribution of idiosyncratic shocks and to adjustment costs. Second, as emphasized earlier, a mean preserving spread increases investment. This reflects the timing in our model as well as the structure of adjustment costs. In contrast to models with irreversibility and other forms of non-convexities, there is no option-to-wait in our model with Calvo style adjustment costs. Third, there are no adjustment costs to labor. Finally, as already emphasized, higher order moments matter for the planner and generate an underlying dynamic. In contrast, Bloom *et al.* (2012) exclude higher order moments in their approximation. As indicated earlier, there is a dynamic to these higher order moments that underlies the serial correlation in the Solow residual.

Case	$C(y,c)$	$C(y,i)$	$C(y,\tilde{A})$	$C(i,c)$	ρ_c	ρ_i	$\frac{\sigma_c}{\sigma_i}$	$\frac{\sigma_c}{\sigma_y}$
			Frictions					
stochastic A	0.91 (0.01)	0.94 (0.01)	0.93 (0.01)	0.71 (0.02)	0.95 (0.02)	0.88 (0.02)	0.53 (0.03)	0.80 (0.03)
stochastic π	0.77 (0.04)	0.90 (0.01)	0.90 (0.002)	0.42 (0.04)	0.95 (0.01)	0.91 (0.01)	0.46 (0.06)	0.80 (0.05)
stochastic λ	0.72 (0.04)	0.93 (0.01)	0.89 (0.01)	0.42 (0.03)	0.97 (0.01)	0.82 (0.01)	0.34 (0.04)	0.66 (0.05)
stochastic π, λ	0.782 (0.02)	0.898 (0.008)	0.915 (0.003)	0.427 (0.02)	0.96 (0.003)	0.86 (0.006)	0.46 (0.03)	0.80 (0.03)
RBC	0.981 (0.002)	0.913 (0.01)	0.986 (0.002)	0.818 (0.01)	0.954 (0.01)	0.890 (0.013)	0.633 (0.04)	0.919 (0.02)

Table 2.5: Endogenous Capital Accumulation - Macroeconomic Moments

Results from 1000 simulations are reported with standard deviations in parentheses below. Here $C(x,y)$ are correlations, ρ_x is an autocorrelation and σ_x is a standard deviation. The variables are: output (y), consumption (c), investment (i) and the Solow residual (mis-measured TFP) ($\tilde{A}$).

Table 2.5 presents standard aggregate moments for a number of cases. These are the traditional macroeconomic moments: the correlations of output (y), consumption (c), investment (i) and TFP($\tilde{A}$). Here the TFP measure is the one constructed from the data as if plants were homogeneous, i.e. mis-measured TFP. The serial correlations of consumption and output as well as relative standard deviations are reported, too.

The rows are the various cases explored before, using the baseline parameters. The last row, "RBC" is the standard stochastic growth model with productivity shocks and without adjustment costs.[24] Here the productivity shocks come from fitting an AR(1) process to the mis-measured TFP series, $\tilde{A}$, generated by the stochastic (π, λ) case. We obtain an AR(1) parameter $\rho_{\tilde{A}} = 0.9183$ and standard deviation of the residual $\sigma_{\tilde{A}} = 0.0132$. This process is fed into the model without adjustment frictions to produce the "RBC" moments.

All of the models match the standard business cycle properties of positively correlated movements of consumption and investment with output. All of these variables are positively correlated with (mis-measured) TFP. So, in the case of shocks to λ, the Solow residual, investment and output all increase when there is a mean preserving spread in the distribution of shocks. The models exhibit consumption smoothing. The aggregate moments are all positively serially correlated.

Further, the models with stochastic π and/or λ create considerably lower comovement between consumption and investment compared to the RBC case. As in models with intermediation shocks, such as **?**, and discussed further for the case of stochastic λ in Bachmann and Bayer (2013), when returns to investment are large, say due to a high value of λ, consumption is reduced to finance capital accumulation.

The key to this lower correlation is the immediate inverse relationship between consumption and investment when there is a shock to λ. After the impact, consumption and investment move together in the transition dynamics. So, overall there is a positive correlation but one that is reduced due to the negative comovement in response to the innovation. This can be see in the impulse response functions for our model, Figures 2-3 and 2-4.

This effect appears in other models of shocks to the variance of productivity shocks. Looking at the impulse response functions in Bloom *et al.* (2012), Figures 7 and 8, and Bachmann and Bayer (2013), Figure 3, this negative comovement at impact is apparent. Further, though this negative comovement is not evident in unconditional data moments, it does appear in impulse response functions. In Figure

---

[24]The RBC moments are produced using our model without adjustment frictions. The only stochastic shocks occur to A.

3 of Bachmann and Bayer (2013), the immediate response in the data to an increase in idiosyncratic risk is for output and investment to increase and consumption to fall.[25] Output and investment fall subsequently.

2.6   Conclusion

The goal of this paper was to understand the productivity gains from capital reallocation in the presence of frictions. To study this we have looked at the optimization problem of a planner facing frictions in capital accumulation and shocks to productivity, adjustment costs and the distribution of plant specific shocks.

The heterogeneity in plant-level productivity provides the basis for reallocation. The frictions in adjustment prevent the full realization of these gains. The model can generate cyclical movements in reallocation and in the cross sectional distribution of the average productivity of capital.

There are three key findings in this paper. The first is the cyclical behavior of reallocation and the distribution of capital productivity. When shocks to either adjustment frictions or the distribution of plant-level shocks are present, then reallocation is pro-cyclical. In fact, even if there are no direct shocks to TFP, the reallocation process creates fluctuations in output and investment. These effects are not present when the only shock is to TFP. Further the standard deviation of the cross sectional distribution of average capital productivity is counter-cyclical, as in Eisfeldt and Rampini (2006) and Kehrig (2011).

Second, in some, though not all environments, the plant-level covariance of capital and profitability shocks matters for characterizing the planner's solution. This is important for a few reasons. It is indicative of state dependent gains to reallocation and our economy is an example of one where moments other than means are needed in the planner's problem.

Third, the model with shocks to adjustment costs and the cross sectional distribution of productivity shocks can reproduce many features of the aggregate economy. A researcher would interpret the data as generated by a model with TFP shocks even though it is actually constant. That is, the researcher could certainly misinterpret the variations in the Solow residual driven by the reallocation of capital as variations in TFP.

---

[25]These results are for German data. Bloom *et al.* (2012) do not report impulse response functions to uncertainty shocks in US data.

Chapter 3

The Employment and Productivity Effects of Short-Time Work in Germany

3.1 Motivation

The question whether 'short-time work' (STW) can save jobs during a recession has found renewed interest following the recent global economic downturn. STW describes a policy response whose alleged effect is to reduce the negative impact of demand shocks on the labor market. More specifically, the policy generates incentives for firms to reduce employment through a reduction in the number of hours worked per employee, instead of through adjustments in the number of employees. A number of OECD countries have implemented short-time work schemes in order to prevent massive layoffs during times of economic distress. In the present study we will focus on the largest such scheme, the German *'Kurzarbeit'*. While the economic press has often attributed the stability of the German labor market during the global economic recession to the extensive use of *Kurzarbeit* in German firms (see Brenke *et al.* (2011), Rinne and Zimmermann (2011), and references therein), a number of recent studies takes a more critical stand (see e.g. Burda and Hunt (2011), Balleer *et al.* (2013), Möller (2012)).

Figure 3-1 uses macroeconomic indicators to illustrate that the global economic recession only had a minor impact on the German labor market. Similar to other OECD countries, Germany was hit by the financial and sovereign debt crisis. In response, output in Germany dropped sharply, resulting in a drop in GDP of 5.13 per cent in 2009. This was mainly caused by weak domestic demand and a reduction

in exports to the US and the rest of the Euro Zone. At the same time hours worked declined by 3.21 per cent in 2009.[1] Even though German GDP growth was the second lowest in the OECD, the labor market in Germany seemed almost unaffected by this shock: the unemployment rate continued to fall over the crisis period. This experience stands in sharp contrast to the United States - where the economic recovery in the post-crisis period has been described as 'jobless' - and other OECD countries.
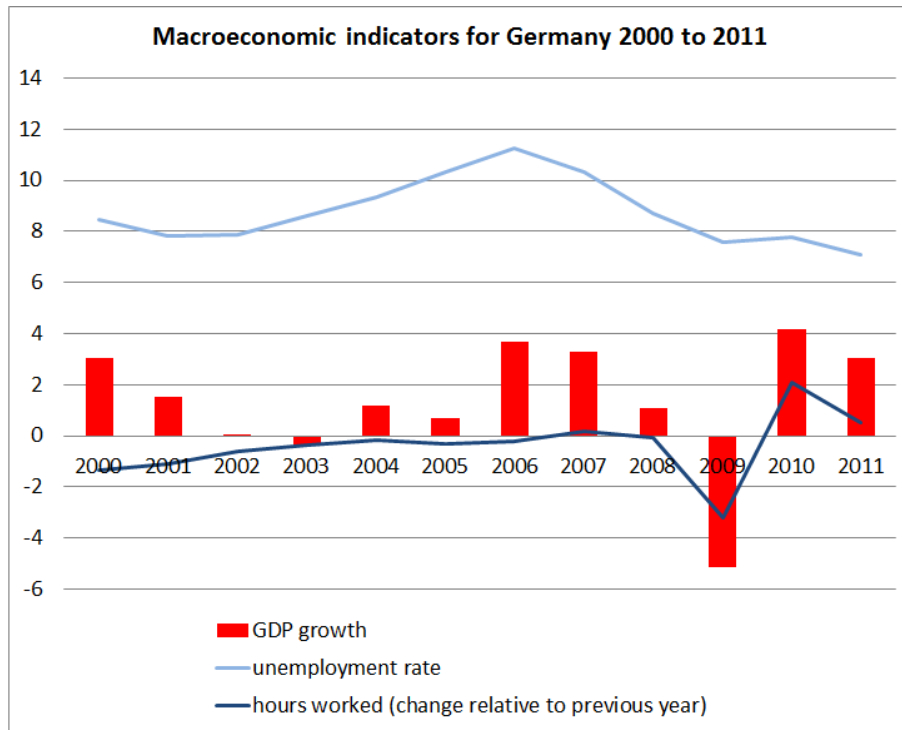
Figure 3-1: Macroeconomic indicators for Germany between 2000 and 2011. Source: data provided by OECD statistics (accessed March 2013).

Existing studies have in common that they focus on the question whether or not *Kurzarbeit* can save jobs. An arguably equally important question regards the longer-term implications of this policy for productivity, output, and employment. Besides the additional financial burden faced by the social security system, the costs of STW also include the effects on output and employment which stem from the government's intervention into the allocation of factors. In market economies the efficiency of the allocation of factors across production sites has been shown to play an important role for aggregate productivity (see Hsieh and Klenow (2009), Bartelsman *et al.* (2013), Cooper and Schott (2014)). While STW may decrease the level of unemployment

---

[1]The figure also suggests that adjustments in the year 2008 when the crisis first hit Germany were mainly done through other instruments such as flexible time accounts.

in the short run and mitigate the consequences of economic downturns, the negative long-term effects on productivity and employment could potentially outweight the short-term gains. By preventing the reallocation of factors across production sites, the policy may unintentionally induce a potentially long-lasting misallocation of labor across firms. Because labor is partly prevented from flowing towards the most productive firms, STW can generate adverse long term effects on GDP through this 'reallocation channel'.

In this paper we devise a dynamic, structural model of labor demand. The model is calibrated to the German economy, which allows us to evaluate the costs and benefits of *Kurzarbeit* quantitatively. Representative firm models ignore the large amount of heterogeneity which is present between firms. Such a framework is therefore unable to capture the effects of a policy on labor reallocation between firms. Our micro data collected at the firm level allows for a clear identification of the policy impact on firm behavior. We exploit the panel structure of the data (AFiD-Panel Industriebetriebe, Germany) and make use of information on medium and large firms in the manufacturing sector to better understand the short-term and long-term consequences of labor market interventions during the recent recession. Although the AFiD-Panel does not contain information on whether or not a firm applied for short-time work, the following section produces ample evidence that the effect of the policy is evident in our data.

3.2 *Kurzarbeit* in Germany

Kurzarbeit has a long tradition in Germany.[2] Firms' eligibility is typically restricted to specific economic conditions and imposes strict rules regarding the workforce affected by a reduction in working time.[3] During the recent global recession the German government dramatically loosened the *Kurzarbeit* eligibility criteria for firms and significantly expanded the scope of the policy (Burda and Hunt (2011)). Figure 3-2 illustrates the increase of short-time work among German firms between 1990 and 2010.[4] At its height in mid-2009 the Kurzarbeit program included around 60,000 establishments and approximately 1.5 million workers. The figure also shows that

---

[2]cite missing

[3]The German Social Security Code (SGB) III defines the legal framework for the use of Kurzarbeit. Altogether, there are three different forms of short-time work: (1) Due to economic distress (§170), (2) seasonal fluctuations (§175) and (3) transfer payments mainly during the German reunification (§216b).

