

The London School of Economics and Political Science

**Essays on Frictional Labour Markets
with Heterogeneous Agents**

Markus Riegler

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

London, May 2015

Declaration

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

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent.

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

I declare that my thesis consists of 39,846 words.

Statement of conjoint work

I confirm that chapter 2 was jointly co-authored with Professor Wouter J. Den Haan and Dr Pontus Rendahl and I contributed at least 50% of this work.*

Statement of inclusion of previous work

I confirm that chapter 3 includes revised results of previous study for the degree of MRes I undertook at the London School of Economics and Political Science. That degree was awarded in 2011.

*Financial support from the ESRC grant to the Centre for Macroeconomics is gratefully acknowledged. We also would like to thank Chris Carroll, Greg Kaplan, Kurt Mitman, Rana Sajedi, Victor Rios-Rull, and Petr Sedlacek for useful comments.

Abstract

The first chapter discusses the effects of uncertainty shocks on the labour market. Using US data, I show empirically that increases in unemployment are due to both an increase in the separation rate and a decrease in the job-finding rate. By contrast, standard search and matching models predict an increase in the job finding rate in response to an increase in the cross-sectional dispersion of firms' productivity levels. To explain observed responses in labour market transition rates, I develop a search and matching model in which heterogeneous firms face a decreasing returns to scale technology, firms can hire multiple workers, and job flows do not necessarily coincide with worker flows. Costly job creation is key to obtaining a decrease in the job-finding rate after an increase in uncertainty. Standard numerical solution techniques cannot be used to obtain an accurate solution efficiently and I propose an alternative algorithm to overcome this problem.

The second chapter studies business cycles when markets are incomplete, nominal wages do not respond one-for-one to price level changes, and labour markets are characterized by matching frictions. During recessions, idiosyncratic labour income risk increases as workers worry about being unemployed. This induces workers to save more. We allow such precautionary savings – in principle – to end up in both an unproductive asset (money) and a productive assets (firm ownership). The increased demand for money puts deflationary pressure on prices. If nominal wages are not sufficiently responsive to deflationary pressures, wage costs and the unemployment rate are pushed up, which in turn intensifies the fear of becoming unemployed. Unemployment benefits improve welfare in our economy.

The third chapter analyses the inefficiencies created in a search and matching model that allows for on-the-job search. First, the Hosios rule for the efficient level of the worker's bargaining power is adapted in a simple model. As the average gain of a new match is lower when some job seekers already have a job, the efficient level of labour market tightness should be lower and the worker's bargaining power higher than in a model devoid of on-the-job search. Second, the decision of when to perform on-the-job search is endogenised. Too much on-the-job search is taking place, because workers do not fully incorporate their current firms' loss when they quit. Partial wage commitment improves the efficiency of the on-the-job search decision and the efficient level can be obtained.

Acknowledgements

I am very grateful to Wouter den Haan and Philipp Kircher for supervising me at LSE. My knowledge and this thesis has benefited substantially from working together with Wouter. I also want to thank colleagues and seminar participants for providing comments on my research. I gratefully acknowledge financial support by the Economic and Social Research Council (ESRC).

It has been a great experience for me to live in London and study at LSE, for which I want to especially thank Johannes Boehm, Alex Clymo, Sergio de Ferra, Bracha Feit Katz, Michael Geppert, Jennifer Helgeson, Stephanie Lackner, Andrea Lanteri, Sam Marden, Lukas Mayr, Ana McDowall, Francesco Palazzo, Joao Pessoa, Jonathan Pinder, Pedro Pinto, Pedro Souza, and Fabian Winkler.

I am greatly indebted to my family that has shaped my life until now: my parents and my brother, as well as my grandparents, aunts, and uncles.

Contents

1 The Impact of Uncertainty Shocks on the Job-Finding Rate and Separation Rate	11
1.1 Introduction	11
1.2 Effects of uncertainty in US data	15
1.3 Model	17
1.3.1 Setup	17
1.3.2 Solution to firm problem	20
1.3.3 Characterisation of the optimal firm behaviour using three cutoffs	22
1.3.4 Market clearing	24
1.4 Parametrization	24
1.5 Effects of uncertainty in the model	28
1.5.1 Comparative statics	29
1.5.2 Decomposition of effects	32
1.5.3 IRFs of uncertainty shocks	37
1.6 Comparison with US time series	40
1.7 Algorithm to solve the model	43
1.7.1 General description of the residual method	43
1.7.2 Application of the residual method in this model	45
1.8 Conclusion	46
Appendices to Chapter 1	48
1.A Empirical part	48
1.A.1 Data description	48
1.A.2 Approximation of changes in unemployment	49
1.A.3 Robustness checks	50
1.B Comparative statics	53
1.B.1 Sensitivity analysis	53
1.C Decomposition exercise	60
1.C.1 Derivation of change in cutoffs for the wait-and-see effect	60
1.C.2 Asymmetry of wait-and-see effect	63
1.D Comparison with US time series	63
2 Unemployment (Fears) and Deflationary Spirals: The Role of Sticky Nominal Wages	66
2.1 Introduction	66
2.2 Model	69
2.2.1 Households	69

2.2.2	Firms	70
2.2.3	Government	71
2.2.4	Firm creation and equity market	71
2.2.5	Money market	72
2.2.6	Equilibrium and model solution	73
2.2.7	Discounting firm profits correctly with heterogeneous ownership	73
2.3	Empirical motivation	74
2.3.1	Deflationary pressure and sticky nominal wages	74
2.3.2	Inability to insure against unemployment risk	77
2.3.3	Savings and idiosyncratic uncertainty	78
2.4	Calibration	79
2.5	Agents' consumption, investment and portfolio decisions	81
2.5.1	Post-displacement consumption	81
2.5.2	Investment decisions	85
2.6	Business cycle properties	88
2.6.1	Benchmark: Business cycles with sticky nominal wages	89
2.6.2	Business cycles with flexible nominal wages	91
2.7	Government policy	94
2.7.1	Unemployment-insurance policies	94
2.7.2	Monetary policy	101
	Appendices to Chapter 2	103
2.A	Data Sources	103
2.B	Solution algorithm	103
2.B.1	Solution algorithm for heterogeneous agent model	103
2.C	Equivalence with standard matching framework	109
3	Efficiency of On-the-Job Search in a Search and Matching Model with Endogenous Job Destruction	111
3.1	Introduction	111
3.2	Simple model	113
3.2.1	Setup	113
3.2.2	Social planner's problem	114
3.2.3	Decentralised economy	115
3.2.4	Efficiency condition	117
3.3	Model with endogenous job destruction and on-the-job search	118
3.3.1	Social planner's problem	119
3.3.2	Decentralised equilibrium without commitment	121
3.3.3	Comparison of the decentralised equilibrium with the efficient outcome	125
3.4	Decentralised equilibrium with (partial) commitment	126
3.4.1	Setup	126
3.4.2	Wage bargaining when the bargaining set is non-convex	129
3.4.3	Determining the equilibrium	130
3.5	Conclusion	132

Appendices to Chapter 3	134
3.A Solution of social planner's problem in the full model	134
3.B Derivation of jointly optimal OTJ search decision with (partial) commitment . . .	136
3.C Proof of Lemma 3	136
3.D Derivation of equilibrium with (partial) commitment	138
References	140

List of Figures

1.1	Index of implied volatility constructed from US stock options.	15
1.2	Impulse response functions after a volatility shock in US data.	16
1.3	Marginal value of a position.	20
1.4	Optimal choice of positions for a given level of idiosyncratic productivity. . . .	23
1.5	Cutoffs for a given aggregate state.	24
1.6	Responses after a permanent 10% increase of idiosyncratic volatility.	30
1.7	Cumulative distribution function of continuation value of a firm.	31
1.8	Wait-and-see effect: change in cutoffs after a 10% increase in uncertainty.	33
1.9	Decomposition of increase in uncertainty into wait-and-see, realized volatility, and general equilibrium effects.	34
1.10	General equilibrium effect (change in market tightness).	36
1.11	Impulse response functions for idiosyncratic volatility shock in the model. . . .	38
1.12	State dependence of IRF for idiosyncratic volatility shock in the model.	39
1.13	State dependence of employment changes for idiosyncratic volatility shock. . . .	40
1.14	Comparison of model without uncertainty shocks with US Time Series.	41
1.15	Comparison of full model with US Time Series.	42
1.16	Accuracy of algorithm.	46
1.17	Appendix: IRF after a shock to labour productivity in US data	49
1.18	Appendix: IRF after a volatility shock in US data using alternative identification. .	50
1.19	Appendix: IRF after a volatility shock in US data using VIX.	51
1.20	Appendix: IRF after a volatility shock in US data using a different filter.	51
1.21	Appendix: IRF after a volatility shock in US data using a different filter.	52
1.22	Appendix: Sensitivity analysis using creation costs.	55
1.23	Appendix: Sensitivity analysis using different size of shock.	56
1.24	Appendix: Sensitivity analysis using different cyclicalities of H	56
1.25	Appendix: Sensitivity analysis using higher filling cost.	57
1.26	Appendix: Sensitivity analysis with respect to exogenous turnover.	57
1.27	Appendix: Sensitivity analysis with respect to wage equation.	58
1.28	Appendix: Sensitivity analysis with respect to curvature in production function. .	58
1.29	Appendix: Sensitivity analysis with respect to arrival rate of shock.	59
1.30	Appendix: Sensitivity analysis with respect to distribution of shock.	59
1.31	Appendix: Comparison of cutoffs with perfect insurance economy.	61
1.32	Appendix: Comparison of model without uncertainty shocks and increased wage stickiness with US Time Series.	64

1.33 Appendix: Comparison of full model backing out volatility shock with US Time Series.	65
2.1 Key Eurozone variables before and after the financial crisis.	75
2.2 Evolution of consumption drop over the unemployment spell.	82
2.3 Distribution of the unemployed.	82
2.4 Financial assets at the beginning of an unemployment spell.	83
2.5 Consumption drop upon unemployment.	84
2.6 Portfolio shares in liquid asset.	85
2.7 Money demand (real).	87
2.8 Post displacement equity holdings.	88
2.9 Post displacement money holdings.	89
2.10 Impulse responses with sticky nominal wages.	90
2.11 Impulse responses with flexible nominal wages.	92
2.12 Switch to higher unemployment insurance at the start of the recession.	98
2.13 Welfare gains from changing unemployment insurance	99
2.14 Switch to cyclical UI policy at the start of the recession.	101
3.1 Non-convex bargaining set.	129

List of Tables

1.1 Forecast error variance decomposition.	17
1.2 Summary of firm's optimal action.	23
1.3 Model parameters.	25
1.4 Recalibrated parameters.	25
1.5 State dependence of impulse response functions.	39
3.1 Optimal decision and worker surplus with (partial) commitment	131

Chapter 1

The Impact of Uncertainty Shocks on the Job-Finding Rate and Separation Rate

1.1 Introduction

There has been recent interest in the importance of uncertainty shocks on macroeconomic variables. In particular, it has been suggested that an increase in uncertainty during the Great Recession has contributed to higher unemployment. There are two channels through which uncertainty shocks can affect the unemployment rate. First, it is possible that higher uncertainty increases the job separation rate so that employed workers are more likely to lose their jobs. Second, higher uncertainty could reduce the job-finding rate, which makes it harder for the unemployed to find a job. In my empirical analysis, I confirm that higher uncertainty reduces employment. This reduction is driven by both a *higher separation rate* of workers and a *lower job-finding rate* of unemployed. The reduced job-finding rate contributes more to higher unemployment than the separation rate after an increase in uncertainty. Existing search and matching models of the labour market cannot account for this significant contribution of the job-finding rate.

In the literature, there are several measures of uncertainty both on the micro-level and on the macro-level. In the context of this chapter, higher uncertainty means a higher expected and higher realized dispersion of idiosyncratic productivity across firms. Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) use microdata to show that plant-level shocks become more dispersed in recessions. This makes recessions, and in particular the Great Recession, appear to be a combination of a negative first-moment shock and a positive second-moment shock at the establishment-level.

An increase in dispersion of productivity across firms implies that more firms get hit by large negative shocks that lead to the layoff of workers. This so-called *realized volatility effect* raises the separation rate in times of higher uncertainty. At the same time it also means that more firms experience large positive shocks, which can increase hiring. Therefore, a reduction of the job-finding rate as seen in the data cannot be explained by the realized volatility effect. In contrast, what is known as the *“wait-and-see” effect* has the potential to reduce the

job-finding rate: if firms are more uncertain about future realizations of productivity, they might become reluctant to hire workers in the presence of adjustment costs. This happens if higher uncertainty reduces the expected continuation value of a newly hired worker.¹ For a model to explain the fall of the job-finding rate, it is important that a wait-and-see effect that makes firms more reluctant to hire exists and that it is sufficiently strong relative to the realized volatility effect.

The standard search and matching model assumes constant returns to scale, which rules out meaningful heterogeneity in firm size. If there are no idiosyncratic shocks, firm size is irrelevant and not pinned down. If there are idiosyncratic productivity shocks, as in Mortensen and Pissarides (1994), firms cannot be allowed to employ an unrestricted amount of workers. Because if it were possible to hire more workers at a given productivity level, the most productive firm would hire all workers. An implication of this restriction to hire additional workers is that the value of a matched worker to a firm increases without bounds with idiosyncratic productivity. In contrast, if the firm experiences low productivity in a future period, it has the option to destroy the match. Therefore, the lower bound for the value of the match is 0. This asymmetry turns out to be important for the effect of uncertainty on labour market variables. If volatility of idiosyncratic productivity increases, more matches are destroyed by sufficiently negative realizations of productivity. Therefore, the realized volatility effect increases the separation rate. In addition, the absence of an upper bound implies that the expected value of a match increases. Hence the wait-and-see effect makes firms more reluctant to fire but it actually *increases* the expected value of a new match. This last effect increases vacancy posting and the job-finding rate in response to higher volatility, which is at odds with patterns observed in the data.

I develop a more general framework which does allow for meaningful heterogeneity across firms. In particular, I assume that a firm's production function exhibits decreasing returns to scale, which means that there can be idiosyncratic productivity shocks without restricting firms in the number of workers they can hire. In equilibrium, firms with various levels of idiosyncratic productivity will hire workers if they are small relative to their productivity, because their marginal product is high. To be precise, firms hire workers until the marginal value of an additional worker equals the hiring cost.

In this model, it is possible for the wait-and-see effect to reduce hirings if higher uncertainty reduces the marginal value of a worker. The marginal value of a worker consists of the marginal profit in the current period plus the expected continuation value the firm gets from the worker being at the firm in the next period. To understand the expected continuation value of a match it is important to consider the contribution of this one worker in future periods for different realizations of the idiosyncratic productivity shocks. Imagine a firm becomes very productive in the next period. In the standard framework, this means that the match becomes very valuable for the firm. In a framework in which the firm is allowed to hire more than one worker, however, the marginal value of one additional worker can never exceed the hiring cost of an additional worker. In this case, the continuation value of a worker hired in the current period will be equal to the discounted hiring cost that the firm saves in the next period. The ability to hire additional workers, when productivity is high, is of course valuable to the

¹By the same logic, firms become more reluctant to fire workers if the expected continuation value increases in times of higher uncertainty because large improvements of productivity become more likely. The firm can wait for such a shock and fire the worker in the future if productivity deteriorates.

firm. But what matters for a firm's optimal choice is the value of the marginal match and this value is bounded by the option to hire additional workers. This puts an upper bound on the continuation value of a worker when the firm becomes highly productive.

The upper bound of the marginal value function is a key difference compared to the standard framework. It is an implication of the option to hire additional workers, comparable to the lower bound of the continuation value implied by the option of destroying matches. The upper bound introduces a concave part to the marginal value function, which has important implications for the question of how uncertainty affects the job-finding rate. Due to the concavity, it is possible that higher idiosyncratic volatility reduces the value of the marginal worker. In turn, this reduces the number of vacancies posted by the firm and the job-finding rate for the unemployed.

In my model, I distinguish between workers and the positions they fill. There are sunk costs of creating a position. They do not have to be paid again when a firm wants to replace a worker that leaves. For example, if a worker turns out to be not a good match for the firm, or if he finds a job at another firm, his workplace still exists at the previous firm. It is then less costly for the firm to hire a new worker who does the same task. The creation cost could be interpreted as the share of capital needed for a position that cannot be recovered if it is shut down, or as the cost of finding new clients when a firm increases its size. As a result of this sunk cost, some firms in my model will fill existing empty positions even though they do not want to create new positions. These firms shrink over time, because some of their positions become obsolete. But they still hire workers, because the quit rate of existing workers exceeds the rate of obsolescence. This is consistent with differences between worker and job flows as documented by Davis, Faberman, and Haltiwanger (2012): using JOLTS establishment data, they calculate that hires amount to around 10% of employment at shrinking establishments on a quarterly basis.

I show that the creation cost is important for the response of the job-finding rate to uncertainty shocks, because it strengthens the wait-and-see effect: as explained above, the marginal value of a position is bounded from above by the cost of creation, because the firm always has the option to create another position. The marginal value is at the upper bound, if a firm creates positions in the current period. Therefore, it does not increase if that firm becomes more productive in the future. In contrast, negative productivity shocks reduce the marginal value of a position. Higher uncertainty thus reduces the expected marginal value, making firms more reluctant to create positions and hire workers. But in order for the wait-and-see effect to be quantitatively important, it is necessary that interior continuation values are sufficiently likely.²

Due to the cost of creating positions in my framework, a firm becomes less likely to destroy positions after a negative shock. It will rather keep the existing positions and workers, and it might even want to fill empty positions when workers quit. In these cases, the marginal value of a position takes an interior value. Then higher uncertainty can have significant negative effects on the expected marginal value of a position created today, and this strong wait-and-see effect can reduce the number of posted vacancies. As a result, the job-finding rate falls as observed in the data. The wait-and-see effect also becomes stronger for those firms that

²If the continuation value was at its upper bound after each positive productivity shock and at its lower bound after each negative productivity shock, increasing the variance of the productivity shock would not affect the expected value. The reason is that the continuation value would neither be strictly concave nor convex but flat at the realized values.

want to fire workers. This means that they reduce the number of workers that they fire which dampens the increase of the separation rate due to the realized volatility effect. As a result, the separation rate contributes less to changes in unemployment in accordance with the data.

In the absence of the creation cost, the job-finding rate hardly reacts to changes in uncertainty. In this case, higher uncertainty increases unemployment only because of a higher separation rate. When the cost is sufficiently large, however, the model can generate a significant drop of the job-finding rate in times of higher uncertainty. Then my model is not only able to explain the changes of the unemployment rate in the last decade, but it can also account for the steep drop in the job-finding rate during the Great Recession.

My model is related to the literature on multi-worker firm models with decreasing returns to scale and idiosyncratic productivity shocks. Examples of random search models include Cooper, Haltiwanger, and Willis (2007), Elsby and Michaels (2013), and Fujita and Najajima (2014), whereas Schaal (2012) and Kaas and Kircher (2014) assume directed search. The difference in my model is that it is costly for firms to create positions, which is an important assumption to explain the falling job-finding rate in times of higher uncertainty. With the exception of Schaal (2012), these papers do not consider uncertainty shocks. Schaal (2012) shows that it is difficult to explain the changes of the unemployment rate during the Great Recession while matching observed labour productivity, when only aggregate productivity shocks are used. His model with uncertainty shocks can explain the increase of the unemployment rate better. But it cannot explain the reduction of the job-finding rate during the Great Recession which means that the increase of the unemployment rate is primarily driven by a higher separation rate. Because of the creation cost present in my framework, the job-finding rate falls in response to higher uncertainty, which fits the observed data better.

This chapter also contributes to the literature of solving heterogeneous agents models with aggregate shocks by extending the Krusell and Smith (1998) algorithm. In these models, aggregate outcomes can depend on the distribution of agents, which is an infinite-dimensional object. In my model, a firm's optimal choice depends on current and future values of market tightness, which is the ratio between vacancies and unemployed. Therefore, a law of motion for market tightness is needed. The distribution of firms over idiosyncratic productivity and firm size becomes more dispersed due to the creation cost. As a result, approximating the distribution of firms only using its first moment is not accurate for explaining the demand of firms for new hires conditional on aggregate productivity and uncertainty. Higher moments cannot capture the characteristics of the distribution well, either. Instead of adding many higher moments, I add the observed residual between estimated market tightness and its market clearing value to the state space.³ It captures the information that is lost by approximating the distribution with its first moment. As the residual is highly autocorrelated, it is useful in predicting future values of market tightness. Intuitively, if firms underestimate the current value of market tightness due to the omitted information, they also expect market tightness to be below its predicted values in the next periods.

In addition to reducing forecast errors, using the residual method also means that market clearing can easily be imposed in each period of a simulation: in contrast to higher moments of the distribution, the residual is not predetermined at the beginning of the period. Demand

³Predicting the demand of firms for new hires is equivalent to predicting market tightness because the number of (un)employed is in the state space. It is important for firms to accurately predict market tightness because it determines their cost of filling vacant positions.

of firms for new hires is downward sloping in the residual, because a higher residual means a higher market clearing value of market tightness by construction. This allows solving for the residual in each period such that aggregate demand of new hires equals aggregate supply exactly.

The chapter is structured as follows: the next section provides evidence on the effects of uncertainty shocks on the job-finding rate and the separation rate in the US. Section 1.3 describes the model and derives the optimal choice of firms. Section 1.4 describes the calibration that is used. In section 1.5, I show the effects of uncertainty shocks in my model. I analyze how good the model can fit the US data from 1998-2013 in section 1.6. Section 1.7 describes the algorithm that I develop to solve the model. Section 1.8 concludes.

1.2 Effects of uncertainty in US data

In this section, I construct a measure for uncertainty based on the implied volatility of US stock options. I then use it to estimate impulse response functions of labour market transition rates following uncertainty shocks. In the theoretical model, uncertainty refers to the second moment of productivity on the firm level. Therefore, I want to construct a measure based on the volatility of *individual* firms as opposed to the volatility of an index of firms.^{4,5}

I use data on implied volatility of US stocks, available from 1996 to the second quarter of 2013. For each firm, I calculate the quarterly average of log implied volatility of its call options with a maturity of 30 days. Allowing for firm fixed effects, I estimate time dummy variables for each quarter. The resulting time-series is shown in figure 1.1.

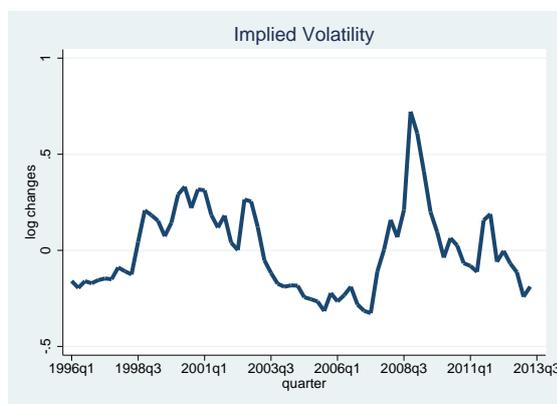


Figure 1.1: Index of implied volatility constructed from US stock options.

The constructed measure for uncertainty is used to estimate a VAR including volatility, labour productivity, and quarterly averages of the job-finding rate and separation rate. The following Cholesky ordering is used to identify volatility shocks: I assume that labour productivity does not react instantaneously to volatility shocks and shocks to the labour market transition rates. Furthermore, volatility may only react to shocks in the transition rates with a

⁴For example, the VIX measures the implied volatility of the S&P 500 index. To the extent that idiosyncratic firm shocks are uncorrelated, higher volatility at the firm level does not translate into higher volatility of the index because of diversification.

⁵Leahy and Whited (1996) and Bloom, Bond, and Van Reenen (2007) use the volatility of stock returns as a measure of uncertainty. An advantage of using the implied volatility of options is that it is a forward-looking measure, but it is not available as far back in time.

one period lag.⁶

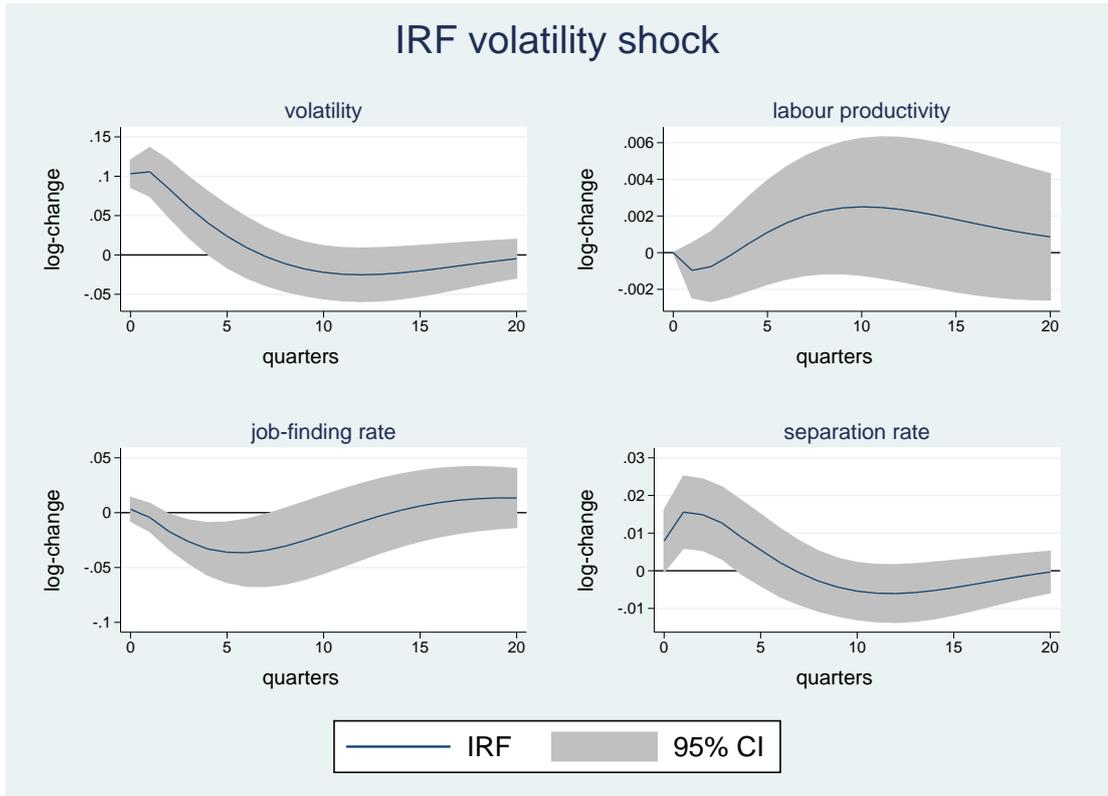


Figure 1.2: Impulse response functions after a volatility shock in US data.

Figure 1.2 shows the impulse response functions after a one standard deviation shock to volatility. The job-finding rate falls and the separation rate increases. Both effects contribute to an increase in unemployment following an increase in volatility. Note that the response of the separation rate peaks earlier than the job-finding rate, whose fall is more persistent. Its relative change is initially stronger but after a couple of quarters the job-finding rate changes more. These relative changes determine how much the unemployment rate is affected by the shock. As shown in appendix 1.A.2, the following approximation holds:

$$\log \frac{u}{\bar{u}} \approx (1 - \bar{u}) \left(\log \frac{s}{\bar{s}} - \log \frac{f}{\bar{f}} \right). \quad (1.1)$$

This means that a 1% increase of the separation rate affects the unemployment rate by as much as a 1% decrease of the job-finding rate. Intuitively, this holds because flows into and out of unemployment are equal on average. For the changes of the job-finding rate and separation rate shown in figure 1.2 this means that the peak response of the unemployment rate is slightly less than 40% times the increase in uncertainty.

Table 1.1 lists the forecast-error variance decomposition of the job-finding rate and the separation rate. The contribution of uncertainty shocks to the job-finding rate is comparable to productivity shocks for horizons of 8 or more quarters. For the separation rate, uncertainty shocks contribute about 3 times as much as productivity shocks, except for the first quarter.

⁶Appendix 1.A.3 shows that the results are robust to changing the Cholesky order, as well as using a different measure for uncertainty, and different detrending of the data.

Forecast Horizon (quarters)	Job-finding Rate		Separation Rate	
	Uncertainty	Productivity	Uncertainty	Productivity
1	0.5%	2.7%	4.8%	5.2%
4	10.7%	18.4%	26.0%	9.4%
8	24.8%	24.4%	28.1%	9.3%
12	27.3%	29.6%	28.2%	9.0%

Table 1.1: Forecast error variance decomposition of job-finding rate and separation rate due to uncertainty and productivity shocks.

The negative effect of higher uncertainty on the job-finding rate that I find is in line with other results in the literature. Recent empirical studies by Guglielminetti (2015) and Mecikovsky and Meier (2015) use different measures for uncertainty and find similar effects. Guglielminetti (2015) uses a survey-based measure of uncertainty. In a trivariate VAR, she finds that the job-finding rate falls after an increase in uncertainty. Mecikovsky and Meier (2015) use the measure of macroeconomic uncertainty in the US proposed by Jurado, Ludvigson, and Ng (2015). They find that job creation falls, whereas job destruction increases in response to higher uncertainty.

1.3 Model

This section describes the framework and derives the equations that are needed to solve the model.

1.3.1 Setup

There is a constant unit mass of workers, who search for jobs when they are unemployed and earn wages when they are employed. They are assumed to be risk-neutral. Workers can lose their job for two reasons. First, they can be fired when their employer experiences a sufficiently negative productivity shock that makes the match unprofitable. Second, it is assumed that not all separations are due to low productivity of the job. With a certain exogenous probability workers quit their job at the end of each period. In both cases, workers that become unemployed enter the unemployment pool in the next period, and can potentially find a new job.

I consider a constant mass of firms that can each hire multiple workers.⁷ Firms post vacancies and matching takes place according to a constant returns to scale matching function. The matching function can be fully characterised by the probability $q(\theta)$ that a vacancy is filled, where labour market tightness θ is the ratio $\frac{v}{u}$ of posted vacancies and unemployed workers. There are two types of costs that a firm incurs when it wants to increase its size. First, it has to pay a cost H for each position that it creates. Second, it has to pay a cost c for each vacancy it posts in order to fill empty positions. The firm takes labour market tightness and the implied probability of filling each vacancy as given.