[4]The spike in STW to the beginning of the 1990s can mainly be attributed to labor market adjustments during the transition of Eastern Germany from a planned economy to a market based economy; see case (3) in previous footnote.

absent recessions, the role of STW is negligible.[5] For this reason we choose to focus exclusively on the policy changes that were implemented during 2009 and 2010.

Figure 3-2: The use of short time work in Germany. Source: Bundesanstalt fuer Arbeit, Germany (accessed March 2013).

The structure of the German economy plays an important role in evaluating the impact of the global recession on the labor market. The use of STW was not evenly distributed across industries and regions. It was mainly firms in the manufacturing sector with strong dependence on exports which applied for STW in response to economic distress. More specifically, many of these firms were part of the automobile industry and included suppliers. These firms are heavily concentrated in some regions in southern and western Germany. It is important to keep in mind that a sequence of labor market reforms during the early 2000s, the so-called *Hartz reforms* offered firms more flexibility in adjusting the workforce. Even though these reforms undoubtedly changed the underlying structure of the German labor market, transition dynamics were mainly concluded by the time the global recession affected Germany.[6]

We now describe the details of the short-time work policy which was put into place in 2009 and 2010. The policy change consisted of an amendment to an ex-

---

[5]Seasonal fluctuations are heavily driven by the use of STW in the construction and agricultural sector to mitigate the impact of periods of bad weather; see case (2) in previous footnote.

[6]see Krause and Uhlig (2012) who evaluate the impact of those reforms on employment.

isting law which governs the use of short-time work by firms.[7] *Kurzarbeit* during the recession consisted of a state-subsidy of 60% (67% for workers with children) for the net earnings difference due to a working hours reduction. Hours worked were paid as usual. The employers' contribution to social security was initially paid in full by the firm. In a subsequent modification, the employers' contribution was made proportional to the hours worked (Crimmann *et al.* (2010), Rinne and Zimmermann (2011)). *Kurzarbeit* therefore provided incentives for firms wanting to cut back their labor demand to reduce the number of hours worked per employee (intensive margin) instead of reducing the number of employees (extensive margin). According to Burda and Hunt (2011) *Kurzarbeit* constituted the most common source of changes in hours per worker between 2008 and 2009.
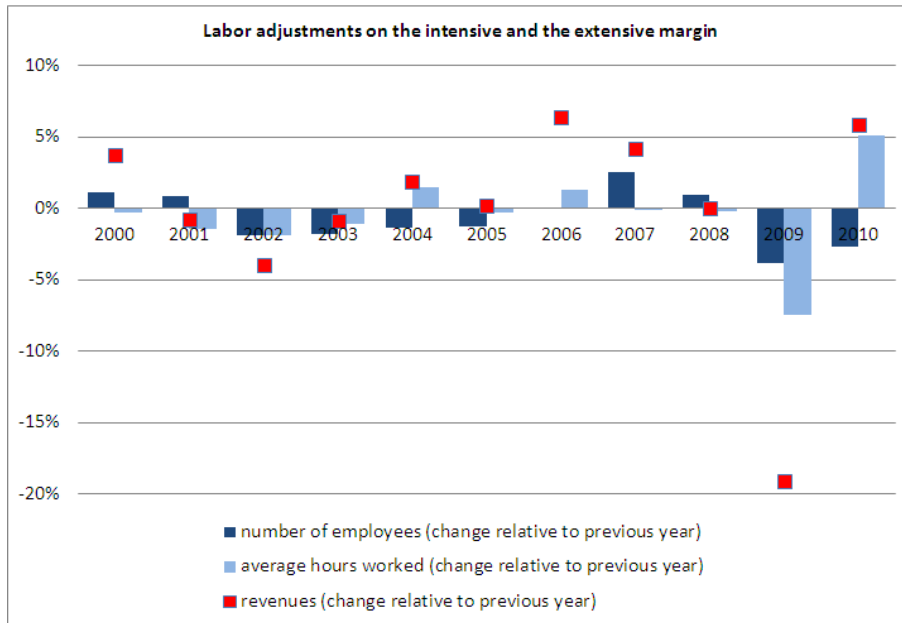
Figure 3-3: Labor adjustments on the intensive and the extensive margin. Source: author's calculations from the micro data from the AFiD-Panel Industriebetriebe, Germany.

Figure 3-3 uses our German micro-data to show the impact of the recession on revenues and labor adjustments. The red squares indicate year-to-year revenue growth. The blue bars show the year-to-year growth in the number of employees (dark blue) and the average number of hours worked (light blue). All numbers represent yearly averages of firm-level variables. Three things stand out. First and most importantly, the years 2009 (2010) saw an unprecedented decline (increase) in the average hours

---

[7]See (German Social Security Code (SGB) III, §169 ff) Further, we refer to Möller (2012) and Burda and Hunt (2011) for additional details.

worked, while the 2009 adjustment in the total number of employees appears small in proportion to the fall in revenues. Second, given the overall pattern, changes in revenues are more pronounced than changes in employment or average hours worked. Third, changes in employment appear to react to changes in revenues with a lag, while changes in average hours respond contemporaneously.

The fact that *Kurzarbeit* created an unprecedented flexibility in average hours is central to this study. We therefore present further evidence in Figure 3-4, where we show the distribution of changes in average hours for the years 2000-2008 (dark blue) vs. the recession years 2009 and 2010 (respectively blue and light blue).[8] The data is computed using year-to-year changes on the establishment-level. For example, the bin '-20%' shows the percentage of firms that changed the average number of hours worked by their employees by between -10 and -20%, while the bin '-30%' shows the percentage of firms that adjusted average hours worked by more than -30%. Three things stand out: First, in general the hours change distribution is characterized by many small adjustments of less than 2.5%, but virtually no inaction. Second, there exists a significant number of firms each year in which average hours worked are adjusted by more than 20% in absolute value, implying that large changes in average hours are prevalent. The third observation regards the changes in 2009-10 with respect to the previous years. We see a clear shifting of firms towards the tails of the distribution, where the large negative (positive) adjustments stem from 2009 (2010). In fact, the distribution in 2010 is almost a mirror image of 2009, as adjustments in average hours are reverted after the end of the STW policy. The fraction of firms that reduced average hours by over 20% more than tripled starting in 2009. In 2010, on the other hand we observe larger-than-average positive adjustments in average hours worked. We conclude that data we are using to compute the moments the model is ultimately trying to match thus clearly shows the impacts of the STW policy. The reduction in average hours can be taken to stem from the STW policy, while as the policy faded out in 2010, hours adjustments were reversed.

3.2.1 Relation to Existing Studies

In this paper we estimate a structural model of the employment behavior of heterogeneous firms. In contrast to much of the existing literature on STW we tackle the question of the effectiveness of STW from a labor demand perspective. Because the decision about STW lies with the firm we find this the right perspective to take. In our model, firms differ in their idiosyncratic productivity and consequently have

---

[8]Excluding the recession years 2000 and 2001 leaves the results virtually unchanged. In the Appendix we show the distribution by year. There is very little variation between 2000 and 2008.
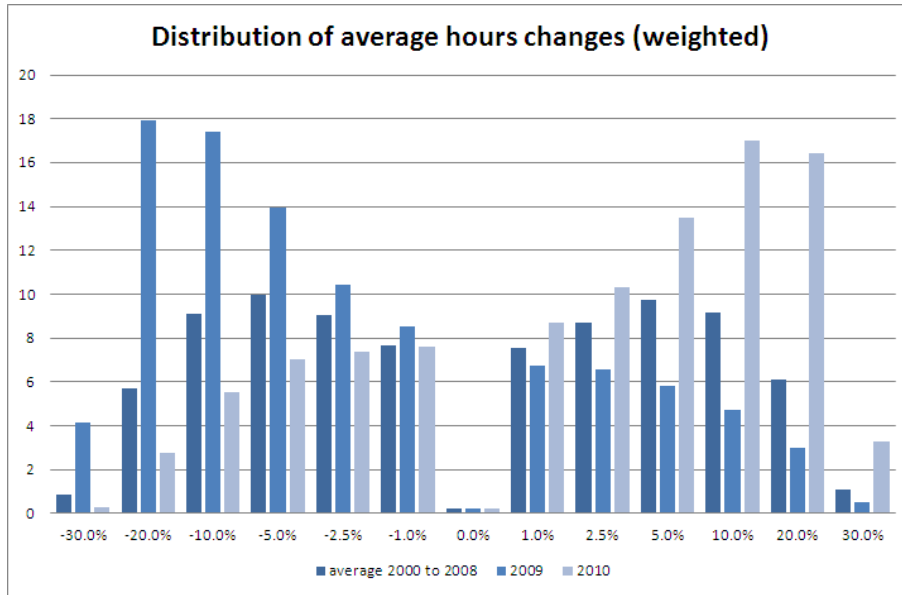
Figure 3-4: Distribution of average hours changes. Source: author's calculations from the micro data from the AFiD-Panel Industriebetriebe, Germany.

differing labor demand. Firms simultaneously choose the number of employees and the number of hours worked per employee. We then calibrate the model to match key moments of the German microdata. The calibrated model is used to understand and evaluate the effects of the policy reform introduced by the German government in 2009.

We share the skepticism of other authors (add refs) about whether *Kurzarbeit* was effectively able to save jobs. This skepticism is founded on the results of a structural model which we use to evaluate the *Kurzarbeit* policy. Using a structural model is preferable for at least three reasons. First, we derive the macroeconomic implications of STW directly from the microeconomic data. This is important because there exist important non-linearities in the labor adjustment decisions of heterogeneous plants which can have macroeconomic implications. For example, firms that increase their labor demand at the end of a recession might not be the same firms that decreased their labor demand at the recession's onset. Our approach puts the firm-level decision rules for hours and employment at the heart of the analysis, allowing us to study the effect of STW on firms with different characteristics such as size or profitability. Second, the calibrated economy enables to us evaluate counter-factual scenarios. We can ask what the effect on the German economy would have in the absence of *Kurzarbeit*. Third, one of the most fundamental questions about the effectiveness of STW cannot be answered adequately in a reduced-form setting. Given recent findings about the

macroeconomic importance of the allocation of factors (Hsieh and Klenow (2009), Bartelsman *et al.* (2013)), one of the main concerns about STW is that it constitutes a goverment intervention into the allocation mechanism. In order for such an intervention to be welfare-enhancing one would have to identify a market failure this policy can help to overcome. As we show further below, we do not find convincing evidence for this. Going beyond this, we are able to quantify the short- and long-run effects on output and employment which stem from this 'reallocation channel'.

Recent papers which study STW in Germany include Burda and Hunt (2011), Balleer *et al.* (2013), Krause and Uhlig (2012), and Cahuc and Carcillo (2011). Burda and Hunt (2011) provide an excellent overview of the institutional framework in Germany and describe the 2009/10 policy in great detail. Using a reduced form model Balleer *et al.* (2013) do not find a significant employment effect of STW, but they do find a positive, albeit small effect, on output. They devise a search and matching model which can rationalize those effects. In their model heterogeneous workers can be put on STW. If an STW 'policy shock' is persistent, it generates positive output and employment effects because the firm can reduce the working times of unprofitable workers and the value of a match is increased. Krause and Uhlig (2012) also use a search-and-matching model with heterogeneous workers and skill depreciation to analyze the German *Hartz-reforms*, a major overhaul of the unemployment benefits system effectuated during the early 2000s. Regarding STW they conclude that the transition towards a post-reform steady state was mainly achieved prior to 2008 and that STW played an important role in keeping umemployment low. Our model differs significantly from Balleer *et al.* (2013) and Krause and Uhlig (2012). It features heterogeneous firms and homogeneous workers. Firms are subject to idiosyncratic productivity shocks which evolve persistently over time. In Balleer *et al.* (2013) STW is the *only* possibility for firms to adjust hours per worker. In this sense, STW is simply a reduction in the marginal cost of adjusting labor demand along the intensive margin to a value less than ∞. This unrealistic description of labor adjustment costs also makes the resulting positive effects of STW little surprising. Krause and Uhlig (2012) do not model STW directly. In contrast to this, we use a simulated method of moments (SMM) technique to estimate adjustment costs that are able to replicate the distributions of hours and employment changes in German firms.

Other related papers include Burdett and Wright (1989) an Braun et al (2013). Burdett and Wright (1989) study how unemployment insurance and STW distort labor inputs. Their main results are than UI causes inefficient layoffs, while STW induces inefficient hours. Braun et al (2013) build on this framework to study the

welfare effects of short-time work. They find that the effectiveness of STW depends on the degree to which firms are insured against idiosyncratic profitability shocks. Similarly to us they find that STW can be poorly targeted and benefit the 'wrong' firms.

3.3 Data

The German 'AFiD-Panel Industriebetriebe' collects annual data ranging from 1995 to 2010 on the universe of German manufacturing plants with more than 20 employees.[9] The underlying data is collected on a monthly basis and the aggregated by year. The data covers approximately 50,000 plants per year. In the case of Germany where employment in the manufacturing sector is heavily concentrated in medium sized and family owned firms (German Mittelstand), the focus on the impact of labor market policies on firm behavior during the global economic recession requires the inclusion of medium-sized manufacturing firms.[10] Using official micro data prepared and made available trough the German Federal Statistical Office in cooperation with the statistical offices of the German Laender we can drastically reduce the bias due to sampling problems because of of over- and under-sampling firms with certain characteristics.

Descriptive statistics on the most important variables can be summarized as follows. This paper makes use of a balanced panel which includes a sub sample of 19,522 firms observed for every year between 1995 and 2010. To abstract from extreme observations this sub sample of firms also excludes observations which report revenues, hours or employment levels above the 95th percentile or below the 5th percentile of the initial distribution. On average (over all years and all firms) each firm in the final sample employs 184 workers. Average hours worked per employee are 135.47 hours per month (including overtime). Mean revenues reported per firm are 3.65 million Euros (deflated to 2005). The mean wage bill is reported to be 0.40 million Euro (deflated to 2005). These numbers suggest that workers compensation accounts for on average 10.1 per cent of revenues.

---

[9]The AFiD-Panel, Industriebetriebe includes smaller plants if they belong to a firm with at least 20 employees. For the years 2007 to 2010 the cutoff to be included into the production survey is 50 employees.

[10]Alternative data provided by commercial providers such as Amadeus often only collects annual information on publicly listed companies with a focus on financial variables.

3.4 Model

The economy consists of firms, workers, and a government. Firms choose their optimal level of employment and hours each period and produce a homogenous output good. Similar to Rogerson (1989) workers are part of a large, risk-sharing household which owns the firms. Employed workers receive a wage income to compensate them for the total number of hours worked. This remuneration can include overtime compensation or short-time work subsidies. If workers are unemployed they receive unemployment benefits b. Unemployment benefits and short-time subsidies are paid by the government which finances itself through taxation.[11] A period in the model refers to one month.

The institutional arrangement on the German labor market is characterized by collective wage agreements. Those agreements are binding for trade union members and employers who are part of the employer's association which led the negotiations with the trade union. The standard workweek $(\bar{h})$ as well as the hourly base wage (ω_0) are subject to these negotiations. We do not explicitly model the contract which determines $\bar{h}$. We assume that firms are small enough such that they take the level of $\bar{h}$ as given (see Hunt 1999). Positive and negative deviations from $\bar{h}$ are costly for the firm: Positive deviations require the payment of an overtime premium of typically around 25%. While small weekly fluctuations in hours worked can be compensated through working time accounts (Burda and Hunt (2011)), reductions of average hours worked to a level below $\bar{h}$ are typically only possible through traditional short-time work (absent the modifications introduced during the recent recession). We therefore assume that absent the policy choices of $h < \bar{h}$ are not feasible (see Kydland & Prescott (1991) or Hunt (1999)). The social security contributions are a function of h, the actual number of hours worked.

3.4.1 Firms

We now present the dynamic problem of the firm.[12] The timing of events for a firm in period t is as follows. In the beginning of the period all plants observe the current profitability draw A.[13] Profitability can entail a common shock as well as a plant-specific shock. Given A and the firm's beginning-of-period level of employment e_{-1} the firm chooses the level of employment e and average hours worked $h \geq \bar{h}$ to

---

[11]Social security contributions in Germany are paid in part by the firms and in part by the workers.

[12]The dynamic optimization problem builds from the specification in Cooper *et al.* (2007) and Cooper *et al.* (2011).

[13]Time subscripts are dropped for notational convenience.

maximize its value $V(A, e_{-1})$. If $e \neq e_{-1)}$ the firm pays adjustment costs to adjust its stock of workers. Adjustments in hours occur without adjustment costs.

The value of the firm in state (A, e_{-1}) is given by

$$V(A, e_{-1}) = \max_{h,e} R(A, e, h) - \omega(e, h) - C(e_{-1}, e) + \beta E_{A'|A} V(A', e) \qquad (3.1)$$

for all (A, e_{-1}). Here $R(A, e, h)$ is the revenue flow of a plant with e workers, each working h hours in profitability state A. $\omega(e, h)$ refers to the compensation of workers and $C(e_{-1}, e)$ describes the adjustment costs to labor. The revenue function depends on hours per worker (h) and the number of workers (e).[14] Factors of production other than labor, such as capital and energy, are freely adjustable within a period. With constant returns to scale and constant elastic demand, the revenue function takes the form $R(A, e, h) = A(eh)^\alpha$. The coefficient α reflects the curvature of the production function along with the elasticity of demand. Note that this particular production function implies that the elasticities of R with respect to e and h are identical. We discuss the implications of this assumption further below (see also Burdett and Wright (1989)). Current profits are defined as revenues minus compensation paid to workers and minus costs of adjusting the workforce.

Compensation without STW

A key part of our analysis concerns the compensation function $\omega(e, h)$. Due to transfers from the government the compensation paid by firms and the compensation received by workers differ. Prior to the introduction of STW the compensation paid by the firm to workers is given by:

$$\omega(e, h) = e\omega_0[\bar{h} + \lambda(h - \bar{h})](1 + \delta). \qquad (3.2)$$

Here ω_0 plays the role of a base wage. The social security contribution is parameterized by δ. λ determines the overtime premium if actual hours worked h exceed the standard workweek $\bar{h}$. Absent the short-time work policy the average number of hours worked cannot lie below the standard work week, so $h \geq \bar{h}$. The length of the average work week $\bar{h}$ is taken as given by the firm.

---

[14]This function comes from the product of a production function and the demand function facing the plant.

Compensation with STW

After the STW reform, the compensation function is more complicated. In addition to overtime also $h < \bar{h}$ is allowed. For $0 \leq h \leq \bar{h}$, the compensation paid by the firm is

$$\omega(e, h) = e\hat{\omega}_0[h(1 + \delta) + \mu(\bar{h} - h)\delta] \tag{3.3}$$

The firm pays social security contributions for $\bar{h}$ but receives a subsidy of $(1 - \mu)$ from the government for the difference in hours between h and $\bar{h}$. The introduction of μ was an essential part of the policy reform during the crisis. We parameterize the social security contributions for the $\bar{h} - h$ hours paid by the firm by μ. An experiment below will be how μ influences the hours and employment choices of the firm.

For $h > \bar{h}$, i.e. when the firms chooses overtime after the reform, the only thing which changes in the compensation function with respect to (3.2) is that the new base wage is given by $\hat{\omega}_0$.

$$\omega(e, h) = e\hat{\omega}_0[\bar{h} + \lambda(h - \bar{h})](1 + \delta). \tag{3.4}$$

Labor Adjustment Costs

The cost of adjusting the stock of workers is given by $C(e_{-1}, e)$. This function captures the various inputs into the process of hiring a worker, including search, recruitment and training costs. It may contain both convex and non-convex forms of adjustment costs. A general cost of adjustment function would be

$$C(e_{-1}, e) = F^+ + \gamma^+(e - e_{-1}) + \frac{\nu}{2}\left(\frac{e - e_{-1}}{e_{-1}}\right)^2 e_{-1} \tag{3.5}$$

if there is job creation $e > e_{-1}$. Similarly

$$C(e_{-1}, e) = F^- + \gamma^-(e_{-1} - e) + \frac{\nu}{2}\left(\frac{e - e_{-1}}{e_{-1}}\right)^2 e_{-1} \tag{3.6}$$

if there is job destruction $e < e_{-1}$. If $e = e_{-1}$, i.e. when there are no *net* changes in employment, then $C(e_{-1}, e) \equiv 0$.

There are three types of adjustment costs, with differences allowed for the job creation and job destruction margins. The first is the traditional quadratic adjustment cost, parameterized by ν. A fixed cost of adjustment is parameterized by F^+ for job creation and F^- for job destruction. Finally, there are linear adjustment costs. The linear firing cost γ^- captures severance payments to workers. One of the key features plant level data is inaction in employment adjustment. The fixed cost and linear costs

are both capable of creating inaction.

A crucial question in our setup regards the substitutability between hours and workers, the intensive and extensive margins of labor adjustment. Given the pre- and post-policy compensation functions (3.2) and (3.3) above we can now study the changes in firms' incentives to reduce employment along either one of those two margins *before* and *after* the policy.

Clearly, absent labor adjustment costs, with $C(e, e_{-1}) = 0$, a firm would never choose $h > \bar{h}$ due to the overtime premium. After the introduction of the STW policy all downwards adjustments of labor would go through reductions in e since $\mu > 0$. The calibration of the labor adjustment costs are therefore important in order to generate movements in average hours.

3.4.2 Households

The household consists of a continuum of workers who can either be employed or unemployed. We normalize the size of the continuum to one. Preferences of an individual agent are given by $u(c - g(h)) - \zeta I(h > 0)$. Here c denotes consumption, with $u(\cdot)$ being a strictly increasing and strictly concave function. The function $g(\cdot)$ determines the disuility incurred from hours worked. We assume $g(\cdot)$ to be strictly increasing and strictly convex. The last term in the individual agent's utility function describes a fixed cost of providing positive hours in the labor market. It is parameterized by ζ.

For this specification of compensation, the worker receives the amount given in (3.2) without the social contribution: i.e. the worker receives $\omega_0[\bar{h} + \lambda(h - \bar{h})]$. As we shall develop further, this compensation can be related to the worker's utility function.

(After the reform) From the worker's perspective, for $\underline{h} \le h \le \bar{h}$, the worker receives compensation directly from the firm of $\hat{\omega}_0 h$ and also obtains compensation of $\phi\hat{\omega}_0(\bar{h} - h)$ from the government. Here ϕ is a replacement rate for compensation lost due to hours reductions. For $h > \bar{h}$, the worker receives $\frac{\omega(e, h)}{(1 + \delta)}$ (for e = 0) from (3.4).

Note that the notation allows the base wage itself to change with the new policy towards hours variation. We will look at those changes both empirically (using household data) and in theory, using a condition for labor market equilibrium.

3.4.3 Government

Prior to the introduction of the policy the government collects tax revenues in order to finance unemployment benefits b for the $(1-N)$ unemployed workers. Taxes are raised through the social security contributions of firms. We assume that the government's budget must be balanced at all times.

$$(1-N) \cdot b = \delta \int e\omega_0[h + \lambda(h - \bar{h})]d\Sigma, \tag{3.7}$$

where Σ is the joint distribution over firms' profitability, hours, and employment. After the policy the tax revenues also have to be used to finance the short-time work scheme. The second term on the left-hand side is the integral over all firms who use short-time work, i.e. where average hours have been reduced below $\bar{h}$. For each employee in such a firm the government pays a fraction of the loss in income, $\phi(\bar{h} - h)$ plus a subsidy for the firms' social security contributions. Instead of $\delta\bar{h}$ the firms only pay a fraction μ for every hour between h and $\bar{h}$, the rest is paid by the government, $\delta(1 - \mu)(\bar{h} - h)$.

$$(1-N)\cdot b + \int \mathbb{1}_{\bar{h}>h}(e\cdot(\bar{h}-h)(\phi+\delta(1-\mu)))d\Sigma = \delta \int e\hat{\omega}_0[\max(h,\bar{h})+\lambda(h-\bar{h})]d\Sigma \tag{3.8}$$

3.5 Parameterization and Estimation

The first set of parameters is not calibrated, but taken directly from the data. Those include the overtime premium λ, the social security contribution rate of the employer σ, the share of σ paid by the firm in the case of STW μ, the replacement rate for the foregone compensation due to hours reduction ϕ, and the average workweek $\bar{h}$. The parameter values are summarized in Table 3.1. We provide additional information on the parameter choices in the Appendix.

The base wage pre-policy will be set to match average employment size.

3.5.1 Revenue Function & Productivity Estimation

The revenue function, $R(A, e, h) = A(eh)^\alpha$, is estimated first. We make use of plant-level data to estimate the curvature of this function, α. The shock to the revenue function is the residual from the estimation. From that residual we can obtain a representation of the stochastic process for the revenue shock. This process is then used in the solution of the dynamic optimization problem in equation (3.1).