In each period, a share δ of positions becomes obsolete. In this case also the workers who filled the positions, leave the firm. In addition, a share λ of workers quit without rendering

⁷There is no firm entry and exit in my model. This simplification can be justified by the low cyclicality of job creation at start-ups as documented by Coles and Kelishomi (2011). Similarly, Fujita and Najajima (2014) argue that cyclical fluctuations of job flows are mostly accounted for by the expansion or contraction of existing establishments.

the respective positions obsolete. In this case, the firm can hire another worker for a now vacant position and it only incurs the hiring cost $\frac{c}{q(\theta)}$ but not the creation cost H .⁸ A firm can also endogenously fire workers and shut down their positions without any cost, if it wants to reduce its size.⁹

For simplicity, I assume that a law of large numbers holds when the firm posts vacancies. That is the firm can fill each empty position for sure by posting $\frac{1}{q(\theta)}$ vacancies. Therefore, the cost of filling an empty position is $\frac{c}{q(\theta)}$.¹⁰ Firms are not allowed to keep an empty position idle. If it does not want to fill an empty position, it has to close it, and pay the cost H again if it wants to reopen it in the future. Even though firms cannot mothball positions, the distinction between positions and workers makes the model more interesting. Now each firm has to decide at the beginning of the period whether it wants to fill its empty positions. Firms that are relatively productive will decide to fill these positions whereas other firms will close at least some of them. As a result, some firms hire workers, while not creating new positions or even shrinking due to obsolescence. This is in line with the empirical results of Davis, Faberman, and Haltiwanger (2012). They find for shrinking establishments that hires amount to around 10% of employment on a quarterly basis.

The production function takes the form $zxF(n)$, where z is the aggregate productivity and x the idiosyncratic productivity component. Both follow a Markov process. The total wage bill is denoted by W . The wage equation is specified in equation (1.18). In general, it depends on both firm-specific and aggregate variables. The total wage bill is increasing in the number of workers and in firm productivity. In addition, it is increasing in aggregate productivity and market tightness. These are typical properties of wage equations in search and matching models, when workers and firms bargain over the surplus of their match. First, workers get a higher wage when the firm is more productive. Second, they get a higher wage when market tightness is high, as it would be easier for a worker to find a new job when he becomes unemployed, whereas it is more difficult for the firm to replace the worker.

The timing in each period is as follows: at the beginning of the period, a firm with n_{-1} positions has $(1 - \lambda)n_{-1}$ workers that fill them. The remaining λn_{-1} positions are empty. Then, the shocks of the aggregate and idiosyncratic productivity processes materialise. Afterwards, the firm has to decide, whether it wants to create new positions. In addition, it has to decide

⁸Fujita and Ramey (2007) also assume a one-off cost of creating positions, which has to be paid before posting vacancies. After separations that are not due to obsolescence, empty positions can be filled again. The important difference with my framework is that firms are not subject to idiosyncratic productivity shocks and operate at constant returns to scale in Fujita and Ramey (2007). As a result, the value of each empty position always equals its cost of creation, which makes the model very tractable. But the lack of meaningful firm heterogeneity would not allow to analyze the effects of increases in idiosyncratic volatility. In contrast to Fujita and Ramey (2007), I do not assume that the cost of creating positions increases in the total number of positions created.

⁹I assume throughout the chapter that a cost is incurred when a position is created. It is a minor change to the model to let firms pay a cost when it closes positions instead. The results when a destruction cost is used instead of the creation cost, are very similar for the following reason. When a firm opens a position, it knows that with certain probabilities it will destroy it in future periods when it becomes sufficiently unproductive. Therefore, when there is a destruction cost instead of the creation cost, the firm takes into account the net present value of this future destruction cost upon creating positions. The resulting optimal creation of positions becomes similar even though the firm only incurs the cost when it destroys the position. Figure 1.22 in appendix 1.B.1 shows the results of the comparative statics exercise when a destruction cost is used.

¹⁰In an earlier version of this chapter, I assumed that vacant positions only get filled gradually. This means that both the number of filled and the number of empty positions become state variables of each firm. Adding another state variable makes the numerical solution of the firm's problem more costly. The impact of assuming that firms can fill positions immediately on the results are likely to be small for the following reasons. First, it does not take firms long to fill vacancies on average. For example, den Haan, Ramey, and Watson (2000) use a quarterly job-filling rate of 0.71, which is in line with the filling probability found by van Ours and Ridder (1992) for establishments in the Netherlands. Second, a firm that creates new positions is at an interior solution to its optimization problem. Therefore, small deviations from its optimal size do not have large effects on its value.

how many of its empty positions it wants to fill.¹¹ The firm shuts down all vacant positions that it does not want to fill. In the next step, the matching process takes place. Afterwards, the firm can fire workers and close their positions.¹² Then production takes place. The number of workers in the production phase is denoted by n^* . After production, a fraction δ of positions becomes obsolete and the respective workers become unemployed. Finally, a fraction λ of workers quits. Hence, the firm begins the next period with $(1 - \delta) n^*$ positions, filled with $(1 - \lambda) (1 - \delta) n^*$ workers.

Let s denote the aggregate state. It summarizes the current realization of the aggregate shocks as well as the distribution of positions and idiosyncratic productivity across firms. Let $V(n_{-1}, x, s)$ denote the value of a firm with n_{-1} positions at the beginning of the period and idiosyncratic productivity x , when the aggregate state is s .¹³ It can be written recursively as

$$V(n_{-1}, x, s) = \max_n zx F(n) - W(n, x, s) - (n - (1 - \lambda) n_{-1})_+ \frac{c}{q(\theta(s))} - (n - n_{-1})_+ H + \dots + \beta \mathbb{E} [V((1 - \delta) n, x', s') | x, s], \quad (1.2)$$

where I use the shorthand notation $(y)_+ := \max\{y, 0\}$. The firm's per period profit is given by the difference between production and the total wage bill reduced by the cost of filling $(n - (1 - \lambda) n_{-1})_+$ positions and the cost of creating $(n - n_{-1})_+$ positions. The continuation value is given by the expected discounted value of a firm with $(1 - \delta) n$ positions at the beginning of next period.

Note that there are two kinks in the objective function that the firm wants to maximise. The first kink appears when $n = (1 - \lambda) n_{-1}$. Then the firm does not fill any of its empty positions, but it also does not fire any of its existing workers. The second kink occurs when $n = n_{-1}$. In this case, the size of the firm stays exactly the same within the period. It fills all its existing empty positions, but it does not create any new positions. The firm's optimal choice could be either at one of these two kinks or in one of the following three regions, into which the two kinks separate the possible number of positions. First, the optimal n^* could be below the first kink at $(1 - \lambda) n_{-1}$. This means that the firm fires existing workers. Second, n^* could be between the two kinks. In this case, the firm fills some but not all of its empty positions, and it does not create any new positions. Third, if n^* is above the second kink at n_{-1} , the firm creates new positions and posts vacancies to fill all empty positions. In the next subsection, I will distinguish these five cases for the optimal choice of the firm. I derive an equation for the optimal choice of positions n^* and the marginal value of a position in each case.

Figure 1.3 shows an example for the resulting marginal value function. The maximum value of an additional position at the beginning of the period is given by the cost the firm saves by not having to create this position and not having to hire the respective worker filling it. The minimum value for the marginal value of a position is 0, because the firm can always destroy it and fire the worker without cost. These upper and lower bounds play an important role when uncertainty changes. In particular, the upper bound only exists, because firms have the option to hire additional workers in my framework. The result is a concave part of the

¹¹It is always optimal for a firm that creates new positions to also fill its empty positions.

¹²Firing takes place after the matching process such that only workers who were unemployed at the beginning of the period can be matched in that period. When I calibrate the model, one period corresponds to one week, such that this assumption is quantitatively not important.

¹³A firm with n_{-1} positions at the beginning of the period employs $(1 - \lambda) n_{-1}$ workers, because all positions were filled in the last production phase and a share λ of workers has quit afterwards.

marginal value function. Then, higher uncertainty can potentially reduce the expected future marginal value.

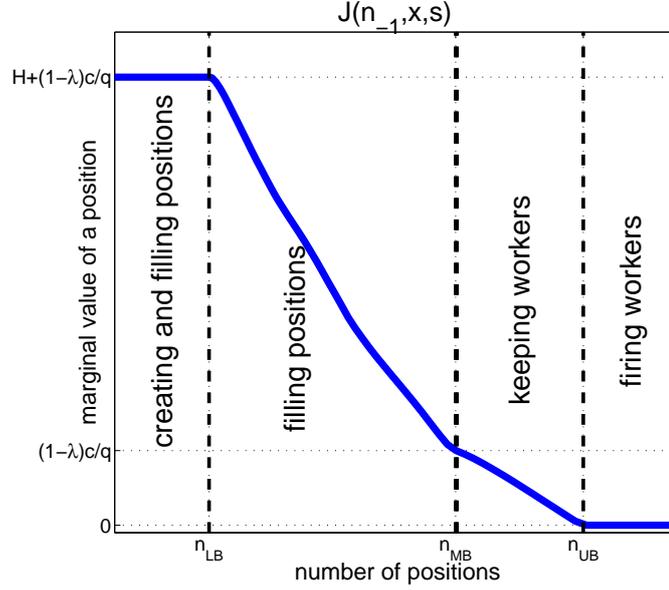


Figure 1.3: Marginal value of a position (including $1 - \lambda$ workers) for a given level of idiosyncratic productivity x and aggregate state s .

1.3.2 Solution to firm problem

Let $J(n_{-1}, x, s) := \frac{\partial V(n_{-1}, x, s)}{\partial n_{-1}}$ denote the marginal value of an additional position at the beginning of the period. To keep notation simpler, let $\Pi(n, x, s) := zx F(n) - W(n, x, s)$ denote the firm's revenue minus the total wage bill. Because of the kinks in the hiring costs, one has to distinguish five possible outcomes, when characterising the optimal employment decision n^* and the marginal value of an additional position:

1. $n^* < (1 - \lambda) n_{-1}$. The firm fires workers. In this case, the number of positions is reduced until marginal profit in the current period and the expected future marginal profit sum to zero. The optimal number of positions n^* is then implicitly given by

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] = 0, \quad (1.3)$$

$$J(n_{-1}, x, s) = 0. \quad (1.4)$$

In this case the marginal value of a position J reaches its lower bound 0.

2. $n^* = (1 - \lambda) n_{-1}$. The firm keeps all its existing workers, but does not fill any empty positions. In this case, the marginal value of a worker must be positive but less than the hiring cost $\frac{c}{q}$ so that the firm neither wants to fire workers nor fill empty positions. The marginal value of beginning the period with an additional position created in the past, being filled with $(1 - \lambda)$ workers, is given by the sum of current marginal profit of these workers and the expected future marginal value of additional $(1 - \lambda)(1 - \delta)$ positions

in the next period:

$$n^* = (1 - \lambda) n_{-1}, \quad (1.5)$$

$$J(n_{-1}, x, s) = (1 - \lambda) \left\{ \frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] \right\} \quad (1.6)$$

$$\in \left(0, (1 - \lambda) \frac{c}{q(\theta(s))} \right).$$

In this case the firm is inactive with respect to both creating and filling positions. As long as it stays in this inaction region, its size is reduced by the factor $(1 - \lambda)(1 - \delta)$ each period because of exogenous quits of workers and obsolescence of positions.

3. $n^* \in ((1 - \lambda)n_{-1}, n_{-1})$. The firm fills some of its empty positions and shuts down the others. The optimal number of positions is determined by the condition that the marginal current and future expected profit of an additional worker equals its hiring cost $\frac{c}{q}$. The marginal value of an additional position at the beginning of the period is equal to the saved cost of hiring $(1 - \lambda)$ workers:

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] = \frac{c}{q(\theta(s))}, \quad (1.7)$$

$$J(n_{-1}, x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (1.8)$$

This means that the marginal values of a position is constant between $(1 - \lambda)n_{-1}$ and n_{-1} . The reason is that the firm does not want to fill all its empty positions. Then, the marginal value of an additional position at the beginning of the period only stems from the $(1 - \lambda)$ workers that it is filled with. They reduce the hiring cost by $(1 - \lambda) \frac{c}{q}$.

4. $n^* = n_{-1}$. The firm keeps and fills all existing empty positions, but it does not create new positions. In this case, the marginal value of a position must exceed the hiring cost for filling but be less than $H + (1 - \lambda) \frac{c}{q}$ so that it lies below the cost of creating and filling positions. The marginal value of an additional position created in the past is given by the sum of current marginal profit of the marginal worker (including the filling cost for the empty positions) and the expected future marginal value of additional $(1 - \delta)$ positions in the next period:

$$n^* = n_{-1}, \quad (1.9)$$

$$J(n_{-1}, x, s) = \frac{\partial \Pi}{\partial n}(n^*, x, s) - \lambda \frac{c}{q(\theta(s))} + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] \quad (1.10)$$

$$\in \left((1 - \lambda) \frac{c}{q(\theta(s))}, H + (1 - \lambda) \frac{c}{q(\theta(s))} \right).$$

In this case, the firm is inactive only with respect to creating positions. As long as it stays in this inaction region, its size is reduced by a factor $(1 - \delta)$ each period because of obsolescence of positions.

5. $n^* > n_{-1}$. The firm opens new positions and fills all empty positions. Its optimal size is determined by the condition that the cost of creating and filling an additional position equals its marginal current and expected future marginal profit. The marginal value of a

position created in the past, filled with $(1 - \lambda)$ workers, is the sum of the saved creation cost and filling cost:

$$\frac{\partial \Pi}{\partial n} (n^*, x, s) + \beta (1 - \delta) \mathbb{E} [J((1 - \delta)n^*, x', s') | x, s] = H + \frac{c}{q(\theta(s))}, \quad (1.11)$$

$$J(n_{-1}, x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (1.12)$$

Note that in this case the marginal value of a position J reaches its maximum value $H + (1 - \lambda) \frac{c}{q}$. This upper bound is important in explaining why the job-finding rate can fall in response to higher uncertainty.

In summary, the marginal value function has to satisfy

$$J(n_{-1}, x, s) = \min \left\{ H + \frac{(1 - \lambda)c}{q(\theta(s))}, \max \left\{ J_F(n_{-1}, x, s), \min \left\{ \frac{(1 - \lambda)c}{q(\theta(s))}, \max \{ J_K(n_{-1}, x, s), 0 \} \right\} \right\} \right\}. \quad (1.13)$$

where J_K and J_F denote the right hand side values of equations (1.6) and (1.10): J_K is the marginal value that results, if the firm keeps all its workers but is inactive with respect to filling positions ($n^* = (1 - \lambda)n_{-1}$), and J_F is the marginal value, if the firm fills all empty positions but is inactive with respect to creating new positions ($n^* = n_{-1}$).

Note that the third case, in which some but not all of the empty positions are refilled, is relatively unimportant when the length of the period is chosen to be short. In the limit of a continuous time variant, the interval $((1 - \lambda)n_{-1}, n_{-1})$ collapses, and one would only distinguish between firms that fire workers, those that do not fill empty positions and thus shrink at rate $\lambda + \delta$, those that refill empty positions and shrink at rate δ , and those that create new positions.