Table 3.1: Parameters for the German economy.

parameter	characteristics	value
λ	overtime premium	0.25
δ	employer social security contribution rate	0.21
μ	share of social contribution paid by firm in case of STW	0.66
ϕ	replacement rate for compensation lost due to hours reduction	0.60 to 0.67
$\bar{h}$	contracted hours	38 hours

We estimate the production function using both OLS and the approach presented in Wooldridge (2009), which instruments for labor using its first lag. In all specifications we control for time and industry dummies. In our micro data the total number of hours worked in a firm is our independent variable of interest, as it corresponds to $e \cdot h$ in our model. The estimated coefficients of α are larger with the OLS approach, but the estimated TFP levels do not differ much in magnitude. We find $\alpha = 0.56$. The autocorrelation of idiosyncratic TFP is estimated to be 0.92.

3.5.2 Adjustment Costs Estimation

The remaining parameters are estimated via a simulated method of moments approach (SMM). This approach finds the vector of structural parameters, denoted Θ, to minimize the weighted difference between simulated and actual data moments:

$$\pounds(\Theta) \equiv (M^d - M^s(\Theta))W(M^d - M^s(\Theta))'. \tag{3.9}$$

The parameters estimated by SMM are $\Theta \equiv (\omega_0, \omega_1, \hat{\omega}_0, \hat{\omega}_1, \zeta, \bar{h}, \nu, F^+, F^-, \gamma^+, \gamma^-, \beta)$.

The estimation method starts by solving the dynamic programming problem in (3.1) for a given value of Θ. The decision rules are calculated as part of this solution. Shocks to profitability are then drawn in a manner consistent with the process estimated in the first stage. Given these shocks and the decision rules at the plant level, a simulated panel data set is created and the simulated moments are calculated. The weighting matrix, W, is obtained by inverting an estimate of the variance-covariance matrix obtained from bootstrapping the data.

The key to the procedure is the choice of moments. These moments must be statistically informative about the key parameters we wish to estimate. The adjustment cost parameters can be identified from variations in job creation and destruction rates as well as through the comovement of employment with revenues. The moments

corresponding to the labor adjustment costs come from the employment change distribution shown in Figure 3-5
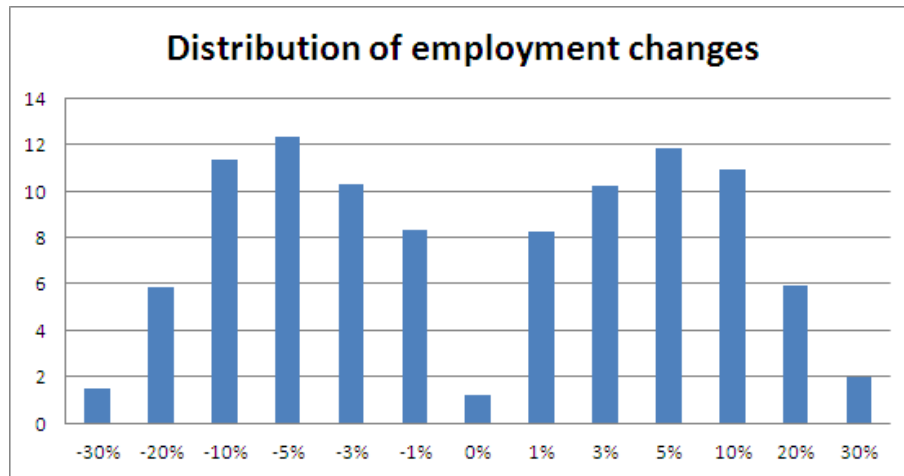
Figure 3-5: Distribution of employment changes. Weighted by total employment. Source: own calculations from the micro data from the AFiD-Panel Industriebetriebe, Germany.

3.6 Robustness

In this section we will check the results with respect to a number of assumptions, e.g.:

- Hours reduction not possible below $\bar{h}$ prior to the policy

- How would results change if technology was biased against short-time work? I.e. include fixed cost of employment for the firm.

- Whether or not employment takes time-to-build

- The (relative) revenue elasticities of hours and employment

3.7 Policy Simulations

Once we have estimated the structural parameters, $\hat{\Theta}$, we can study the introduction of the labor market policy. To a large degree this is just the substitution of (3.3) and (3.4) for (3.2).

3.8   Conclusion

This paper studies the employment and productivity implications of short-time work in Germany. This policy was intended to provide incentives for firms during the recent recession to adjust labor input through hours per worker (intensive margin) instead of firing workers (extensive margin). Using confidential German firm micro data we will estimate a model of costly labor adjustment and use it to simulate the effects of the policy.

Appendix to Chapter 1

A.0 Data

The main dataset I use for this paper is the Business Dynamics Statistics (BDS) dataset published by the Census. This annual dataset is derived from the Longitudinal Business Database (LBD) and covers both firm size, firm age, as well as firm- and establishment level data. A unique feature of the BDS is its longitudinal source data that permit tracking establishments and firms over time. A strength of data is its robustness to ownership changes because the age of a firm is determined by the age of its oldest establishment.

I complement the analysis by considering alternative data sources obtained from the Bureau of Labor Statistics (BLS). Virtually all of my qualitative results can also be obtained with the 'Business Employment Dynamics' (BED)series by the BLS. The BED is derived from a quarterly census of all establishments under state unemployment insurance programs, representing about 98 percent of employment on nonfarm payrolls. The data frequency is quarterly. It includes data on firm age and firm size. A caveat is the limited comparability between the age and size series as the age data is based upon establishment-level data, while the size class tabulations use firm-level data instead. For this reason I present most of the trends using the BDS data.

Another source released by the BLS is the Current Employment Statistics (CES) program. This is a monthly survey of about 145,000 firms and government agencies, representing roughly 557,000 establishments. Despite its high frequency the survey-nature of the CES and its limited representation of the US economy make this data source less useful for the purpose of the present paper.

The series for house prices come from the Federal Housing Finance Agency (FHFA), which provides national and state-level house price indeces from 1991 onwards. The unemployment rate was obtained from the BLS. The quarterly series of state-level personal income was obtained from the Bureau of Economic Analysis (BEA).

Table 2: Summary Statistics for Variables used in Regression

	N	mean	min	p25	p50	p75	max
empc	3,276	-0.000	-4.4e+04	-1.2e+03	-88.366	975.876	8.1e+04
hpi	3,738	-0.000	-34.978	-2.381	-0.210	1.509	48.872
pi	3,738	0.000	-8.4e+04	-1.4e+03	-98.793	1144.323	8.0e+04
ue	3,738	0.011	-2.460	-0.408	-0.025	0.400	4.135
N	3738						
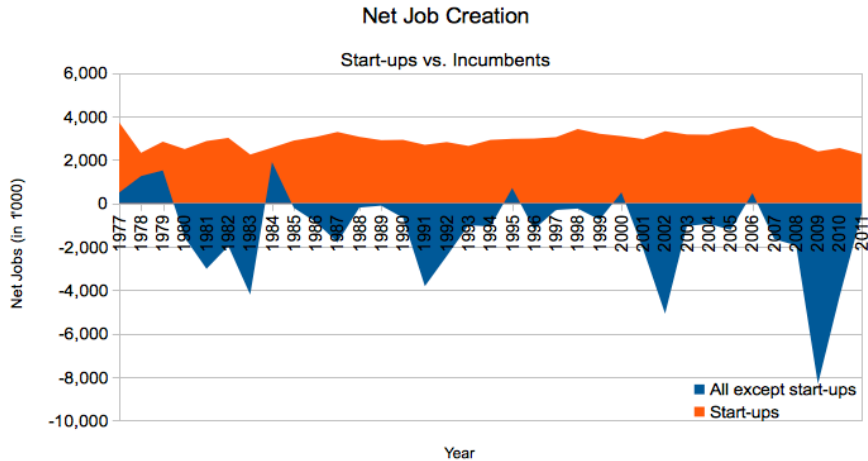
Figure -6: Net job creation by start-ups vs. incumbents. Source: Census, BDS

A.1 Additional Material to support the Stylized Facts

This Appendix includes figures referenced to in the main text. It is intended to give further evidence for the stylized facts presented in the main text.

The Importance of Start-Ups Figure -6 plots an updated version of a graph used in Coles and Kelishomi (2011). It shows net job creation by start-ups and incumbent firms. Net job creation by incumbent firms is typically negative. This is related to the life cycle of a typical firm. Figure -7 shows gross job creation and destruction between 1977-2011.

The left panel shows the raw data, while the right panel shows HP-filtered data. In both cases we see that while job destruction spiked during the 2007-09 recession, the spike was *less* pronounced than during the 2001 recession. Furthermore the graphs show that compared to all previous recessions, there has been an unusually sharp decline in job creation rates. Figure -8 compares the cyclicality of employment in entrants and incumbents. The standard deviation of the plotted series are 0.10 for
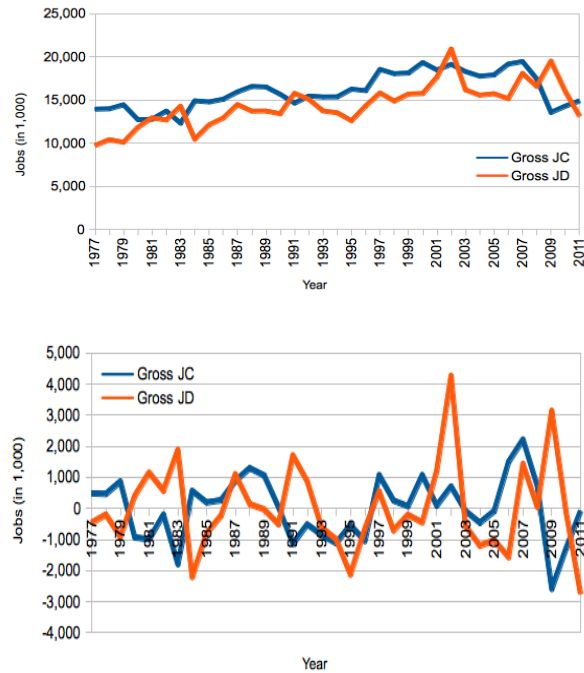
Figure -7: Gross job creation and destruction 1977-2011. The HP-filtered cyclical component is depicted in the right panel. Source: Census, BDS

start-ups and 0.02 for incumbent firms.

Comparing different Recession episodes   Other studies, e.g. Sanchez and Liborio (2012) have used alternative data sources such as the Business Employment Dynamics (BED) from the BLS to show the decline in startup activity.

Indicators for Credit Supply, Interest Rates, and Home Equity Extraction
On the one hand, the lending environment has become tighter during the last recession. Many studies point to the idea that the decrease in credit supply is the result of illiquid funding markets faced by commercial banks and a reassessment of bank lending practices and business strategies (see literature review). Banks whose balance sheets have been more severely affected by increased loan defaults may either have insufficient capital to make additional loans, or may choose to conserve capital instead of making loans to entrepreneurs ?. Other than during previous post-WWII recessions the percentage of institutions reporting negative quarterly net income increased to over 30% in 2009.[15] According to the Federal Reserve's 'Senior Loan Officer Opin-