1.3.3 Characterisation of the optimal firm behaviour using three cutoffs

Given the marginal value function derived above, the optimal behaviour of a firm can be summarised by three cutoffs, that determine how many positions firms want to create, how many positions they want to fill, and how many workers they want to fire. These cutoffs are functions of idiosyncratic productivity x and the aggregate state s . If the number of positions at the beginning of the period is small, a firm will create positions until the number of positions is equal to the lower bound $n_{LB}(x, s)$. This cutoff is implicitly given by the condition that the marginal value J_F of a firm that fills all position is exactly equal to the upper bound of the marginal value function:

$$J_F(n_{LB}(x, s), x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (1.14)$$

If the initial number of positions lies above the lower cutoff, the firm will be inactive with respect to creating new positions but it will fill all empty positions as long as the marginal value J_F is greater than $(1 - \lambda) \frac{c}{q}$. This condition implicitly defines the middle cutoff $n_{MB}(x, s)$:

$$J_F(n_{MB}(x, s), x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (1.15)$$

The third cutoff, $n_{UB}(x, s)$, determines the upper bound of positions in the firm. The firm does not fire existing workers, as long as the marginal value J_K is positive. The upper bound

is implicitly defined by

$$J_K(n_{UB}(x, s), x, s) = 0. \quad (1.16)$$

With these three cutoffs, the optimal firm behaviour can be summarised as follows: if a firm with productivity x has fewer than n_{LB} positions at the beginning of the period, it creates the difference $n_{LB} - n_{-1}$ and fills all of them. If its initial positions are between n_{LB} and n_{MB} it is inactive with respect to the number of positions and it fills the λn_{-1} empty positions it initially has. If the number of initial positions lies between n_{MB} and $\frac{n_{MB}}{1-\lambda}$, it fills its empty positions only up to n_{MB} and shuts down the remaining ones. If the number of initial positions lies between $\frac{n_{MB}}{1-\lambda}$ and n_{UB} , the firm closes all open positions but keeps its workers. This means that the number of positions is reduced to $(1-\lambda)n_{-1}$. If the firm initially has more positions, it reduces its size to $(1-\lambda)n_{UB}$, which also involves firing workers. Table 1.2 summarises the optimal behaviour of a firm and figure 1.4 illustrates the optimal choice of positions n^* given positions at the beginning of the period n_{-1} .¹⁴

n_{-1}	n^*	workers hired	workers fired
$< n_{LB}$	n_{LB}	$n_{LB} - (1-\lambda)n_{-1}$	0
$\in [n_{LB}, n_{MB})$	n_{-1}	λn_{-1}	0
$\in [n_{MB}, \frac{n_{MB}}{1-\lambda})$	n_{MB}	$n_{MB} - (1-\lambda)n_{-1}$	0
$\in [\frac{n_{MB}}{1-\lambda}, n_{UB})$	$(1-\lambda)n_{-1}$	0	0
$\geq n_{UB}$	$(1-\lambda)n_{UB}$	0	$(1-\lambda)(n_{-1} - n_{UB})$

Table 1.2: Summary of firm's optimal choice of positions n^* , and the number of workers hired and fired, depending on number of positions at the beginning of the period n_{-1} .

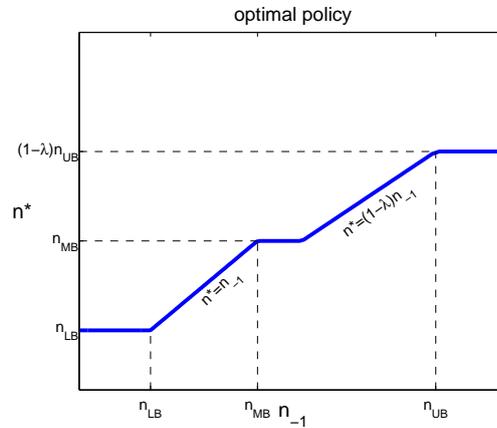


Figure 1.4: Optimal choice of positions for a given level of idiosyncratic productivity.

The cutoffs are increasing in x as higher idiosyncratic productivity makes the marginal position more valuable in a given aggregate state. Figure 1.5 provides an example of the three cutoffs and the optimal decision of firms.

¹⁴Note that this describes the optimal choice of positions within a period. Between periods, a fraction δ of positions is destroyed. Therefore, the number of positions at the beginning of the next period is $(1-\delta)n^*$.

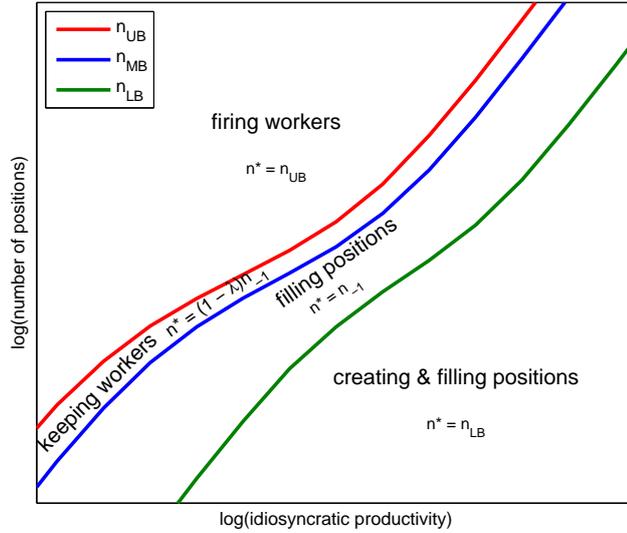


Figure 1.5: Cutoffs $n_{LB}(x)$, $n_{MB}(x)$, and $n_{UB}(x)$ for a given aggregate state.

1.3.4 Market clearing

In equilibrium, the following market clearing condition has to hold: the mass of workers that firms want to hire must be equal to the number of newly matched unemployed according to the matching function. Let $F_s(n_{i,-1}, x_i)$ denote the distribution of positions and idiosyncratic productivity across firms in the aggregate state s . Then

$$\int_{n_{i,-1}} \int_{x_i} (n^*(n_{i,-1}, x_i, s) - (1 - \lambda)n_{i,-1})_+ dF_s(n_{i,-1}, x_i) = M(u(s), v(s)) \quad (1.17)$$

must hold in all possible states s , where $M(u, v)$ is the number of matches when the mass of unemployed is u and the mass of posted vacancies is v . It is important to ensure that this condition is satisfied when the economy is simulated in the numerical solution. In section 1.7, I describe the algorithm that I develop to ensure market clearing exactly.¹⁵

1.4 Parametrization

The model is parametrized to US data. In section 1.6, I compare the results from the model with the corresponding time series in the US. Because the probability for unemployed to find a job within a month is close to 50%, I choose a high frequency for my model. This allows unemployed to find a job in less than a month. In particular, I use a weekly parametrization like Elsby and Michaels (2013) with a discount factor $\beta = 0.999$.¹⁶

I calibrate the model for various values of the cost of creating positions H . In the next sections, I will compare the results for those different values of H . Most of the parameters are held constant across calibrations, and they are summarised in table 1.3. There are three parameters that I recalibrate in order to meet the following three targets across calibrations: in the absence

¹⁵This is important because differences between aggregate supply and aggregate demand might accumulate over time in a simulation.

¹⁶This is equivalent to a yearly interest rate of 5.3%.

of aggregate shocks, the unemployment rate should be 6.5%, the job-finding rate should be 11.25%, and the cost of filling a vacancy should be 14% of the average quarterly wage.¹⁷ In order to meet these targets, I choose ω_3 in the wage equation, the standard deviation $\bar{\sigma}_x$ of the idiosyncratic shock, and the mass of firms, as described below. The resulting parameters for different values of H are listed in table 1.4.

Parameter	Value	Description
β	0.999	Discount factor
η	0.6	Elasticity of matching function ($M(u, v) = \mu u^\eta v^{1-\eta}$)
μ	0.128	Scale of matching function (job-finding rate $a(0.72) = 0.1125$)
c	0.156	Vacancy posting cost: normalised such that $\frac{c}{q(0.72)} = 1$
ω_1	0.8	Firm's share of production
ω_2	0.5c	Part of the wage that is proportional to market tightness
ω_3	depending on H	Constant in wage equation
δ	0.001	Depreciation of positions
λ	0.003	Exogenous worker turnover
α	0.75	Elasticity of production function
H	0 – 10	Cost of creating a position (relative to filling it)
$mass$	depending on H	Mass of firms
λ_x	0.043	Arrival rate of idiosyncratic shock (Elsby and Michaels (2013))
$\bar{\sigma}_x$	depending on H	Standard deviation of idiosyncratic shock
σ_σ	0.08	Unconditional standard deviation of $\log \bar{\sigma}_x$
ρ_σ	0.988	Autocorrelation of $\log \bar{\sigma}_x$
σ_z	0.02	Unconditional standard deviation of $\log z$
ρ_z	0.996	Autocorrelation of $\log z$

Table 1.3: Model parameters based on a weekly calibration. The model is solved for different values of creation cost H . The parameters that are recalibrated depending on H can be found in table 1.4.

H	0	2	5	10	Cost of creating a position
Calibrated parameters					
ω_3	0.354	0.350	0.346	0.338	Constant in wage equation
$\bar{\sigma}_x$	0.131	0.231	0.351	0.502	Standard deviation of idiosyncratic shock
$mass$	0.197	0.182	0.156	0.116	Mass of firms
Not-targeted statistics without aggregate risk					
	0%	22%	52%	99%	Creation cost relative to quarterly output
	0%	1.8%	4.1%	7.4%	Aggregate creation costs relative to output
	1.2%	1.2%	1.1%	1.1%	Aggregate filling costs relative to output
	0.79	0.77	0.74	0.71	Labour share
	0.40	0.43	0.41	0.37	Standard deviation of annual employment growth

Table 1.4: List of parameters that are recalibrated depending on H to keep the unemployment rate, job-finding rate, and filling cost relative to average wage the same across calibrations. The lower part reports some statistics of the model without aggregate risk. The aggregate expenditures on filling costs c/q and on creation costs H , both relative to output. The last row reports the cross-sectional standard deviation of yearly employment growth calculated as $std\left(\frac{n_{i,t} - n_{i,t-52}}{0.5(n_{i,t} + n_{i,t-52})}\right)$.

Matching function and vacancy costs. The matching function is assumed to take the Cobb-Douglas form $M(u, v) = \mu u^\eta v^{1-\eta}$. The elasticity η is set to 0.6 as in Elsby and Michaels (2013). This implies that the job-finding probability for an unemployed as a function of market tightness is $a(\theta) = \mu\theta^{1-\eta}$. The probability of a vacancy being filled is $q(\theta) = \mu\theta^{-\eta}$. Without loss of generality, I target market tightness to be equal to 0.72 in the absence of aggregate

¹⁷These targets are also used by Elsby and Michaels (2013).

shocks.¹⁸ Then the scale of the matching function $\mu = 0.128$ is implied by the targeted job-finding probability $a(0.72) = 0.1125$. Vacancy posting cost c is chosen to normalise the cost of filling positions to 1 ($c = q(0.72)$). Then the creation cost H can be interpreted as a multiple of the filling cost in steady state.

Production function. The production function takes the form $F(n) = n^\alpha$. I set $\alpha = 0.75$ in between the values used by Elsby and Michaels (2013) and Schaal (2012). Creation costs could be interpreted as the value of capital that is lost if a position is shut down and the respective share of capital is sold. Therefore, α should not be interpreted as just the curvature of the production function when increasing the number of workers, holding capital constant. More precisely, $F(n)$ is the revenue function rather than the amount of goods produced by the firm. A downward sloping demand curve, for example due to monopolistic competition, reduces α relative to the curvature in the pure production function.

Wage equation. I impose the following functional form for the total wage bill:¹⁹

$$W(n, x, s) = (1 - \omega_1)zxn^\alpha + (\omega_2\theta + \omega_3)n. \quad (1.18)$$

This form is common in search and matching models: workers get a share of the marginal product and a payment that compensates them according to their outside option. This depends on the value of home production and on the option value of finding another job. The latter is proportional to $a(\theta) \frac{c}{q(\theta)} = c\theta$, where a denotes the probability of finding a job. This is because each firm posts vacancies (if any) until the marginal benefit equals the hiring cost $\frac{c}{q}$.²⁰

The firm's share of the marginal product is set to $\omega_1 = 0.8$. It is not crucial that the wage per worker differs across firms, i.e. that ω_1 is less than 1. More generally, it acts in the same way as a change in average productivity. Therefore, it scales all firms' sizes, but it does not affect the dynamics of positions or vacancies.²¹ A result of the relatively high value of ω_1 is that wage dispersion between firms is lower. This avoids that a worker would want to leave an unproductive firm voluntarily to get a potentially higher wage when finding a new job.

¹⁸Note that the scale of the matching function μ and the vacancy posting cost c can be adjusted in response to a proportional change in market tightness such that $a(\theta)$ and $c\theta$ remain unchanged. Then, the job-finding rate remains the same and the firm's problem is not affected because the filling cost $\frac{c}{q(\theta)} = \frac{c\theta}{a(\theta)}$ and the term $c\theta$ in the wage equation are unchanged.

¹⁹The total wage bill is specified as opposed to the wage per worker to simplify notation. The derivative of W with respect to n gives the effect of hiring an additional worker on a firm's wage bill. This differs from the wage per worker because the wage per worker is decreasing in firm size. Therefore, hiring an additional worker reduces the wage a firm has to pay to its existing workers, dampening the increase of the total wage bill.

²⁰This form can be derived by imposing Stole and Zwiebel (1996) bargaining in the absence of creation costs for positions ($H = 0$) like in Elsby and Michaels (2013). There is no analogous analytical solution to the bargaining problem in my full model when $H > 0$. This is because some workers are hired at firms that newly create positions, and others are hired to fill positions that have become empty due to exogenous worker turn-over. The surplus of a match is not equal across these different positions. Therefore, one needs to know the distribution of vacancies across firms, if one wants to calculate the expected value of an unemployed, for instance. As a result there is no analytical solution to the bargaining problem. To make the bargaining problem simpler, it could be assumed that firms have all the bargaining power. Then, the value for a newly employed worker does not exceed the value of an unemployed. This approach was used for example in Cooper, Haltiwanger, and Willis (2007) and in Postel-Vinay and Robin (2002). In my framework, it would imply that all workers are paid the same wage equal to the flow benefits of an unemployed.

²¹The reason is as follows. A firm's production less wages is equal to $\omega_1zxn^\alpha - (\omega_2\theta + \omega_3)n$. From a firm's perspective, changing only ω_1 to be 1 is equivalent to a permanent change in productivity. The firm optimally responds by multiplying its number of positions by the factor $\omega_1^{-\frac{1}{1-\alpha}}$. Then all first order conditions hold exactly as before. This means that all firms become proportionally larger which is perfectly offset by multiplying the mass of firms by $\omega_1^{\frac{1}{1-\alpha}}$. Then the employment dynamics of the model are identical regardless of the value ω_1 even with aggregate risk. The average wage, however, is affected by a different choice of ω_1 . As cost of filling relative to the average wage is targeted in the calibration, the parameter ω_3 has to be recalibrated when ω_1 is changed. Therefore, the results of the calibrated model do depend on the choice for ω_1 . Figure 1.27 in appendix 1.B.1 shows that the main results are robust to choosing $\omega_1 = 1$.

The parameter $\omega_2 = 0.5c$ is set to the respective values derived from Stole and Zwiebel (1996) bargaining when $H = 0$ and the worker's bargaining power is 0.5. The constant ω_3 is calibrated depending on H such that the calibration targets are met.

Attrition. I choose the depreciation rate of positions δ to be 0.001 on a weekly basis, which corresponds to 5% on a yearly basis. The choice of λ determines the exogenous rate of workers leaving a firm. It should represent the rate at which employees quit for reasons other than low productivity of the firm. Thus it also determines the share of vacancies that are posted by firms to replace these workers relative to vacancies posted to fill newly created positions. Given that the total separation rate is implied by the calibrated values for unemployment and the job-finding rate, λ also determines how much job destruction is exogenous and endogenous, respectively. The probability of a worker leaving for exogenous reasons is approximately $\lambda + \delta$. The average separation rate in the model is 0.008. I set $\lambda = 0.003$ such that about half of separations are exogenous in the absence of aggregate shocks. This corresponds to quits being approximately half of total separations in the US.²²

Idiosyncratic shocks. Like in Elsby and Michaels (2013), idiosyncratic shocks arrive at rate $\lambda_x = 0.043$. When a firm is hit by a shock, the newly drawn idiosyncratic productivity follows a log-normal distribution:

$$\log x_t^{new} \sim N\left(0, \sigma_{x,t}^2\right). \quad (1.19)$$

Standard deviation of idiosyncratic productivity in steady state ($\bar{\sigma}_x$) is recalibrated when H is varied. This is necessary because a higher cost of creation makes firms more reluctant to destroy positions and fire workers. Hence, endogenous job destruction falls. Table 1.4 reports the resulting standard deviation that is increasing in H . Note that this does not necessarily mean that the number of positions within a firm becomes more volatile. The last row of table 1.4 shows that for all values of H , the standard deviation of annual employment growth is close to the estimate of 0.416 calculated by Elsby and Michaels (2013) using data on continuing establishments in the Longitudinal Business Database from 1992 to 2005.²³

Uncertainty shocks. The standard deviation of idiosyncratic productivity follows a log-normal process:

$$\log \sigma_{x,t} - \log \bar{\sigma}_x = \rho_\sigma (\log \sigma_{x,t-1} - \log \bar{\sigma}_x) + \varepsilon_{\sigma,t}. \quad (1.20)$$

When idiosyncratic volatility changes, I also adjust aggregate productivity to compensate for the following effects. As idiosyncratic productivity follows a log-normal distribution, increases of $\sigma_{x,t}$ increase average productivity across firms.²⁴ In addition, if firms wanted to keep their marginal product the same, the specified production function implies that labour demand is a convex function of idiosyncratic productivity x .²⁵ Therefore, higher dispersion

²²Monthly quits are on average 52% of total separations from 2001-2013 in the JOLTS dataset produced by the BLS. Note that while the probability of an exogenous quit in my model is time-invariant, the share of quits in total separations is procyclical because the separation rate is countercyclical. It would be a simple extension of the model to make exogenous separations dependent on the aggregate productivity and uncertainty states.

²³The annual growth rate is calculated as $\frac{n_{i,t} - n_{i,t-52}}{0.5(n_{i,t} + n_{i,t-52})}$ as in Davis and Haltiwanger (1992).

²⁴The expected newly drawn level of productivity is given by $\mathbb{E}x_t^{new} = \exp\left(\frac{\sigma_{x,t}^2}{2}\right)$.

²⁵This can be seen in a simple static model, in which the marginal product of a worker equals the wage:

$$\alpha z x_i n_i^{\alpha-1} = w.$$

Then labour demand by firm i is given as

$$n_i = \left(\frac{\alpha z x_i}{w}\right)^{\frac{1}{1-\alpha}}$$

of productivity across firms increases aggregate demand for labour. I abstract from this effect by multiplying productivity with \tilde{z}_t , dependent on the cross-sectional distribution of idiosyncratic productivity, such that labour demand would not be affected by changes in uncertainty in a frictionless model:²⁶

$$\int_{n_{i,-1}} \int_{x_i} (\tilde{z}_t x_{i,t})^{\frac{1}{1-\alpha}} dF_s(n_{i,-1}, x_i) = \text{const.} \quad (1.21)$$

Aggregate productivity shocks. Aggregate productivity z follows a log-normal process that is independent from the volatility of idiosyncratic productivity:

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_{z,t}. \quad (1.22)$$

The weekly autocorrelation ρ_z is chosen to be 0.996, which corresponds to a quarterly autocorrelation of 0.95. The standard deviation of the innovation is chosen such that the unconditional standard deviation of log-productivity is 0.02.

Creation cost. I solve the model for various values of H ranging from 0 to 10. Table 1.4 shows that this means that the cost of creating a position is about $10 \cdot H\%$ of quarterly output. Thus, for the highest cost considered ($H = 10$) it is about as costly for firms to create one position as the average quarterly output of each position is. I also report the share of output that firms spend on posting vacancies and on creating positions in table 1.4. This filling cost is a bit more than 1% of output, and the aggregate creation cost ranges from 0 to 7.4% of output. Adding the creation costs, filling costs and the labour share in table 1.4, it can be seen that average firm profits are about 20% of output. They can be interpreted as the return on capital, which is absent in my model. Then, the implied capital per position is about 15 times quarterly output.²⁷ Hence, the largest creation costs $H = 10$, considered in this chapter, could be interpreted in the following way: when a position is created the firm has to pay less than 7% of the value of capital that is needed for the position, and this amount cannot be recovered when the position is destroyed and the capital is sold.²⁸

1.5 Effects of uncertainty in the model

In the first subsection, I study the long run effects of a permanent change in idiosyncratic volatility. It shows that the creation cost H is crucial for the response of the job-finding rate after an increase in uncertainty. When $H = 0$, the job-finding rate is hardly affected, but when $H = 10$, it falls significantly and this is responsible for almost half of the drop in employment.

²⁶Frictionless means that firms are always at the creation cutoff $n_{LB}(x)$. This would happen if it was possible to trade existing positions and workers between firms, until the marginal product at all firms equalises. For the numerical solution of the model, this also means that $m^x \equiv \int_i x_i^{\frac{1}{1-\alpha}} dF(i)$ becomes a state variable.

²⁷This is calculated using the quarterly discount factor, which implies the quarterly interest rate $\frac{\beta^{13}}{1-\beta^{13}}$, and a profit share of 20%. Then the implied ratio of capital and quarterly output is $0.2 \frac{\beta^{13}}{1-\beta^{13}}$.

²⁸The resale price of capital in case of partial irreversibility is usually calibrated to match some business cycle statistics. There is a wide range of values obtained. For example, Khan and Thomas (2013) use an investment resale loss of 4.6%, whereas Bloom (2009) has a resale loss of 33.9% in his full model. An alternative way to compare the creation cost used in this chapter with the capital adjustment cost literature is to calculate the aggregate amount of capital adjustment costs spent in the economy in equilibrium. Cooper, Haltiwanger, and Willis (2004) estimate capital disruption costs that are 15% of revenues minus labour costs. An analogous estimate of the aggregate creation costs in my model is 16% when $H = 5$.

In the second subsection, I decompose the effects of higher uncertainty into changes driven by the wait-and-see effect, the realized volatility effect, and the general equilibrium effect as market tightness changes. It demonstrates that the wait-and-see effect becomes stronger when H is large. This is important, because the wait-and-see effect ultimately reduces the job-finding rate. In subsection 1.5.3, I solve the full model with time-varying aggregate productivity and time-varying uncertainty. The impulse response functions after an uncertainty shock confirm that it is only in case of high creation cost that the job-finding rate falls significantly. I also show that the quantitative effect of an increase in uncertainty depends on the state of the economy. In particular, higher uncertainty increases unemployment by more when the initial unemployment rate was already high.

1.5.1 Comparative statics

In this section, I consider the long run effects of a permanent increase in idiosyncratic volatility in steady state. I demonstrate that the presence of creation costs is necessary to get a fall in the job-finding rate as found in the empirical part. Figure 1.6 shows the effect of a 10% increase in idiosyncratic volatility σ_x for increasing values of cost H .²⁹ This second moment shock is accompanied with a negative first moment shock, as described above, such that labour productivity would not be affected in a frictionless environment, in which the marginal product is equal at all firms.

In response to the shock, the separation rate reacts more strongly with higher creation costs. From the third subplot it becomes evident that without creation costs ($H = 0$), the job-finding rate hardly changes in response to higher uncertainty. When H becomes larger, the job finding rate falls significantly. Both, the higher separation rate and the lower job-finding rate imply a stronger fall in employment after a permanent increase in idiosyncratic volatility for high H . In the sixth subplot I decompose how much of the change in unemployment is due to the higher separation rate and how much is due to the lower job finding rate.³⁰

There are two opposing effects, that change labour productivity in this heterogeneous agents model. First, the wait-and-see effect increases the dispersion of firms conditional on idiosyncratic productivity, because they become more reluctant to create and to destroy positions. This makes the average product of labour more dispersed across firms which leads to a reduction of measured labour productivity in the aggregate.³¹ The second effect is that reduced average employment increases the average product of firms because they operate at decreasing returns to scale. This effect increases labour productivity. In the comparative statics exercise the latter effect dominates. In the full model, however, labour productivity can fall in response to an increase in uncertainty when H is large, as can be seen in figure 1.11.

If we want to explain the effect of uncertainty on the creation of positions, we need to

²⁹Note that due to the calibration strategy, idiosyncratic volatility is increasing in H to get the same level of endogenous job destruction. Therefore, a 10% increase, raises idiosyncratic volatility by more percentage points with higher creation costs, which is a reason for the stronger response of the unemployment rate. My focus in interpreting the results, however, is on the relative contribution of job-finding rate and separation rate. Figure 1.23 in appendix 1.B.1 shows the effects, when idiosyncratic volatility is increased by the same absolute amount. Then the response of employment is comparable across calibrations. Importantly, the contribution of the job-finding rate to the fall in employment is hardly affected by the alternative size of the shock.

³⁰I calculate the change in unemployment if only the job-finding rate was changed, Δu^{JFR} , and the change in unemployment if only the separation rate was changed, Δu^{SR} . Note that $\Delta u \approx \Delta u^{JFR} + \Delta u^{SR}$ does not hold exactly. I report the contribution of the job-finding rate as $\frac{\Delta u^{JFR}}{\Delta u^{JFR} + \Delta u^{SR}}$.

³¹This effect on measured TFP has also been discussed by Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014).

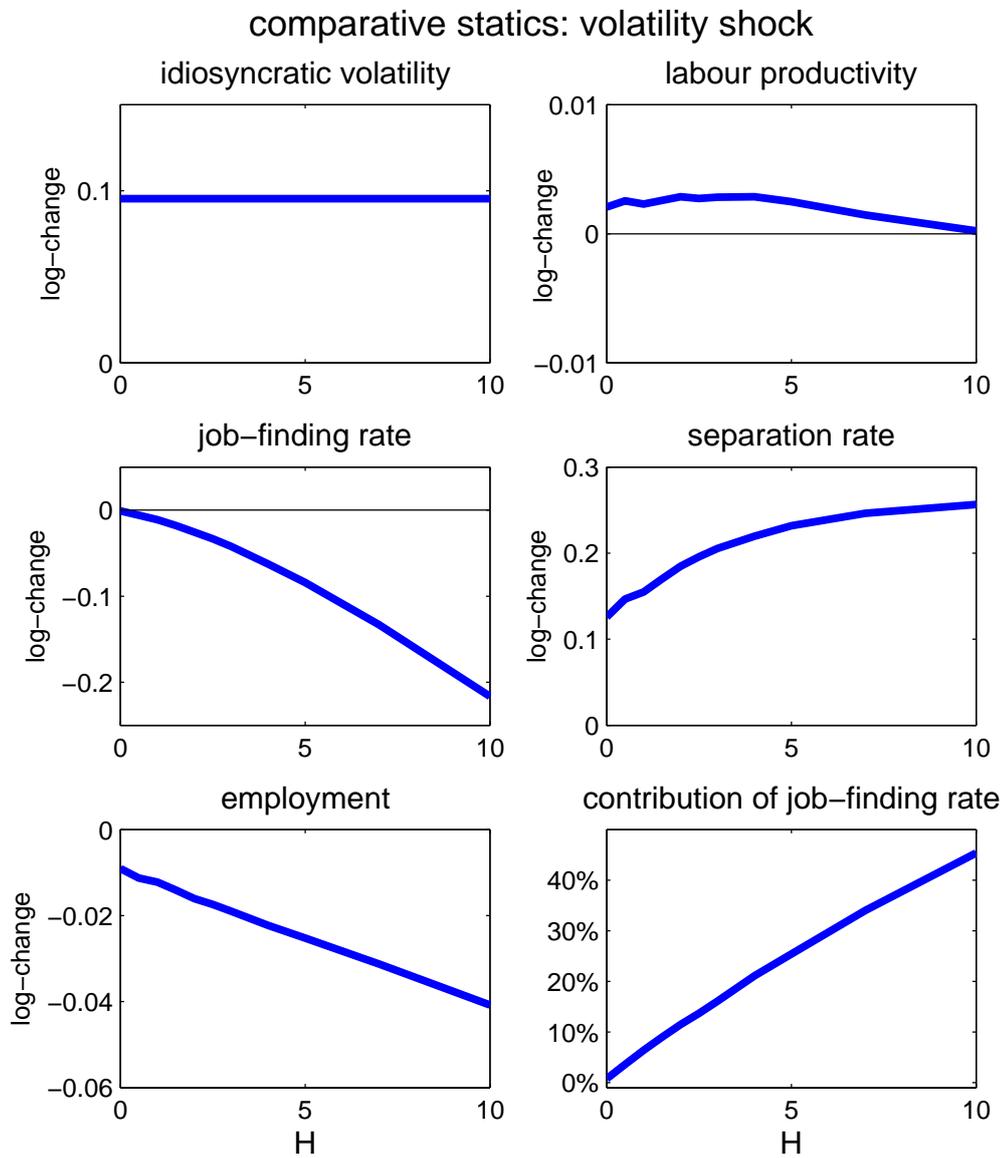


Figure 1.6: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The job-finding rate only falls and thus contributes to lower employment in the presence of creation cost $H > 0$.

calculate the expected continuation value of a firm at the creation cutoff. Consider a firm that is at the creation cutoff n_{LB} in the current period. It creates positions until the marginal value of a position is equal to the cost of creating it. Hence, the marginal value will be at the upper bound. In the absence of aggregate shocks, the marginal value of a position at this firm will still be equal to the cost of creating a position if its idiosyncratic productivity does not change or if it receives a positive shock.³² In the former case it will keep the same number of positions as its optimization problem looks the same as in the current period, whereas after a positive productivity shock, it will create new positions until the marginal value is equal to the cost of creating it. If the firm receives a negative productivity shock, the marginal value of a position falls below the cost of creating positions and, depending on the size of the negative shock, the firm might close positions or even fire workers. If it fires workers, the marginal value reaches its lower bound at 0. The range of interior values for the continuation value is important for the evaluation of the effect of a second moment shock. If the size of the inaction region is such that a firm is unlikely to be in the inaction region after a new productivity draw, increasing the size of the shock does not matter much. In contrast, if there is a large probability that the firm's marginal value will be in the interior region, larger negative shocks reduce the interior continuation value, while larger positive shocks still do not affect the continuation value due to it being equal to its upper bound already. Then, higher uncertainty reduces the expected continuation value.

Creation costs are important to get a negative effect on the job finding rate, because they make an interior solution for the continuation value more likely. The probability of drawing a level of idiosyncratic productivity, for which the firm has a continuation value between the lower bound of 0 and the upper bound of the creation cost, becomes larger in the presence of creation costs.