---

[15]Based on FDIC data. The average number of institutions with negative quarterly income between 1990 and 2006 was 8.39%. During 2001 and 2002 the highest percentage was 14.87%. See also Figure -12 where the increase in interest rates was much less pronounced during 2001 than 2008.
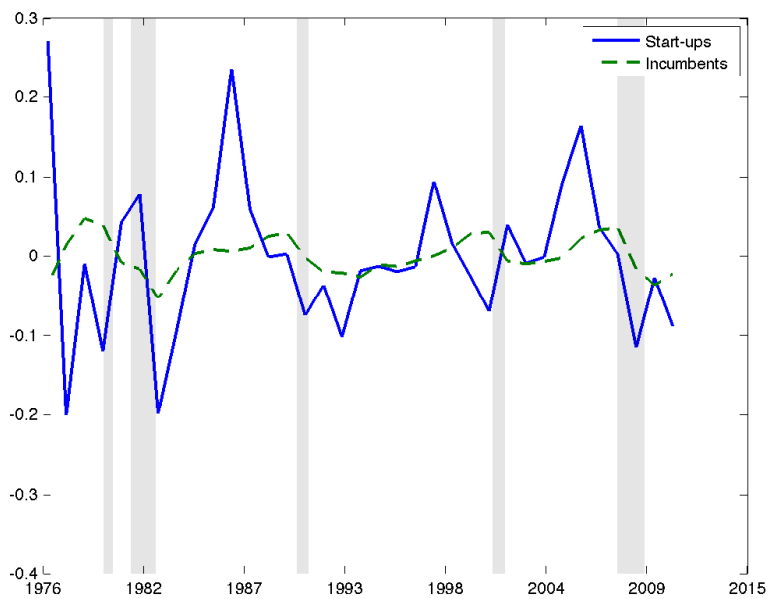
Figure -8: Cyclicality of job creation. Start-ups vs. employment in incumbent firms (dashed line). I HP-filter the annual data with $\lambda = 100$. Plotted is the cyclical component over the trend component. Recession dates are indicated as the shaded areas. Source: Census, BDS
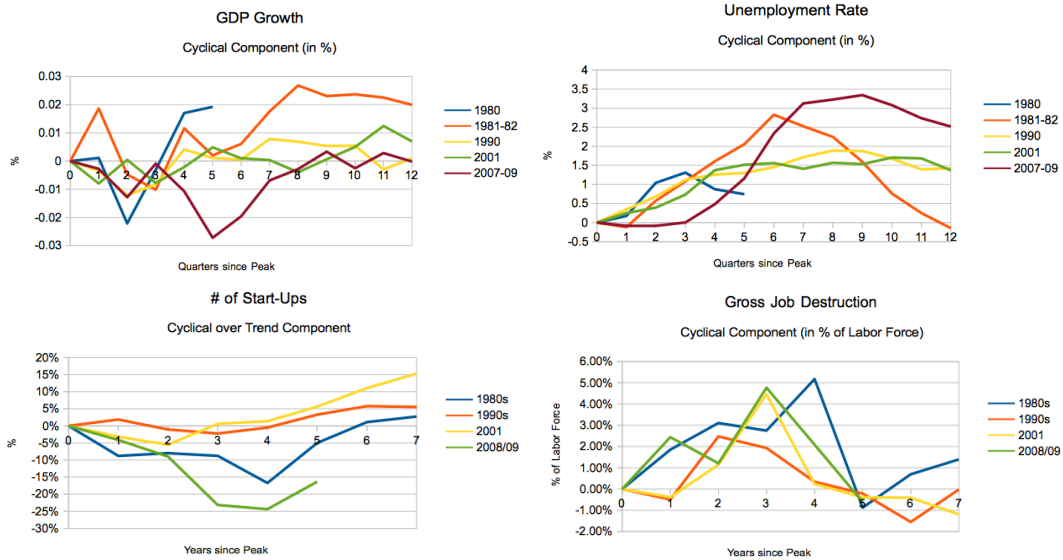
Figure -9: Comparing Recession Episodes: GDP, Unemployment, number of start-ups , and job destruction. GDP and unemployment are quarterly series, start-ups and job destruction are annual. All series are HP-filtered with $\lambda = 100$ for annual and $\lambda = 1600$ for quarterly data. The x-axis shows periods since the respective pre-recession peak, i.e. last period before the official NBER recession date. Unemployment data comes from the BLS and matches the period of Census data publication. For the annual series I treat the 1980 and 1981/2 recession as a single episode.
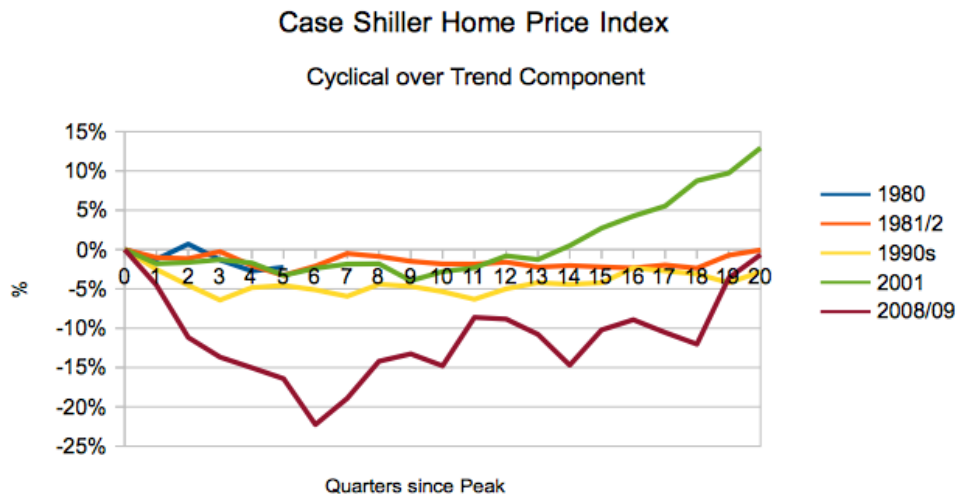
Figure -10: Cash Shiller Home Price Index. HP-filter $\lambda = 1600$. The x-axis shows quarters since the respective pre-recession quarter (based on NBER classification). Inflation-adjusted, not seasonally adjusted. Source: Standard&Poor's. Own computations

99
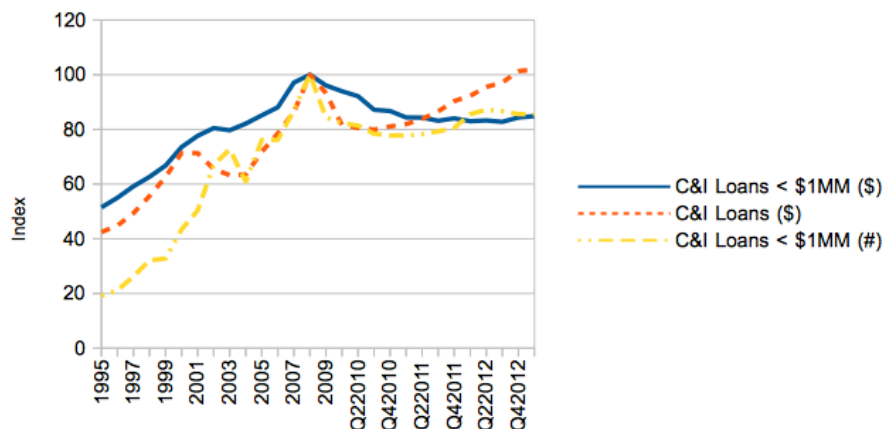
Figure -11: Domestic Commercial and Industrial Loans to U.S. Addressees. The blue solid line is C&I loans under $1Million (in Millions of $). The organe dotted line is all C&I loans (in Millions of $). The yellow dash-dotted line is the number C&I loans under $1Million. Source: FDIC

ion Survey on Bank Lending Practices' by the end of 2008, 69.2% of banks reported that they had tightened credit standards, especially for firms with annual sales less than $50 million (80%). Results are shown in Figure -13.

Decomposing Changes in the Unemployment Rate   Following the methodology developed in Elsby *et al.* (2009) I use data from the Bureau of Labor Statistics (BLS) and decompose changes in the unemployment rate into changes due to variations in the inflow rate and changes due to variations in the outflow rate of unemployment. The data shows that the increase in the unemployment rate was mainly due to decreases in the *outflow* from unemployment, i.e. lower hiring. Using the formula for the evolution of the steady state unemployment level we can write $u_t = \frac{s_t}{s_t + f_t}$, where s_t and f_t describe the unemployment inflow and outflow hazard rates. Log differentiation of this expression then yields $d \log u_t \approx (1 - u_t)[d \log s_t - d \log f_t]$. See Elsby *et al.* (2009) for further details. An increased entry hazard would speak for higher rates of job destruction through layoffs and quits, while a decreased exit probability is related to stalling job creation and/or decreased efficiency of the matching process. While early papers such as Darby *et al.* (1986) suggested that increases in unemployment during recessions are mainly due to increasing number of inflows, the more recent literature has taken the opposite stand. Hall (2005a), Hall (2005b), and Shimer (2012) have made the claim that modern recessions do not share this feature and are characterized by acylcical inflow rates. I use the Q2 2013 Current Population Survey (CPS) by the Bureau of Labor Statistics (BLS). The left panel of Figure -15
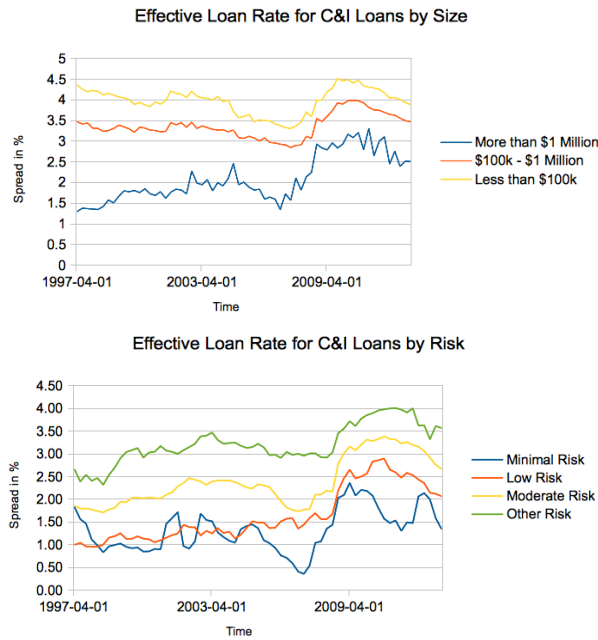
Figure -12: Commercial and Industrial Loan Rates Spreads over intended federal funds rate, by loan size and risk (E2). Source: Federal Reserve
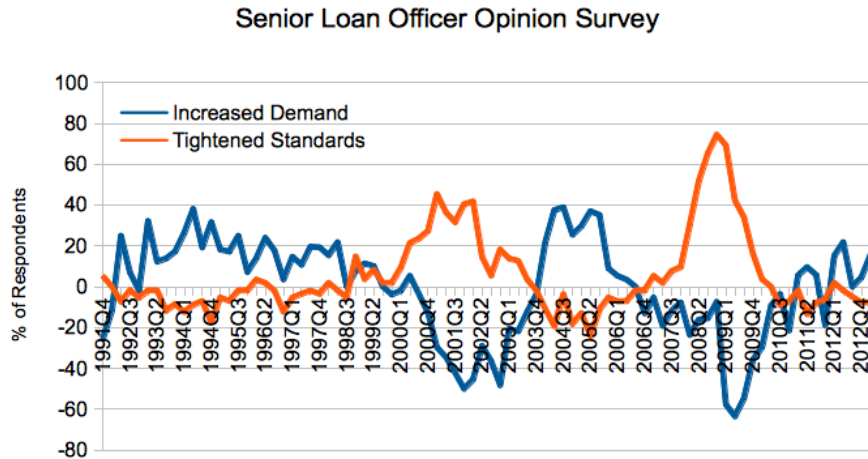
Figure -13: Results from 'Senior Loan Officer Opinion Survey on Bank Lending Practices'. The blue line plots the net percentage of banks reporting tightening standards for C&I loans to firms with annual sales of less than $50 million. The organge line plots the net percentage of banks reporting stronger demand for C&I loans from those same firms. Source: Federal Reserve.

101
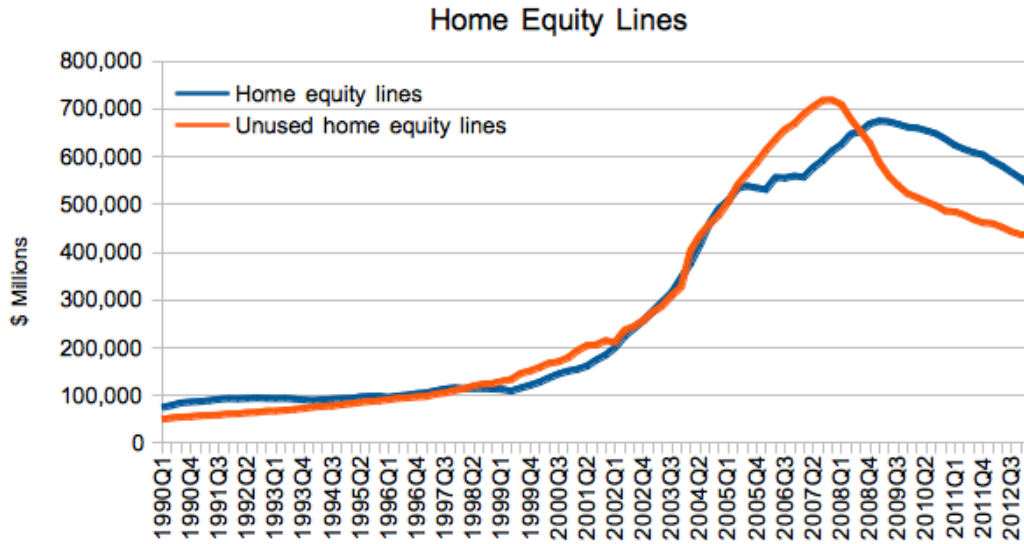
Home Equity Lines

Figure -14: Used and Unused Home Equity Lines. Source: FDIC
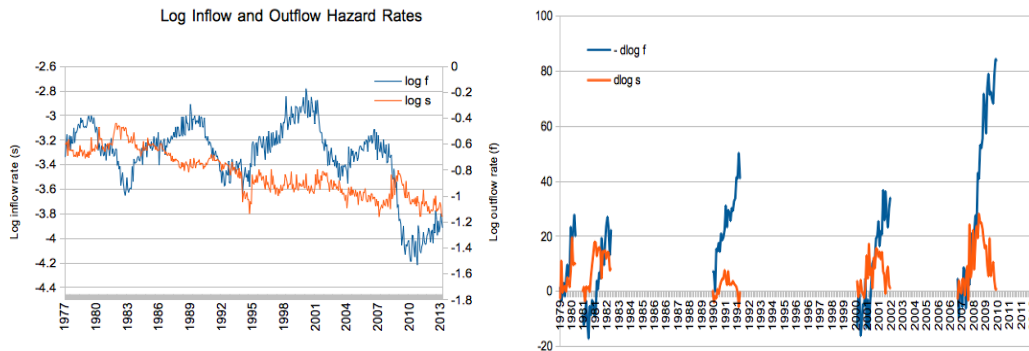
Figure -15: Left: Log inflow hazard rate s (orange, left scale) and log outflow hazard rate f (blue, right scale). Right: Changes in log inflow rates s and log outflow rates f by recession. Changes are shown with respect to start-of-recession values. I follow Elsby *et al.* (2009) in choosing the starting dates as the respective minimum and maximum unemployment rates preceding and following the NBER recession dates. Source: BLS, CPS, own computations.