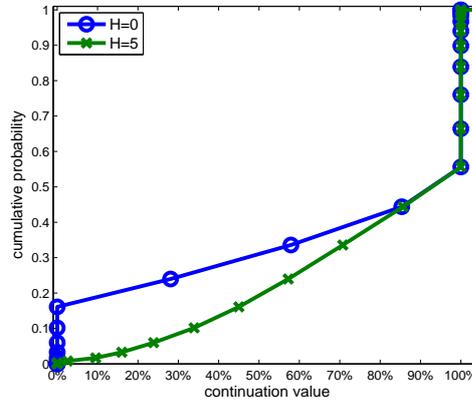


Figure 1.7: Cumulative distribution function of continuation value of a firm with median productivity at the creation cutoff when a shock arrives.

Figure 1.7 draws the cumulative distribution function of the continuation value of a firm at the creation cutoff if it gets hit by a shock. The continuation value is given relative to its maximum value of $H + (1 - \lambda) \frac{c}{q}$, to make comparison between a model without creation costs ($H = 0$) and with creation costs ($H = 5$) easier. If a firm that was at the creation cutoff receives a positive productivity shock, the continuation value is at its upper bound regardless

³²Due to depreciation of positions ($\delta > 0$), this is also true after a small negative productivity shock as the marginal value is decreasing in the number of positions.

of the presence of creation costs. In the figure, this can be seen by the jump of the cumulative distribution function from roughly 0.5 to 1 at the maximum continuation value of 100%. If a negative shock is drawn, the continuation value falls below 100% and eventually reaches 0 for very negative shocks. Note that without creation costs it is more likely that a firm that creates a position in this period will destroy it in the next period, because it could open new positions in later periods without cost. A firm with creation costs, however, would not only incur the cost of filling a newly created position but also the cost of creation in later periods. Therefore, it is less likely to destroy the position when a negative shock hits.

1.5.2 Decomposition of effects

The goal is to decompose the effects of higher uncertainty into three components, namely the wait-and-see effect, the realized volatility effect and the general equilibrium effect due to changes in market tightness.

The wait-and-see effect means that higher uncertainty makes firms more reluctant to create positions, because a future adverse productivity shock becomes more likely. This reduces the expected continuation value and shifts the creation cutoff n_{LB} down. Likewise, firms fire fewer workers as the chance of a future positive productivity shock becomes higher. The resulting increase of the expected continuation value shifts the firing cutoff n_{UB} up. Hence, the wait-and-see effect implies that higher uncertainty leads to fewer positions being created and fewer workers being fired.

The realized volatility effect occurs, because a higher volatility of idiosyncratic firm productivity makes more firms experience large positive or negative shocks. Holding a firm's cutoffs constant, it is thus more likely to get a low enough productivity draw that makes it fire workers, and also more likely to get a high enough productivity draw that makes it create new positions. Therefore, the realized volatility effect leads to more positions being created and more workers being fired in the presence of higher uncertainty.

Thirdly, in general equilibrium changes in market tightness affect a firm's optimal behaviour. For example, if the combination of wait-and-see effect and realized volatility effect imply that the aggregate demand for new hires falls short of supply, market tightness falls to ensure market clearing in the matching market. Then, it becomes cheaper for firms to fill their empty positions. This primarily shifts the intermediate cutoff n_{MB} upwards, which means that they are more willing to fill an empty position as opposed to shutting it down. In the following subsections, I discuss the three components in detail.

1.5.2.1 Wait-and-see effect

The wait-and-see effect measures the change in aggregate hiring and firing that is due to firms adjusting their cutoffs in response to increased uncertainty. It is not useful, however, to just look at the change of cutoffs conditional on idiosyncratic productivity x , like the amount that $n_{LB}(x)$ shifts. This is because firms with the same x are not directly comparable as they are at different percentiles of the cross-sectional distribution when volatility changes. For example, take a firm, whose idiosyncratic productivity x is one standard deviation below the mean. If the standard deviation is doubled, a firm with this same x is only half of a standard deviation below the mean. Consequently, it is more likely that this particular firm gets a new productivity draw that lies below x in the presence of higher uncertainty. Therefore, I will

compare how the cutoffs of firms at the same percentile of the cross-sectional distribution change. More precisely, I measure the change of cutoffs at the same percentile relative to their “perfect-insurance” counterparts. This method takes care of the fact that firms at the same percentile of the distribution do not have the same idiosyncratic productivity x . I explain the method in detail in appendix 1.C.1. The results are cutoff functions $d_{LB}^\sigma(p)$, $d_{MB}^\sigma(p)$, and $d_{UB}^\sigma(p)$. They provide a measure for the log of the three cutoffs as functions of the percentile p , given a certain level of idiosyncratic volatility σ .

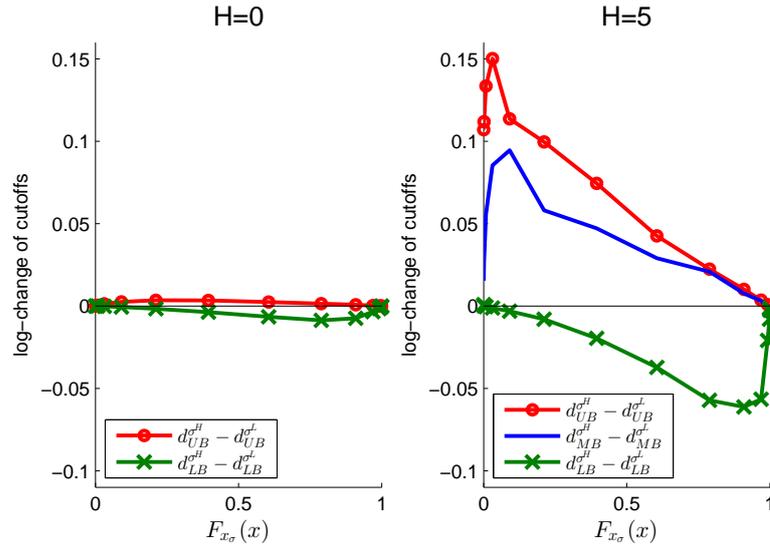


Figure 1.8: Wait-and-see effect: change in cutoffs after a 10% increase in uncertainty.

Figure 1.8 shows how these three cutoff functions change, when idiosyncratic volatility is increased. The cutoffs at the lower bound fall, which means that comparable firms will create fewer positions in the presence of higher uncertainty. In contrast to that, the cutoffs at the upper bound increase, which makes firms fire fewer workers. The reason is that at the lower cutoff the concavity of the marginal continuation value dominates such that higher uncertainty reduces the continuation value. At the upper cutoff instead, the convexity dominates such that higher uncertainty increases the expected continuation value. The direction of the shift at the medium cutoff is ambiguous. As the medium cutoff is determined by the condition whether a firm wants to fill empty positions or not, there are two opposing effects. On the one hand, the firm incurs a cost if it hires new workers. This is similar to a firm at the lower cutoff which creates positions and hires new workers. Higher uncertainty makes firms more reluctant to hire new workers, which would shift down the cutoff. On the other hand, if the firm decides not to hire new workers, it shuts down the empty positions. This is similar to a firm at the upper cutoff, which not only fires workers but also shuts down empty positions. Higher uncertainty makes firms also more reluctant to do so. This would shift the medium cutoff upwards. Whether the medium cutoff shifts upwards or downwards is then determined by the size of the creation cost relative to the filling cost. The higher the cost to create a position, the more reluctant a firm is to shut it down. Therefore, for sufficiently high H , the cutoff will shift upwards as can be seen for $H = 5$ in figure 1.8. In contrast, when there is only a small cost to create positions, the reluctance of the firm to fill empty positions becomes dominant, and the cutoff will shift down. In particular, in the absence of creation cost ($H = 0$), the medium

and lower cutoffs coincide such that they move down as drawn in the left panel of figure 1.8.

Now, it is possible to calculate the implications of the wait-and-see effect isolated from the other effects when uncertainty increases. To do so, I calculate the differences $d_j^\sigma(p)$ for low uncertainty (σ^L) and high uncertainty (σ^H) in partial equilibrium. Then, I shift the cutoffs $n_j(x)$ by the change of these difference to get the resulting cutoffs $n_j^{W\&S}$ for the wait-and-see effect:

$$\log n_j^{W\&S} \left(F_{\sigma_x}^{-1}(p) \right) = \log n_j \left(F_{\sigma_x}^{-1}(p) \right) + \left[d_j^{\sigma^H}(p) - d_j^{\sigma^L}(p) \right],$$

$$j \in \{LB, MB, UB\}.$$

For example, if higher uncertainty increases the difference d^σ between the lower cutoff $n_{LB}(x)$ and the perfect insurance cutoff $\tilde{n}_{LB}(x)$ by 4% for firms at the 75th-percentile, then the wait-and-see effect is calculated by reducing the lower cutoff $n_{LB}(x)$ by 4% at the 75th-percentile. This means that the cutoffs are shifted by the amount shown in figure 1.8, but there is no actual change in volatility when the economy is simulated. The latter effect is the realized volatility effect described in the next subsection.

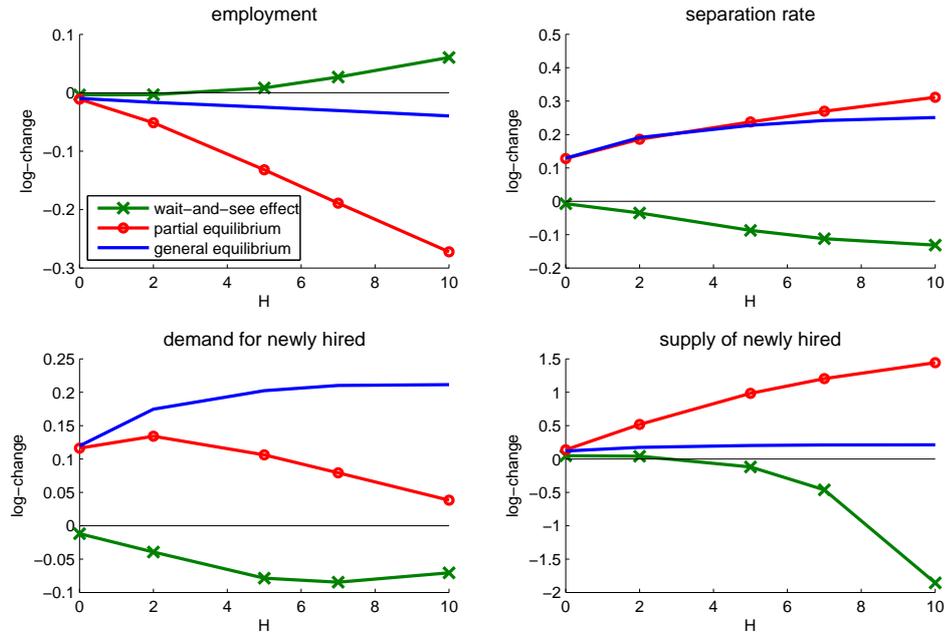


Figure 1.9: Decomposition of a 10% increase in uncertainty into wait-and-see effect (green curves with crosses), realized volatility effect (difference between red curves with dots and green curves with crosses), and general equilibrium effect (difference between solid blue curves and red curves with dots).

Figure 1.9 shows the change of employment, separation rate and demand and supply of newly hired workers after a 10% increase in uncertainty as a function of creation cost H . The separation rate falls unambiguously, because the firing cutoff shifts upwards. The size of the effect is increasing in H , because the wait-and-see effect becomes stronger as argued above. The aggregate demand for newly hired workers falls as well. This is driven by the downward shift of the creation cutoff $n_{LB}(x)$.

Note that the middle cutoff $n_{MB}(x)$ also determines how many firms want to fill their empty positions. This cutoff shifts up for larger H as seen above. The willingness of more

firms to fill their vacant positions increases demand for new hires. But it is dominated by the change in the creation cutoff. Intuitively, any downward shift of the creation cutoff means that fewer new positions are created. Each position that is not created reduces the demand for new hires by one worker. In contrast, an upward shift in the middle cutoff means that only positions that are empty are then filled as well. As only a small fraction of positions is vacant at the beginning of the period, a shift of the middle cutoff has smaller effects on labour demand than a similar shift of the creation cutoff.

Taking the changes of the separation rate and of the number of new hires together, one can calculate the effect on employment. As the two effects run in opposite directions, the resulting change in employment is ambiguous. It turns out that the reduced separation rate dominates for high H and employment rises slightly. In the fourth subplot of figure 1.9, the implied change of the supply of newly hired is drawn. Note that market tightness remains unchanged such that changes in supply are entirely driven by an increase or decrease of the pool of unemployed. Therefore, whenever employment increases, there are fewer unemployed and the matching function implies that the number of new matches changes proportionally to this change in unemployment.

1.5.2.2 Realized volatility effect

The realized volatility effect captures the effects of higher uncertainty through the increase in the cross-sectional dispersion. When volatility increases, there are more firms that get hit by sufficiently negative or positive shocks that lead them to fire workers or to create new positions. Hence, the realized volatility effect captures the changes due to larger shocks occurring, holding the cutoffs fixed, whereas the wait-and-see effect accounts for the change of cutoffs. The combined impact of wait-and-see effect and realized volatility effect is given by the partial equilibrium response after an increase in uncertainty. This means that market tightness is held constant, but firms optimally adjust their cutoffs and experience larger shocks. These partial equilibrium responses are drawn in figure 1.9. The difference between the partial equilibrium effect and the wait-and-see effect is then due to the realized volatility effect. If firms get hit by larger shocks, but did not adjust their cutoffs, they are more likely to be below the creation cutoff or above the firing cutoff at the beginning of the period. This can be seen in figure 1.9 as an increase in the separation rate and an increase in the demand for newly hired workers. The higher separation rate has a negative effect on employment, whereas more newly hired workers would increase employment. It can be seen from the first subplot that the former effect dominates and employment falls for all values of creation cost. This fall is increasing in H . Lower employment increases the pool of unemployed, which in turn increases the supply of new matches, as can be seen in the fourth subplot. Note that for larger values of H , the increase in the supply of new workers exceeds the increase in demand for new workers. Therefore, market tightness has to fall to get market clearing, as explained in the next subsection.

1.5.2.3 General equilibrium effect

Whenever changes in uncertainty do not change supply of new matches as given by the matching function by the same proportion as aggregate demand of firms for new matches, market tightness has to adjust to clear the market. In the subsection above, it could be seen that higher

uncertainty leads to an excess supply of workers except for very small values of H . Therefore, market tightness falls in general equilibrium. This reduces the supply of new matches, because the job-finding rate for unemployed falls. Aggregate demand increases, because it becomes easier and thus cheaper for firms to fill their empty positions. In addition, lower market tightness reduces the wage and thus increases profits.

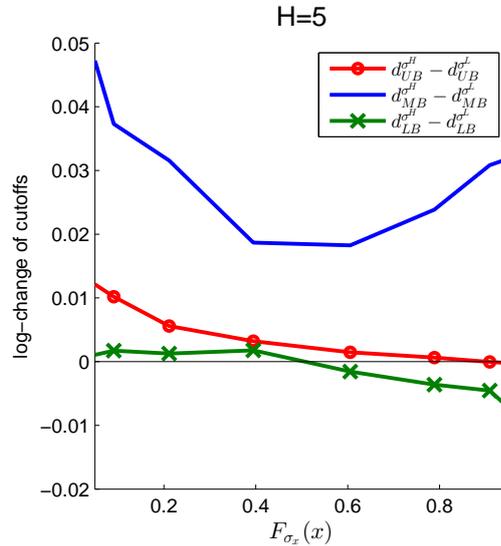


Figure 1.10: General equilibrium effect: change in cutoffs due to a change in market tightness that is necessary to clear markets after a 10% increase in uncertainty.