	1-19	20-99	100-499	500+
Firms	88.4%	9.66%	1.54%	0.35%
Establishments	71.32%	10.48%	4.66%	13.54%
Employment	20.14%	18.02%	13.93%	47.90%
Number of Start-ups	98.1%	1.75%	0.14%	0.01%
Startup Employment	69.36%	20.90%	8.26%	1.47%

Table 3: Size- and Employment Distributions. Source: Census/BDS. Employment is calculated using the DHS-denominator.

	Age 0	Age 1	Age 2	Age 3	Age 4	Age 5
Firms	11.09%	8.54%	7.22%	6.29%	5.55%	4.97%
Employment	3.16%	3.15%	2.87%	2.68%	2.53%	2.42%

	Age 6-10	Age 11-15	Age 16-20	Age 21-25	Age 26+	
Firms	18.67%	12.91%	9.42%	7.18%	8.16%	
Employment	10.36%	8.89%	8.14%	7.94%	47.87%	

Table 4: Firm- and Employment distributions by age. Source: Census, BDS.

plots the log variation in the inflow (s) and outflow rates (f). While the inflow rate increased at the onset of the recent recession, its cyclicality is dwarfed by that of the decrease in the outflow rate. The right panel of the same figure plots the changes in the decomposition of the unemployment rate and leads to the same conclusion: the decreases in the unemployment exit hazard has been the major contributing factor to the continuingly high unemployment rate we observe today. This result strengthens the conclusion summarized in *Stylized Fact 1*.

The Size-Age Distribution of Firms and Establishments

A.2 Model Properties

This Appendix includes proofs, derivations, and details about the properties and fit of the model.

Policy Function for Employment   The policy function for employment is shown in Figure -16.

Proof of Proposition 2

Proof. Part 1: An agent will always choose to collateralize the highest possible fraction of the loan. Denote this fraction as μ. The entrepreneur chooses this fraction to
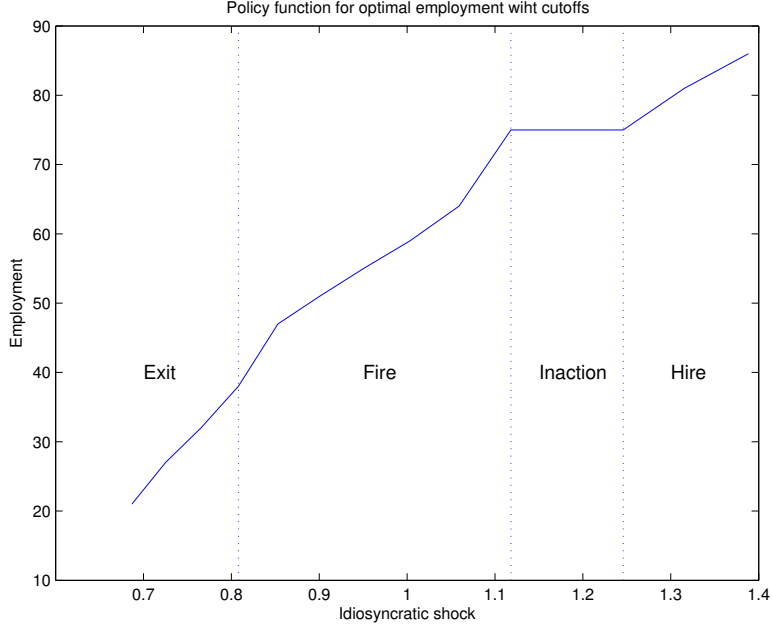
Figure -16: Employment Policy Function for given values of a, θ, and e_{-1}.

minimize his interest payments $\tilde{R} = 1 \cdot \mu + \hat{R} \cdot (1 - \mu)$. The minimization problem reads $\min_{0 \leq \mu \leq 1} c_e - c_e \cdot 1 \cdot \mu - c_e \cdot \hat{R} \cdot (1 - \mu)$ subject to the collateral constraint $\mu \cdot c_e \leq q^h$ and $0 \leq \mu \leq 1$. The collateral constraint says that the value of the secured fraction of the loan, $\mu \cdot c_e$, cannot exceed the value of the collateral. The resulting corner solution is $\mu = \min\{\frac{q^h}{c_e}, 1\}$. If $\frac{q^h}{c_e} \geq 1$ then $\mu = 1$ and $\tilde{R} = 1$. If $\frac{q^h}{c_e} < 1$ we have $\mu = \frac{q^h}{c_e}$ and $\tilde{R} = \frac{q^h}{c_e} + \hat{R}(1 - \frac{q^h}{c_e}) = \frac{q^h}{c_e} + \hat{R}(\frac{c_e - q^h}{c_e})$.

Part 2: In a given period the bank lends an uncollateralized amount x to a mass M_t of ex-ante identical entering entrepreneurs. A fraction $F(\bar{\varepsilon}^x) = \int_0^{\bar{\varepsilon}^x} d\nu$ of the M_t entrants will receive an initial productivity draw below the exit threshold $\bar{\varepsilon}^x$ and hence default on the loan. The remaining fraction $1 - F(\bar{\varepsilon}^x) = \int_{\bar{\varepsilon}^x}^\infty d\nu$ will receive a draw above $\bar{\varepsilon}^x$ and repay the initial loan times the non-default interest rate $\hat{R}$. The zero-profit condition of the bank implies $M_t x - \hat{R} \cdot M_t x \int_{\bar{\varepsilon}^x}^\infty d\nu = 0$ or $\hat{R} = \left(\int_{\bar{\varepsilon}^x}^\infty d\nu\right)^{-1}$. Clearly, $\frac{\partial \int_{\bar{\varepsilon}^x}^\infty d\nu}{\partial \bar{\varepsilon}^x} < 0$, so it follows that $\frac{\partial \hat{R}}{\partial \bar{\varepsilon}^x} > 0$. $\qquad\square$

Proof of Corollary

Proof. We have

$$
\begin{cases}
\frac{\partial \tilde{R}}{\partial q^h} = \frac{1}{c_e} - \frac{1}{c_e} \cdot \hat{R} \leq 0 & \text{if } q^h < c_e \\
\frac{\partial \tilde{R}}{\partial q^h} = 0 & \text{if } q^h \geq c_e
\end{cases}.
$$

104
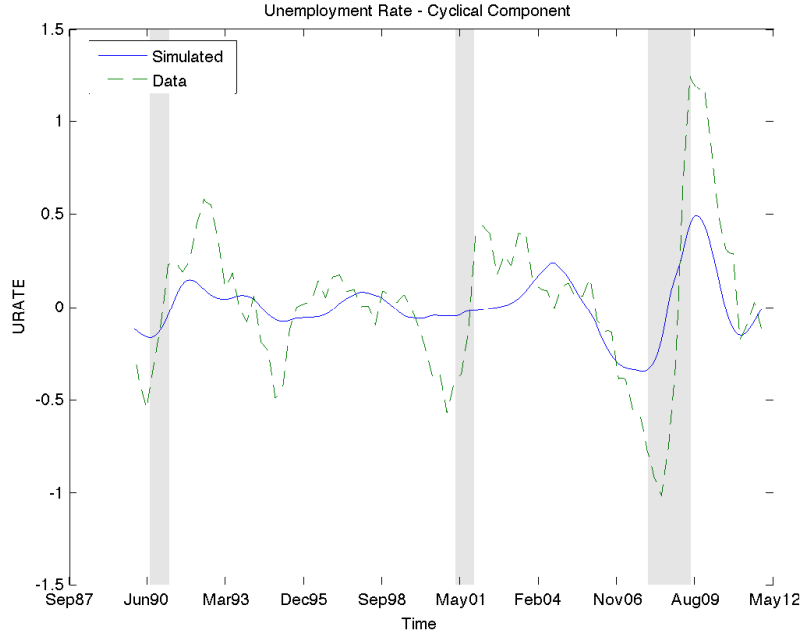
Figure -17: Cyclical component of the unemployment rate. Data vs. simulation using estimated processes only for q^h between 1990 and 2011. Shaded areas correspond to NBER recession dates.

Furthermore

$$
\begin{cases}
\frac{\partial \tilde{R}}{\partial \hat{R}} = 0 & \text{if } q^h \geq c_e \text{ or } \bar{\varepsilon}^x = 0 \\
\frac{\partial \tilde{R}}{\partial \hat{R}} = -\left(\frac{c_e - q^h}{c_e}\right) < 0 & \text{else}
\end{cases}.
$$

From Proposition 1.4.2 $\frac{\partial \bar{\varepsilon}^x}{\partial a} < 0$ and $\frac{\partial \bar{\varepsilon}^x}{\partial \theta} > 0$. Since $\frac{\hat{R}}{\bar{\varepsilon}^x} > 0$ we obtain the results stated in the corollary. $\qquad\square$

A.2.1 Additional Model Simulations

A.2.2 Accuracy of the Solution

Figure ?? shows an accuracy plot which compares the actual values of θ from a simulation of the model with the model forecast based on $\mathbb{H}$. Importantly, the latter series does not include the actual θ as an input, but makes forecasts based on the last-period prediction. This means that errors are allowed to accumulate over time. Figure ?? shows that the two lines are almost indistinguishable. The average percentage difference is 0.002%. The maximum percentage difference is 0.005%. This suggests that $\mathbb{H}$ is succesful in tracking the simulated dynamics of θ.
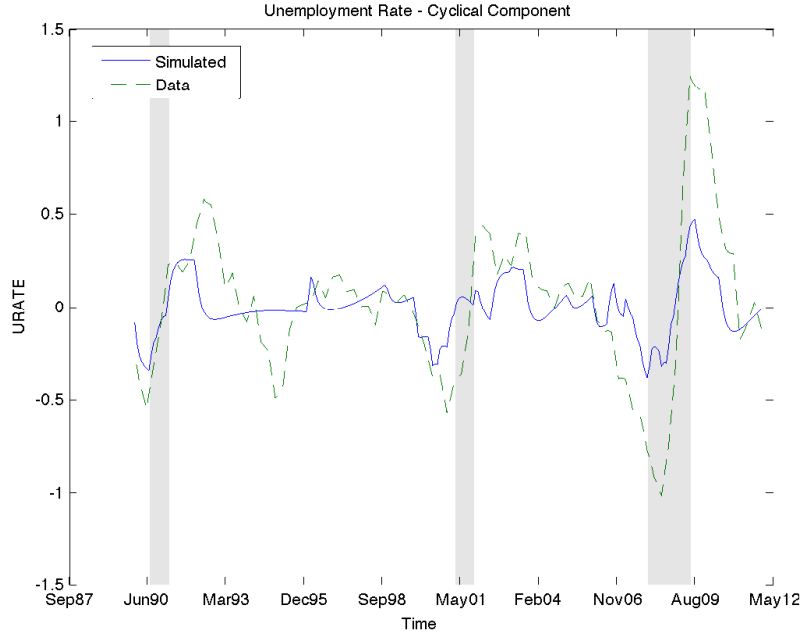
Figure -18: Cyclical component of the unemployment rate. Data vs. simulation using estimated processes only for a between 1990 and 2011. Shaded areas correspond to NBER recession dates.

A.3 Computational Strategy

For the solution of the model I use a non-stochastic grid method. While this method requires finer grids for firm-specific labor and productivity it has the great advantage of eliminating sampling error. As **?** shows, sampling error can lead to severe distortions in the model's results. This is all the more important in my setup, as the mass of entering firms can be small relative to the mass of incumbents. Therefore sampling uncertainty may bias the results even though the overall number of firms is large.

Before beginning the simulation I create fine grids for n and ϵ. Denote the number of grid points by $\#_n$ and $\#_\epsilon$, respectively. I specify an initial distribution over the points $[n_i, \epsilon_j]$, where $i \in [1, 2, \ldots, \#_n]$ and $j \in [1, 2, \ldots, \#_\epsilon]$. This determines the mass of firms with employment n_i and productivity ϵ_j. The simulation then follows this iterative process:

1. At each grid point incumbent firms decide whether to continue operation or exit. The decision is based on equation (1.9) above.

2. New firms enter based on equation (1.11)

3. The aggregate productivity state realizes according to its law of motion specified in (1.13).

106

4. The idiosyncratic productivity state realizes. This implies distributing the mass at each point $[n_i, \epsilon_j]$ to a new point $[n_i, \epsilon_k]$, where $k \in [1, 2, \ldots, \#_\epsilon]$, according to the law of motion specified in (1.14).

5. Apply the employment policy function. This involves distributing the mass at each point $[n_i, \epsilon_k]$ to $[n'_i, \epsilon_k]$, where n'_i is given by the firm's policy rule resulting from the maximization of (1.4).

6. Go back to step 1.

The simulation algorithm takes as given the policy functions for employment (hires, fires, and inaction) ϕ_e, and exit, as well as the laws of motion of all exogenous states, π_ϵ and π_A. To find a solution for a given aggregate state A, it iterates on a distribution over emplyment and idiosyncratic productivity, $\lambda(e, \epsilon)$ and finds its fixed point, where

$$\lambda_{t+1}(\bar{e}_l, \bar{\epsilon}_m) = \sum_{i=1}^{M} \sum_{j=1}^{N} \Pr(\phi_e(\bar{e}_i, \bar{s}_j) = \bar{e}_l | e_t = \bar{e}_i, \epsilon_t = \bar{\epsilon}_j) \pi_{jm} \lambda_t(\bar{e}_i, \bar{\epsilon}_j).$$

The distribution λ has dimensionality $(\#_e \cdot \#_\epsilon \times 1)$, where $\#_e$ and $\#_\epsilon$ respectively refer to the number of grid points for employment and the idiosyncratic shock. In practise the law of motion is set up by combining the policy functions and the law of motion for the idiosyncratic state into a large transition matrix Γ, which has dimensionality $(\#_e \cdot \#_\epsilon \times \#_e \cdot \#_\epsilon)$. This transition matrix Γ may vary for incumbents and entering firms, since entrants are allowed to have a different initial transition matrix for the idiosyncratic shock. The non-zeros in the row associated with $\bar{e}_i, \bar{\epsilon}_j$ are then defined as

$$\Gamma((i-1) \cdot \#_\epsilon + j, (\phi_e(i,j) - 1) \cdot \#\epsilon + 1 : \phi_e(i,j) \cdot \#_\epsilon) = \pi_\epsilon(i, :) \cdot (1 - \phi_x(i,j)).$$

Then we can rewrite the law of motion for λ as

$$\tilde{\lambda}_1 = \tilde{\lambda}'_0 \Gamma,$$

and the solution can be found by iteration or solving $\tilde{\lambda} = \tilde{\lambda}' \Gamma$, where $\tilde{\lambda}$ is the eigenvector of Γ that is associated with its unitary eigenvalue.

In the presence of an aggregate shock the algorithm can obviously not be used to compute a stationary distribution. But the same logic applies and a distribution λ, which then has dimensionality $(\#_e \cdot \#_\epsilon \cdot \#_A \times 1)$ and a transition matrix Γ which then has dimensionality $(\#_e \cdot \#_\epsilon \cdot \#_A \times \#_e \cdot \#_\epsilon \cdot \#_A)$ can be set up. The simulation

then consists of drawing a random sequence of realizations of the aggregate shock and computing $\tilde{\lambda}_1 = \tilde{\lambda}_0' \Gamma$. The code is available upon request.

A.4 Extensions (in progress)

Introducing Financial constraints for all firms I introduce a working-capital assumption into the model. Firms have to pay a fraction λ of their period expenses at the beginning of the period. Those expenses include the wage bill $w \cdot e$ and adjustment costs, including the fixed cost. To finance those costs, firms borrow use their housing value q^h as collateral just like entrants in the benchmark model. For any uncollateralized fraction of the loan the firm has to pay the higher interest rate $\hat{R} \geq 1$. At the end of the period, once profits are realized, the entrepreneur pays back the loan to the bank. I assume that an entrepreneur's realization of ε is perfectly observable by the bank. This modification essentially makes q^h a state variable of the entrepreneur's problem.

Results in progress...

Alternative wage setting To apply the Stole and Zwiebel (1996) framework I assume that the agent's utility function be linear, $Z(c) = c$. As in Elsby and Michaels (2013) this is done to obtain a closed form solution for the problem. Details to follow.

Appendix to Chapter 2

The appendix describes our method of solving the planner's problem. The approach taken for characterizing the law of motion for the joint distribution, Γ, is described in the text. Here we focus on the planner's choice of capital in the reallocation process.

Any vector of capital allocated across adjustable plants $k(\varepsilon)$ will have associated values for μ_a and ϕ_a. Create a grid for potential vectors $k(\varepsilon)$. To do so, define two benchmarks for the planners decision regarding the allocation of capital across those plants that are in F^A. Define k^{MAX} as the vector where marginal products are equalized across plants. This vector was found in (2.9) for the frictionless benchmark case above. In the presence of Calvo adjustment costs, the planner will not reallocate more capital between plants than under the allocation rule k^{MAX}, but possibly less. The second benchmark will be called k^{MIN} and is simply the case where capital is equally distributed across adjustable plants (i.e. no reallocation). The idea behind this procedure is that the planner will choose a vector $k(\varepsilon)$ which is *between* k^{MAX} and k^{MIN}, meaning that the planner will reallocate some capital between plants, but not as much as under the frictionless benchmark. We consider convex combinations of k^{MAX} and k^{MIN}.

Define a variable m, that takes values between zero and one and determines a potential vector of $k(\varepsilon)$'s as follows: $k_m = m \cdot k^{\text{MAX}} + (1 - m) \cdot k^{\text{MIN}}$. For each k_m compute $\mu_m = E(k_m(\varepsilon)^\alpha)$ and $\phi_m = Cov(\varepsilon, k_m(\varepsilon)^\alpha)$ characterizing this vector. This allows the calculation of output associated with m. The planner optimizes over m and this translates into μ_m, ϕ_m.

To check the robustness of this procedure start from a model with the baseline parameters without any exogenous shocks. It turns out that the planner chooses $m = 0.9508$, which means that the optimal vector $k(\varepsilon) = 0.9508 \cdot k^{\text{MAX}} + 0.0492 \cdot k^{\text{MIN}}$, so capital reallocation is about 5% lower compared to the frictionless benchmark. In order to see how good of an approximation the decision rule 'm' is, we apply the following procedure.

We work directly with the planner's value of the steady state (SS) allocation. The simplified version of the value function has only two states, μ_n and ϕ_n, so there will

be a value $V(\mu_n^{SS}, \phi_n^{SS})$ associated to the steady state. This value is equal to forever receiving the output associated with the amount of reallocation 'm' times the fraction of adjustable plants, plus the output associated with the SS state vector times the fraction of non-adjustable plants.

$$V(\mu^{SS}, \phi^{SS}) = \frac{\int_{\varepsilon \in F^A} \varepsilon k(\varepsilon)^\alpha f(\varepsilon) d\varepsilon \ + (1 - \pi)(E(\varepsilon)\mu^{SS} + \phi^{SS})}{1 - \beta} \qquad (10)$$

The planner can now choose any allocation of capital across plants. This allocation implies a mapping into the values of μ_n and ϕ_n. The planner will be allowed to choose the allocation that maximizes the expression for $V(\mu^{SS}, \phi^{SS})$ above. Being bound to the same grid, the resulting vector is identical to the one previously found. We now perturb this vector in order to find profitable deviations that keep the aggregate capital stock constant. The perturbation adds a random vector with mean zero to the k-vector that maximized (27) given the grid. If the resulting vector produces a higher lifetime utility, the k-vector is updated accordingly. This procedure is repeated 1,000,000 times. The results show that our grid for m comes extremely close to the optimal solution. Although profitable deviations *are* possible, they remain very small: the difference in output is around 0.01%.

Appendix to Chapter 3

Parameter choices explained

Overtime premium λ (1) "As a general principle, in cases of overtime exceeding maximum working hours, an employee is entitled to overtime premium pay. By the application of a pay supplement it is usually at least 25 per cent higher than the pay which the employee could have earned if the work had been done in normal working time. Employees can claim overtime premium pay even if such overtime was prohibited by statute or collective agreement. No entitlement exists if the overtime exceeding statutory working hours is done to enable employees to catch up on work that has not been done through their own fault." (see eurofund). (2) Some collective agreements and employment contracts provide for an increased payment for overtime hours. Appropriate are 25 per cent on regular working days and 50 per cent on Sundays and Holidays. (see Eichhorst, Wener "Traditionelle Beschaeftigungsverhaeltnisse im Wandel. Benchmarking Deutschland: Normalarbeitsverhaeltnisse auf dem Rueckzug. table 8). (3) Hunt sets a similar parameter, which characterizes overtime payment, to 0.25. (see page 136). (4) According to Burda and Hunt, additional flexibility in terms of working hours (for instance time accounts) reduced the actual overtime premium paid by firms significantly over the last years. (see footnote 18, page 301).

Employer social security contribution rate δ (1) "Employers' social contributions are payments (either actual or imputed) by employers which are intended to secure for their employees the entitlement to social benefits should certain events occur, or certain circumstances exist, that may adversely affect their employees' income or welfare - sickness, accidents, redundancy, retirement, etc." Statistics from the OECD suggest that employer social security contribution is around 20.6 per cent of the gross wage. (2) An international comparison published by EUROSTAT shows that the contribution rate is around 23 per cent. (3) On the contrary the German statistics office DESTATIS calculates a rate which is 28 per cent. (4) More specifi-

cally, this rate includes contributions to a wide range on insurance schemes: health insurance (7.3 per cent), long term care (0.975 per cent), pension plans (9.8 per cent), unemployment insurance (1.5 per cent), insurance against bankruptcy (0.04 pH).

Share of social contribution paid by firm in case of STW μ

Replacement rate for compensation lost due to hours reduction ϕ "The calculation of the individual STC is just like the calculation of unemployment benefits: for the lost working hours, employees with at least one dependent child receive a compensation of 67 per cent of the net difference to the regular wage, whereas those without dependents get 60 per cent; for a loss of work of 100 per cent, the STC has the same amount as the unemployment benefits." (see Crimmann, Wiesner and Bellmann (2010), page 19). The budget for partial unemployment is financed through the government.

Table 5: Coefficients from a Panel-VAR on Hours growth and Employment growth. Dependent variable in column.

x	Dn	Dhours	Dn before 2003	Dhours before 2003	Dn after 2002	Dhours after 2002
L.Dn	.15862217	.08074814	.14221813	.05946681	.16131513	.07616513
L.Dhours	.01907817	-.00802338	.05100504	.07231959	.00310598	.01567351

Contracted hours $\bar{h}$

Bibliography

Acemoglu, D. and Wlliam B. Hawkins (2013). 'Search with multi-worker firms'. *Theoretical Economics.*

Adda, J. and Russell W Cooper (2003). *Dynamic economics: quantitative methods and applications.* MIT Press. Cambridge, Mass.

Avery, R. B., Raphael W. Bostic and Katherine A. Samolyk (1998). 'The role of personal wealth in small business finance'. *Journal of Banking & Finance* **22**(6–8), 1019–1061.

Bachmann, R. (2011). 'Understanding jobless recoveries'.

Bachmann, R. and Christian Bayer (2013). ''Wait-and-See' business cycles?'. *Journal of Monetary Economics* **60**(6), 704–719.

Balleer, A., Britta Gehrke, Wolfgang Lechthaler and Christian Merkl (2013). Does short-time work save jobs? a business cycle analysis. IZA Discussion Paper 7475. Institute for the Study of Labor (IZA).

Bartelsman, E., J. Haltiwanger and S. Scarpetta (2013). 'Cross-country differences in productivity: The role of allocation and selection'. *American Economic Review.*

Bassett, W. F., Simon Gilchrist, Gretchen C. Weinbach and Egon Zakrajsek (2011). Improving our ability to monitor bank lending. NBER chapters. National Bureau of Economic Research, Inc.

Basu, S. and J. Fernald (1997). 'Returns to scale in u.s. production: Estimates and implications'. *Journal of Political Economy* **105**, 249–83.

Berger, D. (2012). 'Countercyclical restructuring and jobless recoveries'.

Bernanke, B. S., Mark Gertler and Simon Gilchrist (1999). The financial accelerator in a quantitative business cycle framework. Handbook of macroeconomics. Elsevier.

Bloom, N. (2009). 'The impact of uncertainty shocks'. *Econometrica* **77**(3), 623–685.

Bloom, N., Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten and Stephen J Terry (2012). Really uncertain business cycles. NBER Working Paper #18245.

Board, F. R. (2011). 'Survey of consumer finances'.