Figure 1.10 shows the change in cutoffs that is due to the change in market tightness. It turns out that the biggest change in cutoffs occurs at the middle cutoff n_{MB} . This cutoff is determined by the condition that a firm is indifferent between filling their empty positions and shutting them down. Market tightness determines the cost of filling the position, which makes it an important determinant for the cutoff. The creation cutoff n_{LB} instead is determined by the condition that a firm is indifferent between creating and filling a new position or not. When the creation cost is large relative to the filling cost, a lower filling cost makes it more profitable for the firm to create and fill new positions, but the total cost is not affected as much in relative terms as it is for the middle cutoff. Note that the quantitative shift of cutoffs depends on the change in market tightness. For low values of H , the market clearing condition is almost satisfied in partial equilibrium. Then, market tightness does not adjust much, which in turn leads to a small shift of cutoffs.

Figure 1.9 also draws the general equilibrium effects of higher uncertainty. Relative to the partial equilibrium outcome, market tightness falls for all but the very small values of H . This primarily increases demand for newly hired workers, while the separation rate is hardly affected. The result is that employment increases relative to partial equilibrium. Both lower market tightness and lower unemployment reduce the supply of new matches. In general equilibrium, the change of demand and supply of new hires must be the same. Note that the excess supply of workers was increasing in H in partial equilibrium. This leads to a bigger fall of market tightness when H is large. As a result, the job-finding rate falls in the presence of a large creation cost as could be seen in the partial equilibrium exercise in figure 1.6.

It is important to distinguish between the job-finding rate and the number of newly hired workers, which is given by the product of job-finding rate and the number of unemployed.

While higher uncertainty reduces the job-finding rate when H is large, the absolute number of new matches increases, because unemployment increases.³³

1.5.3 IRFs of uncertainty shocks

In this subsection, I compute impulse response functions after an increase in idiosyncratic uncertainty in the full model. I show that higher uncertainty leads to a fall in employment. In the absence of creation costs, this fall in employment is entirely driven by a higher separation rate. The higher the creation cost H becomes, the more the job-finding rate falls similar to the results in the comparative statics exercise above. This means that the job-finding rate contributes more to changes in employment. In the full model with time-varying aggregate shocks, the response of endogenous variables after a shock depends on the distribution of firms when the shock hits. For example, in some periods an increase in uncertainty leads to a fall in employment that is twice as strong as the fall after a shock in other periods. In particular, I show that the response of employment becomes stronger, when a shock hits in periods of high unemployment.

Figure 1.11 shows the median response of labour market variables for different values of the creation cost.^{34,35} Similar to the results in the comparative statics exercise above, the job-finding rate hardly responds in the absence of creation costs. In the presence of higher creation costs H the job-finding rate falls. The increase in the separation rate becomes stronger as well, but the peak change of the job-finding rate becomes larger relative to the peak change of the separation rate. When $H = 10$, the peak change of the job-finding rate is more than 70% of the peak change in the separation rate. Compared to the US data in figure 1.2, the peak fall in the job-finding rate is still too small, as it was found to be about twice as big relative to the increase in the separation rate. The model can capture well the timing of the peak responses. The job-finding rate peaks after around 5 quarters in the model and in the data, whereas the separation rate peaks early after the shock.

Labour productivity increases slightly for low values of the creation cost, while it falls for higher values. Note that the idiosyncratic volatility shock is constructed such that labour productivity would not change in a frictionless economy. It changes in the model, because the distribution of positions across firms changes. When the creation cost is higher, the wait-and-see effect becomes stronger. This increases the dispersion of the distribution of positions over

³³In steady state, unemployment rate u , job-finding rate f , and separation rate s satisfy $u = \frac{s}{s+f}$. Then the change of newly hired workers fu is given by

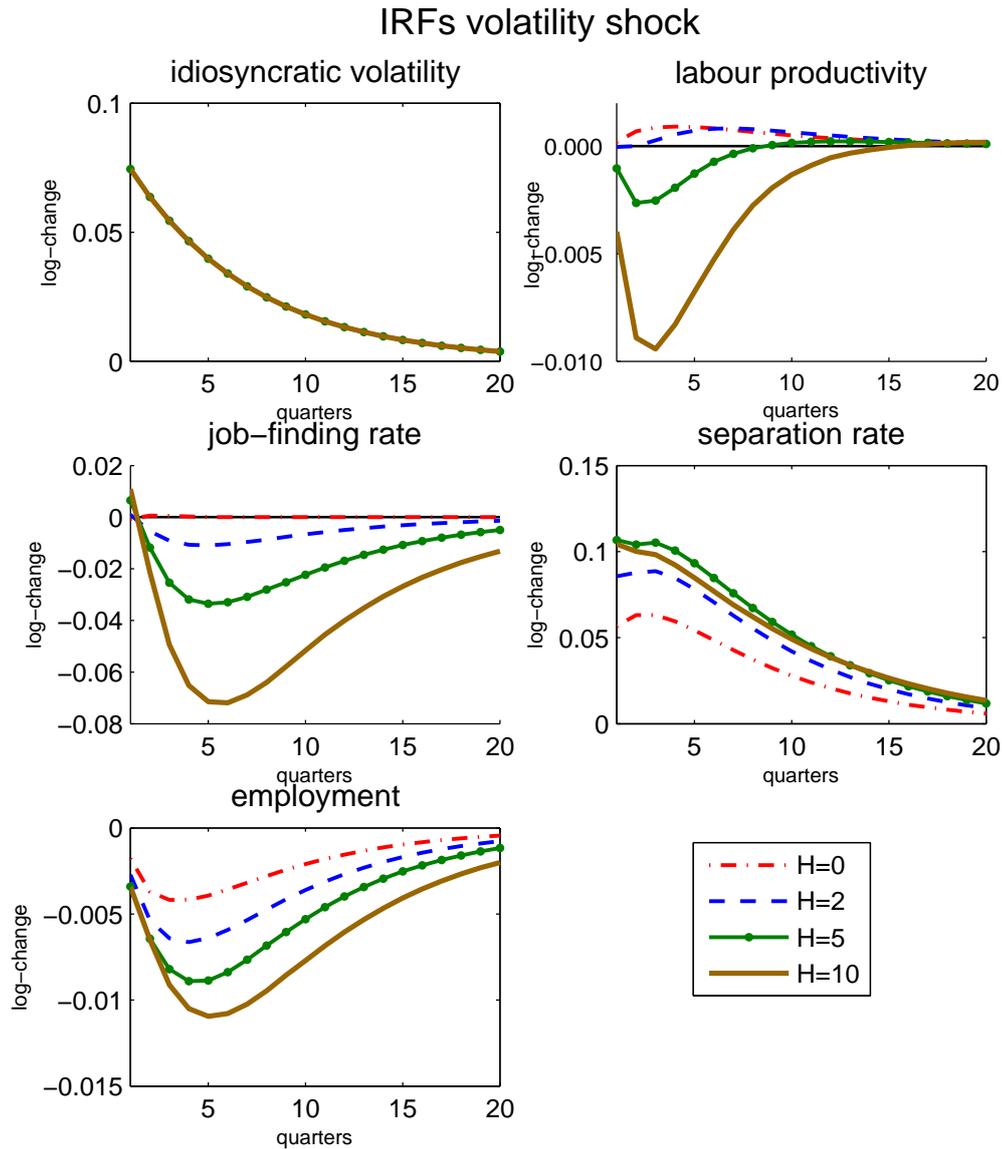
$$d \log (fu) = (1 - u) d \log s + u d \log f.$$

As $(1 - u)$ is considerably larger than u , the number of newly hired will increase unless the job-finding rate falls by more than a multiple of the increase in the separation rate:

$$d \log (fu) > 0, \text{ if } -d \log f < \frac{1 - u}{u} d \log s.$$

³⁴I calculate the responses after a one standard deviation increase of σ_x at 500 different times along a randomly drawn time path of aggregate shocks. Then I take the median value of these realizations of the respective variable at each point in time.

³⁵When comparing the IRFs with the ones estimated in the US data in figure 1.2, note that a change in idiosyncratic volatility has bigger effects on the transition rates than a change in the empirical volatility measure in the data. It is not surprising that the empirical measure obtained from implied volatilities is not directly comparable to changes in the standard deviation of the idiosyncratic shock, which is not observed in the data. Therefore, my focus is on analysing to what extent the separation rate and the job-finding rate drive changes in employment. In section 6, I use a scaled down version of the empirical volatility measure into my model to compare the model's performance with time series data from the US.



firms with the same idiosyncratic productivity, and reduces measured labour productivity. This effect dominates the implications of decreasing returns to scale, which increase the average product when employment falls. Note that this fall is qualitatively different compared to the comparative statics exercise, where labour productivity increased in figure 1.6. Quantitatively, the initial drop in labour productivity for high H is comparable to the IRF estimated in the US data, when the size of the shock is chosen to get the same peak response of employment. The confidence intervals, however, are wide in the empirical IRF, such that the drop is not significant. But also in the model the change in labour productivity can vary much depending on the state in which the shock hits, as can be seen in figure 1.12.

Whereas figure 1.11 plots the median impulse response functions in the model, the ob-

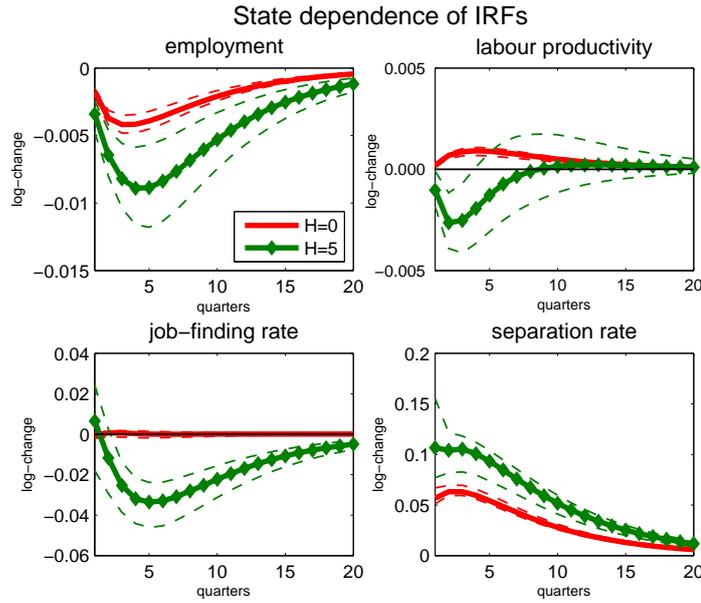


Figure 1.12: Impulse response functions for idiosyncratic volatility shock in the model. Solid lines depict the median change of the respective variable. Dashed lines draw the 5th and the 95th percentiles of the responses.

served responses can vary significantly depending on the initial state. In the model, the cross-sectional distribution of positions and idiosyncratic productivity is a state variable. As it is solved using non-linear methods, the response to shocks can differ depending on the initial distribution. In particular, for higher values of the creation cost H , the dependence on the initial state becomes more relevant. Figure 1.12 repeats the median IRFs from figure 1.11 for $H = 0$ and $H = 5$. In addition, it also shows the 5th and the 95th percentile of the responses in dashed lines. One can see that for $H = 0$ these bands are relatively narrow around the median response. In contrast, they are much wider for $H = 5$. For example, the peak drop in employment can be twice as much in the 95th percentile as it is in the 5th percentile. In theory, the whole distribution at the initial state matters. One moment of this distribution, which is easily observable, is the unemployment rate when the shock hits. On average, employment falls more after an increase in uncertainty when the unemployment rate was already high at the time the shock hits.

H	n^L	$d \log n$ n^M	n^H
0	-0.005	-0.004	-0.004
2	-0.008	-0.007	-0.006
5	-0.011	-0.009	-0.008
10	-0.015	-0.012	-0.011

Table 1.5: State dependence of impulse response functions. The sample of IRFs is divided into three equal parts according to employment when the volatility shock hits. (n^L , n^M , n^H correspond to the lowest, middle, and highest third, respectively.) Then the average peak changes of log-employment in the subsamples are calculated and reported depending on creation cost H .

Table 1.5 reports the peak change in employment, when the sample of impulse responses is divided into three equally sized parts according to initial unemployment. The initial unemployment rate is more important for high values of H . For example when $H = 10$, employment

falls by 1.5% when the unemployment rate was in the highest third, and by only 1.1% when the unemployment rate was in the lowest third. One reason, why shocks affect employment more when the unemployment rate is high, is implied by the matching function, as discussed in Michaillat (2014). When the unemployment rate is low, the same relative change of market tightness affects the unemployment rate less. Equation (1.44) in appendix 1.A.2 shows that changes in log-employment are approximately proportional to the unemployment rate, when the job-finding rate or the separation rate are changed by the same relative amount.

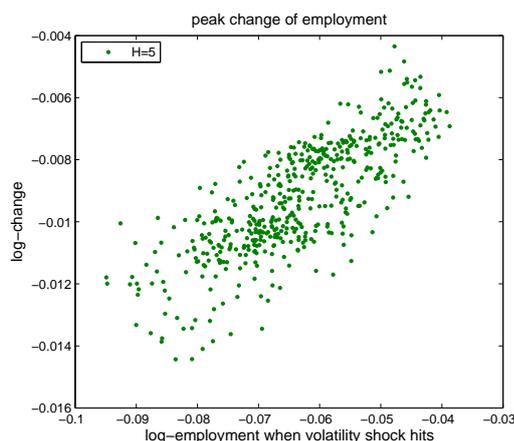


Figure 1.13: Each dot corresponds to the peak response of log-employment in one impulse response function after a positive uncertainty shock for $H = 5$.

But even if the initial unemployment rate is taken into account, some state dependency remains. Figure 1.13 plots the observed peak responses of employment as a function of initial employment for $H = 5$. It can be seen that the range of possible values conditional on initial employment is still large. These differences are due to characteristics of the cross-sectional distribution that are not captured well by aggregate employment. For example, it could be that in some periods there are more firms close to (one of) the cutoffs, such that they are more responsive to additional shocks. Figure 1.13 also gives an indication that the first moment of the distribution is not enough to forecast aggregate outcomes. Section 1.7 explains the algorithm that I develop to get more accurate forecasts of market tightness without adding many characteristics of the cross-sectional distribution as additional state variables.

1.6 Comparison with US time series

In this section, I compare the model with the US time series from 1998 to 2013. First, I argue that a model with productivity shocks alone cannot match the observed unemployment rate after 2004 well. Then, I show that the unemployment rate can be matched much better when idiosyncratic volatility shocks are added. In addition, when the creation cost is large, the model also performs better in matching the observed job-finding rate.

In the exercises of this section, I back out the aggregate productivity shock such that labour productivity in the model equals labour productivity in the data. In the first exercise I assume that idiosyncratic volatility is constant. The resulting series of the unemployment rate, the job-finding rate, and the separation rate are drawn in figure 1.14. At a first glance, the un-

employment rate is not as volatile in the model as in the data.³⁶ This shortcoming could be alleviated if the model was recalibrated to produce a larger volatility of employment. In particular, a lower value of ω_2 in the wage equation could be used. Then, the wage reacts less to changes in market tightness. Hence, the marginal profit of positions become more volatile and firms will react more strongly in response. Similarly, Hall (2005) has shown that wage stickiness can help to overcome too little employment volatility in matching models. The results are shown in figure 1.32 in appendix 1.D. Note, however, that even if the unemployment rate reacted more to productivity shocks, it would be difficult to match the observed unemployment rate after 2004. Whereas it was low in 2007, the model would predict the lowest value for unemployment already in 2004. In addition, unemployment increases in the model from 2011 onward, while a steady decline was observed in the data.