Brenke, K., Ulf Rinne and Klaus F. Zimmermann (2011). Short-time work: The german answer to the great recession. IZA Discussion Papers 5780. Institute for the Study of Labor (IZA).

Burda, M. C. and Jennifer Hunt (2011). 'What explains the german labor market miracle in the great recession'. *Brookings Papers on Economic Activity* **42**(1), 273–335.

Calvo, G. (1983). 'Staggered prices in a utility-maximizing framework'. *Journal of monetary Economics* **12**(3), 383–398.

Chaney, T., David Sraer and David Thesmar (2012). 'The collateral channel: How real estate shocks affect corporate investment'. *American Economic Review* **102**(6), 2381–2409.

Chodorow-Reich, G. (2013). 'The employment effects of credit market disruptions: Firm-level evidence from the 2008-09 financial crisis'. *The Quarterly Journal of Economics*.

Clementi, G. L. and Dino Palazzo (2010). Entry, exit, firm dynamics, and aggregate fluctuations.

Coles, M. G. and Ali Moghaddasi Kelishomi (2011). New business start-ups and the business cycle. CEPR Discussion Paper 8588. C.E.P.R. Discussion Papers.

Cooper, R. and J. Haltiwanger (2006*a*). 'On the nature of the capital adjustment costs'. *Review of Economic Studies* **73**, 611–34.

Cooper, R., G. Gong and P. Yan (2011). Dynamic labor demand in china: Public and private objectives.

Cooper, R., Guan Gong and Ping Yan (2012). Costly labor adjustment: Effects of china's employment regulations. Working Paper 17948. National Bureau of Economic Research.

Cooper, R., John Haltiwanger and Jonathan L Willis (2007). 'Search frictions: Matching aggregate and establishment observations'. *Journal of Monetary Economics* **54**(Supplement 1), 56–78.

Cooper, R. W. and Immo Schott (2014). Capital reallocation and aggregate productivity. Working Paper 19715. National Bureau of Economic Research.

Cooper, R. W. and John C. Haltiwanger (2006*b*). 'On the nature of capital adjustment costs'. *Review of Economic Studies* **73**(3), 611–633.

Costain, J. S. and Anton Nakov (2013). Logit price dynamics. SSRN Scholarly Paper ID 2212601. Social Science Research Network. Rochester, NY.

Crimmann, A., F. Wiebner and L. Bellman (2010). *The German work-sharing scheme: An instrument for the crisis*. ILO, Conditions of Work and Employment Programme.

Darby, M. R., John C. Haltiwanger and Mark W. Plant (1986). The ins and outs of unemployment: The ins win. NBER Working Paper 1997. National Bureau of Economic Research, Inc.

Den Haan, W. J. (2010). 'Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents'. *Journal of Economic Dynamics and Control* **34**(1), 79–99.

Den Haan, W. J., Garey Ramey and Joel Watson (2000). 'Job destruction and propagation of shocks'. *American Economic Review* **90**(3), 482–498.

Dennis Jr., W. J. (2010). 'Small business credit in a deep recession'.

Dennis Jr., W. J. (2011). 'Financing small businesses small business and credit access'.

Dennis Jr., W. J. (2012). 'Small business, credit access and a lingering recession'.

Drautzburg, T. (2013). 'Entrepreneurial tail risk: Implications for employment dynamics'.

Dromel, N. L., Elie Kolakez and Etienne Lehmann (2010). 'Credit constraints and the persistence of unemployment'. *Labour Economics* **17**(5), 823–834.

Eisfeldt, A. L. and Adriano A Rampini (2006). 'Capital reallocation and liquidity'. *Journal of Monetary Economics* **53**(3), 369–399.

Elsby, M. W. L. and Ryan Michaels (2013). 'Marginal jobs, heterogeneous firms, & unemployment flows'. *American Economic Journal: Macroeconomics.*

Elsby, M. W. L., Ryan Michaels and Gary Solon (2009). 'The ins and outs of cyclical unemployment'. *American Economic Journal: Macroeconomics* **1**(1), 84–110.

Fort, T., John Haltiwanger, Ron S. Jarmin and Javier Miranda (2013). How firms respond to business cycles: The role of firm age and firm size. SSRN Scholarly Paper ID 2278454. Social Science Research Network. Rochester, NY.

Foundation, E. M. K. (2013). The kauffman foundation and LegalZoom startup environment index 2012. SSRN Scholarly Paper ID 2226668. Social Science Research Network. Rochester, NY.

Galí, J., Frank Smets and Rafael Wouters (2012). 'Slow recoveries: A structural interpretation'. *Journal of Money, Credit and Banking* **44**, 9–30.

Gilchrist, S. and Egon Zakrajšek (2012). 'Credit spreads and business cycle fluctuations'. *American Economic Review* **102**(4), 1692–1720.

Gilchrist, S., Jae Sim and Egon Zakrajsek (2014). Uncertainty, financial frictions and investment dynamics. NBER Working Papers 20038. National Bureau of Economic Research, Inc.

Gourio, F. and A.K. Kashyap (2007). 'Investment spikes: New facts and a general equilibrium exploration'. *Journal of Monetary Economics* **54**, 1–22.

Gourio, F. and Jianjun Miao (2010). 'Firm heterogeneity and the long-run effects of dividend tax reform'. *American Economic Journal: Macroeconomics* **2**(1), 131–68.

Greenspan, A. and James Kennedy (2005). 'Estimates of home mortgage originations, repayments, and debt on one-to-four-family residences'.

Hall, R. E. (2005a). 'Employment efficiency and sticky wages: Evidence from flows in the labor market'. *The Review of Economics and Statistics* **87**(3), 397–407. ArticleType: research-article / Full publication date: Aug., 2005 / Copyright © 2005 The MIT Press.

Hall, R. E. (2005b). 'Employment fluctuations with equilibrium wage stickiness'. *The American Economic Review* **95**(1), 50–65.

Haltiwanger, J. C., Ron S. Jarmin and Javier Miranda (2010). 'Who creates jobs? small vs. large vs. young'. *National Bureau of Economic Research Working Paper Series.*

Haltiwanger, J., Ron S. Jarmin and Javier Miranda (2012). 'Business dynamics statistics briefing: Where have all the young firms gone?'. *SSRN eLibrary.*

Hawkins, W. B. (2011). 'Do large-firm bargaining models amplify and propagate aggregate productivity shocks?'.

Hopenhayn, H. A. (1992). 'Entry, exit, and firm dynamics in long run equilibrium'. *Econometrica* **60**(5), 1127–50.

Hopenhayn, H. and Richard Rogerson (1993). 'Job turnover and policy evaluation: A general equilibrium analysis'. *Journal of Political Economy* **101**(5), 915–38.

Hsieh, C.-T. and P. Klenow (2009). 'Misallocation and manufacturing tfp in china and india'. *Quarterly Journal of Economics* **CXXIV**(4), 1403–1448.

Jermann, U. and Vincenzo Quadrini (2012). 'Macroeconomic effects of financial shocks'. *American Economic Review* **102**(1), 238–271.

Kaas, L. and Philipp Kircher (2011). Efficient firm dynamics in a frictional labor market. SSRN Scholarly Paper ID 1756684. Social Science Research Network. Rochester, NY.

Kehrig, M. (2011). The cyclicality of productivity dispersion. The University of Texas at Austin, Job Market Paper.

Kennedy, J. (2011). 'A simple method for estimating gross equity extracted from housing wealth'.

Khan, A. and J. Thomas (2003). 'Nonconvex factor adjustments in equilibrium business cycle models: Do nonlinearities matter?'. *Journal of Monetary Economics* **50**, 331–60.

Khan, A. and Julia K. Thomas (2008). 'Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics'. *Econometrica* **76**(2), 395–436.

Khan, A. and Julia K. Thomas (2013). 'Credit shocks and aggregate fluctuations in an economy with production heterogeneity'. *Journal of Political Economy.*

Kiyotaki, N. and John Moore (1997). 'Credit cycles'. *Journal of Political Economy* **105**(2), 211–248.

Krause, M. U. and Harald Uhlig (2012). 'Transitions in the german labor market: Structure and crisis'. *Journal of Monetary Economics* **59**(1), 64–79.

Krusell, P. and Anthony A. Smith (1998). 'Income and wealth heterogeneity in the macroeconomy'. *Journal of Political Economy* **106**(5), 867–896.

Lee, Y. and Toshihiko Mukoyama (2012). Entry, exit and plant-level dynamics over the business cycle. Technical report. Federal Reserve Bank of Cleveland.

Liu, Z., Jianjun Miao and Tao Zha (2013*a*). Land prices and unemployment. Working Paper 19382. National Bureau of Economic Research.

Liu, Z., Pengfei Wang and Tao Zha (2013*b*). 'Land price dynamics and macroeconomic fluctuations'. *Econometrica* **81**(3), 1147–1184.

Lucas, R. E. (1978). 'On the size distribution of business firms'. *Bell Journal of Economics* **9**(2), 508–523.

Mann, R. J. (1998). 'Comment on avery, bostic and samolyk'. *Journal of Banking & Finance* **22**(6-8), 1062–1066.

Midrigan, V. (2011). 'Menu costs, multiproduct firms, and aggregate fluctuations'. *Econometrica* **79**(4), 1139–1180.

Midrigan, V. and Daniel Yi Xu (2014). 'Finance and misallocation: Evidence from plant-level data'. *American Economic Review*.

Möller, J. (2012). 'The german labor market response in the world recession: demystifying a miracle'. *Zeitschrift für ArbeitsmarktForschung - Journal for Labour Market Research* **42**(4), 325–336.

Moon, J. (2009). 'Small business finance and personal assets'. *Federal Reserve Bank of San Francisco: Community Investments*.

Mortensen, D. T. and Christopher A. Pissarides (1994). 'Job creation and job destruction in the theory of unemployment'. *The Review of Economic Studies* **61**(3), 397–415.

Moskowitz, T. J. and Annette Vissing-Jorgensen (2002). 'The returns to entrepreneurial investment: A private equity premium puzzle?'. *American Economic Review* **92**(4), 745–778.

Olley, S. G. and Ariel Pakes (1996). 'The Dynamics of Productivity in the Telecommunications Equipment Industry'. **64**(6), 1263–1297.

Petrosky-Nadeau, N. (2013). 'Credit, vacancies and unemployment fluctuations'. *Review of Economic Dynamics.*

Pissarides, C. A. (2009). 'The unemployment volatility puzzle: Is wage stickiness the answer?'. *Econometrica* **77**(5), 1339–1369.

Pissarides, C. A. and Barbara Petrongolo (2001). 'Looking into the black box: A survey of the matching function'. *Journal of Economic Literature* **39**(2), 390–431.

Restuccia, D. and Richard Rogerson (2008). 'Policy distortions and aggregate productivity with heterogeneous establishments'. *Review of Economic Dynamics.*

Rinne, U. and Klaus F. Zimmermann (2011). Another economic miracle? the german labor market and the great recession. IZA Discussion Papers 6250. Institute for the Study of Labor (IZA).

Robb, A. M. and David T. Robinson (2012). 'The capital structure decisions of new firms'. *Review of Financial Studies* p. hhs072.

Sanchez, J. M. and Constanza S. Liborio (2012). 'Starting a business during a recovery: this time, it's different'. *The Regional Economist.*

Schmalz, M., David Sraer and David Thesmar (2013). 'Housing collateral and entrepreneurship'.

Sedlacek, P. (2011). 'Firm age, business cycles and aggregate labor market dynamics'.

Shimer, R. (2005). 'The cyclical behavior of equilibrium unemployment and vacancies'. *American Economic Review* **95**(1), 25–49.

Shimer, R. (2012). 'Reassessing the ins and outs of unemployment'. *Review of Economic Dynamics* **15**(2), 127–148.

Siemer, M. (2013). Firm entry and employment dynamics in the great recession. SSRN Scholarly Paper ID 2172594. Social Science Research Network. Rochester, NY.

Stole, L. A. and Jeffrey Zwiebel (1996). 'Intra-firm bargaining under non-binding contracts'. *Review of Economic Studies* **63**(3), 375–410.

Sveen, T. and Lutz Weinke (2005). 'New perspectives on capital, sticky prices, and the taylor principle'. *Journal of Economic Theory* **123**(1), 21–39.

Thomas, J. (2002). 'Is lumpy investment relevant for the business cycle?'. *Journal of Political Economy* **110**, 508–34.

Townsend, R. M. (1979). 'Optimal contracts and competitive markets with costly state verification'. *Journal of Economic Theory* **21**(2), 265–293.

{United States Congressional Oversight Panel} (2010). *The Small Business Credit Crunch and the Impact of the TARP*. U.S. Government Printing Office.

{United States Congressional Oversight Panel} (2011). *Congressional Oversight Panel March Oversight Report: The Final Report of the Congressional Oversight Panel*. U.S. Government Printing Office.

List of Figures

List of Tables