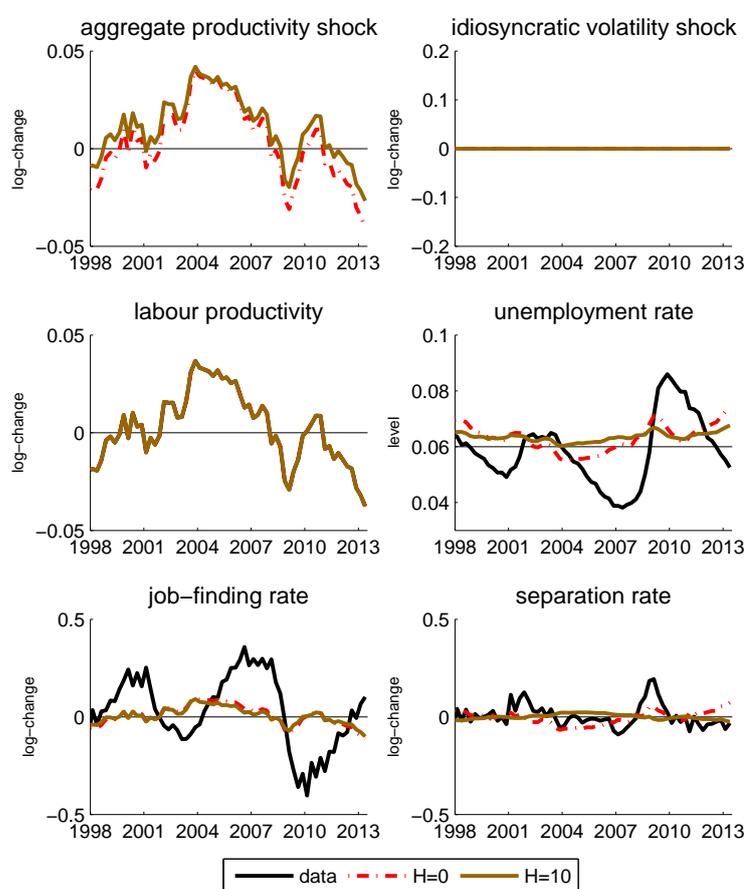


Figure 1.14: Model without uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data (linear trend removed).

Uncertainty shocks seem a good candidate to explain the low unemployment rate in 2007, as well as the decline in unemployment since 2010. In the data, uncertainty was low in the mid 2000s, it peaked in 2008, and it has been falling since then. In the model, low uncertainty reduces unemployment. Therefore, in a second exercise I use the estimated index of implied

³⁶Too little employment volatility has been documented in standard search and matching models, for example, in Costain and Reiter (2008) and Shimer (2005).

volatility, shown in figure 1.1.³⁷ Whereas the idiosyncratic volatility shock is fed into the model according to observed variations in the implied volatility index, the aggregate productivity shock is still estimated in each quarter to match labour productivity in the model to the data.

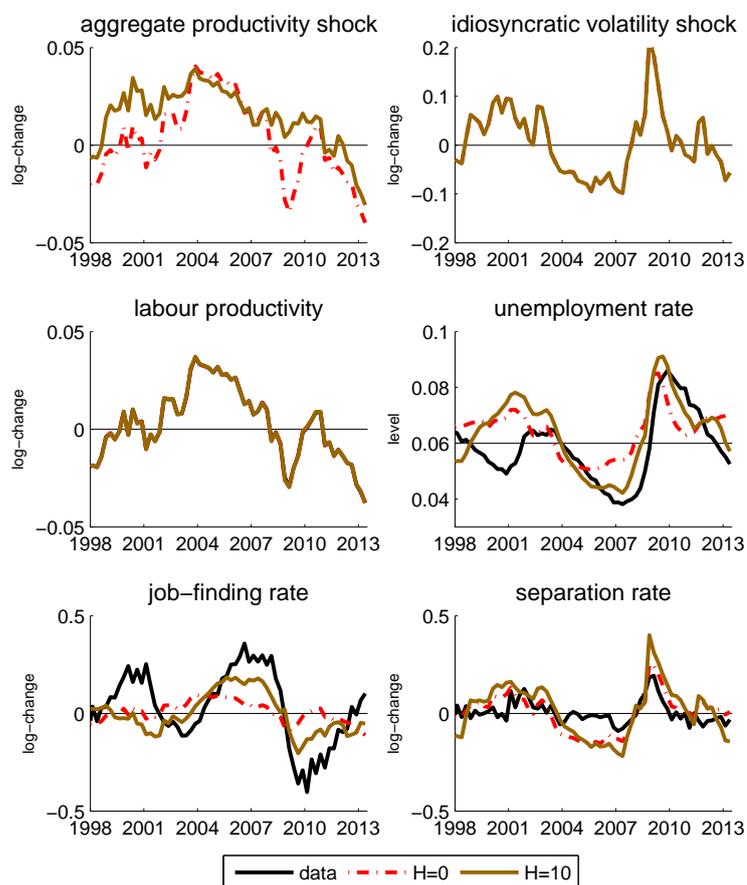


Figure 1.15: Model with aggregate productivity and uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data. Log-changes in idiosyncratic volatility are 20% of changes of the implied volatility index estimated in section 2. Linear trends are removed from US data series.

Figure 1.15 shows the resulting time series for the unemployment rate, the job-finding rate and the separation rate. While the model without creation costs performs better when the uncertainty shock is added, it still cannot reproduce the low unemployment rate in 2007. Fluctuations of the job-finding rate are too small as well. When creation cost $H = 10$ instead, the model closely tracks the unemployment rate in the 2000s. It can explain a steep fall in the job-finding rate after 2007. But the job-finding rate is still a bit less volatile than in the data, whereas the separation rate is too volatile. In terms of the timing, the unemployment rate and the job-finding rate lead their respective empirical counterparts. This indicates that market tightness reacts more sluggishly in the data than in the model. If the creation cost H was made dependent on the aggregate number of positions created in each period as in Fujita and Ramey (2007), the response of aggregate creation to aggregate shocks would become smoother. This

³⁷I use a scaled down version of the log-changes in implied volatility because employment in the model reacts more strongly to the same relative change in idiosyncratic volatility than employment in the data reacts to changes in implied volatility.

could improve the fit of the model.

The model has difficulties to explain the low unemployment rate in the run-up of the 2001 recession. The reason is that observed volatility was above trend and labour productivity was close to the trend. Then, the model predicts the unemployment rate to be above trend. It is possible to back out the idiosyncratic volatility shocks that would be needed to match the unemployment rate. The resulting shocks are in figure 1.33 in appendix 1.D. It can be seen that uncertainty would have to be lower than observed before 2001 to explain the unemployment rate. Nevertheless, over the whole sample the correlation between the backed out volatility shocks and the empirical measure for uncertainty is 0.46 for $H = 0$ and 0.59 for $H = 10$.

1.7 Algorithm to solve the model

In the presence of aggregate shocks, market tightness depends on the distribution of positions and idiosyncratic productivity across firms, which is an infinite-dimensional object, as is common in models with heterogeneous agents. One method to solve the model numerically is to follow Krusell and Smith (1998) by approximating the infinite-dimensional state space using a finite number of aggregate state variables characterizing the distribution. I extend the Krusell and Smith (1998) algorithm by adding *one* specific additional state variable to capture omitted characteristics of the distribution: I add the residual in the law of motion for market tightness as a state variable. Thereby, I can increase the accuracy of the law of for market tightness significantly, and at the same time ensure that the market clears in each period. The next subsection describes the algorithm in a general way before I describe how I apply it to the specific model.

1.7.1 General description of the residual method

In a heterogeneous agents model with aggregate uncertainty, current and future prices can in general depend on the whole cross-sectional distribution of agents' characteristics. Krusell and Smith (1998) suggest to approximate this distribution by using some of its moments, denoted by m . A law of motion for m , which can also depend on exogenous state variables z , describes their transition from one period to the next:

$$m' = \Psi_m(m, z). \quad (1.23)$$

In addition, it might be necessary to approximate a set of prices p as a function of the state variables:^{38,39}

$$p = \Psi_p(m, z). \quad (1.24)$$

These laws of motion are used to find the optimal individual policy functions. But when the model is simulated, the approximation leads to residuals x_i such that for the simulated

³⁸For example this could be the rental rate of capital or the wage. If labour supply is endogenous, the wage can potentially depend on the whole distribution of agents. In my model, I need a law for market tightness.

³⁹Note the following difference between laws of motion for moments m and those for prices p . Laws of motion for m describe the intertemporal transition of moments from one period to the next. In contrast, laws for prices describe a price in the current period depending on the current period's state variables.

moments \tilde{m} and prices \tilde{p} the following equations hold:

$$\tilde{m}' = \Psi_m(\tilde{m}, \tilde{z}) + x_m, \quad (1.25)$$

$$\tilde{p} = \Psi_p(\tilde{m}, \tilde{z}) + x_p. \quad (1.26)$$

The residual captures the information of the cross-sectional distribution that was omitted or is not captured well by the functional form imposed on Ψ_i . Since the omitted moments are typically autocorrelated, also the residual will be autocorrelated. For example, if the residual x_p is positive in one period, this means that the market clearing price is higher than what would have been predicted based on the aggregate state variables (\tilde{m}, \tilde{z}) . Then it is likely that in the next period, the market clearing price will again be above its predicted counterpart based on next period's state variables (\tilde{m}', \tilde{z}') . I propose to make use of this information contained in the residual when forecasting future variables. I do so by adding the variables $x \equiv (x_m, x_p)$ to the set of state variables. The laws of motion that are used in the individual's problem become

$$m' = \Psi_m(m, z) + x_m, \quad (1.27)$$

$$p = \Psi_p(m, z) + x_p, \quad (1.28)$$

$$x' = \Psi_x(m, z, x). \quad (1.29)$$

Given these law of motions, the individual optimization problem can be solved using standard techniques. Adding the variable x_i affects this in the same way as adding any other aggregate state variable would.

Once the optimal policy functions are found, one could simulate the economy like in the standard Krusell and Smith (1998) algorithm. Then, adding a residual as a state variable is equally costly as adding one higher moment. The advantage is that by not being restricted to a specific moment, accuracy could be improved by more, because the residual summarizes the error made due to omitting all higher moments.

A further way to make use of the residuals in the simulation part comes from the additional degree of freedom that is obtained, because the residuals x_i are not predetermined. Usually state variables are directly calculated from the simulated distribution, for example the first or higher moments thereof. The residual x_i , however, can be determined in the period itself. As it affects the agents' policy functions, it can be chosen such that market clearing holds, if it is added to law of motions for prices as in equation (1.28). If it is added to intertemporal law of motions like in equation (1.27), the residual can be used to get next period's value m' to be equal to its forecast value. In both cases, there is a condition that should hold exactly in each period of the simulation. One can meet this condition by solving for the residual. As individual policy functions depend on it, the residual can be adjusted until the conditions are met.

The residual is thus used as a state variable in the individual problem, but it is endogenously determined in the simulation part. This is similar to the determination of prices in general equilibrium. They are taken as given by individuals but they can be chosen such that markets clear. Similarly, the algorithm lets the price vary until markets clear.

1.7.2 Application of the residual method in this model

In my model, I use the residual method in the law for market tightness. Market tightness is an important variable, because it determines how costly it is for firms to fill their vacancies. When aggregating, the market clearing condition (1.17) must hold:

$$\int_{n_{i-1}} \int_{x_i} (n^*(n_{i-1}, x_i, s) - (1 - \lambda) n_{i-1})_+ dF_s(n_{i-1}, x_i) = a(\theta(s)) u(s). \quad (1.30)$$

I denote by x_θ the residual in the law of motion for market tightness. The other aggregate state variables are aggregate productivity (z), volatility of idiosyncratic productivity (x_σ), the $\frac{1}{1-\alpha}$ -moment of idiosyncratic firm productivity (m^x), and the aggregate number of positions at the beginning of the period (N_{-1}).

Note that the laws for z , x_σ , and m^x only depend on the exogenous processes chosen for aggregate productivity, idiosyncratic productivity, and the volatility of idiosyncratic productivity. Hence they are known (z, x_σ), or can be estimated (m^x) without solving the model.

For θ , x_θ , and N , I specify laws that are (log-)linear in all state variables:

$$\log \theta = \Psi_{\theta,0} + \Psi_{\theta,1} \log z + \Psi_{\theta,2} \log N_{-1} + \Psi_{\theta,3} x_\sigma + \Psi_{\theta,4} m^x + x_\theta, \quad (1.31)$$

$$x_\theta = \Psi_{x_\theta,0} + \Psi_{x_\theta,1} \log z_{-1} + \Psi_{x_\theta,2} \log N_{-2} + \Psi_{x_\theta,3} x_{\sigma-1} + \Psi_{x_\theta,4} m_{-1}^x + \Psi_{x_\theta,5} x_{\theta-1}, \quad (1.32)$$

$$\log N = \Psi_{N,0} + \Psi_{N,1} \log z_{-1} + \Psi_{N,2} \log N_{-1} + \Psi_{N,3} x_{\sigma-1} + \Psi_{N,4} m_{-1}^x + \Psi_{N,5} x_{\theta-1}. \quad (1.33)$$

Given these laws of motion (together with the known transition matrices for the aggregate shocks), one can find the optimal cutoffs in each state on the aggregate grid by iterating the marginal value function. From the cutoffs, one can calculate the number of positions each firm wants to fill, which depends on the idiosyncratic and aggregate state variables:

$$(n^*(n_{i-1}, x_i, z, x_\sigma, m^x, N_{-1}, x_\theta) - (1 - \lambda) n_{i-1})_+. \quad (1.34)$$

Then, I simulate the economy using the histogram method developed by Young (2010).⁴⁰ In the simulation, I impose the market clearing condition (1.30) using the residual method. All arguments except x_θ are predetermined in (1.34). I can solve for the unknown x_θ such that the market clearing condition (1.30) holds. Using the simulated values, the coefficients in the laws of motion (1.31), (1.32), and (1.33) can be estimated and updated. This procedure of solving the individual problem, simulating the economy, and updating the coefficients is repeated until the updated coefficients in one step are close to the ones estimated in the previous step.

Figure 1.16 demonstrates the advantage of adding the residual state variable compared to higher moments of the distribution. Using the simulated 2nd to 10th central moments of the distribution of positions, an alternative law of motion for market tightness can be estimated. In the upper part of figure 1.16, it can be seen that even adding these 9 additional characteristics of the distribution leads to a worse fit of the market clearing values of market tightness. The difference between the market clearing values and the fitted values are shown in the lower part of the plot.⁴¹

⁴⁰This means that instead of simulating a large number of firms, the cross-sectional distribution of positions and idiosyncratic productivity is characterized by a histogram over a fine grid. The advantage of this non-stochastic simulation method is that it avoids sampling error.

⁴¹The predicted value of market tightness is based on current state variables (z, N_{-1}, x_σ, m^x) and the predicted residual \hat{x}_θ according to equation (1.31). The predicted residual is given by equation (1.32) using lagged state variables.

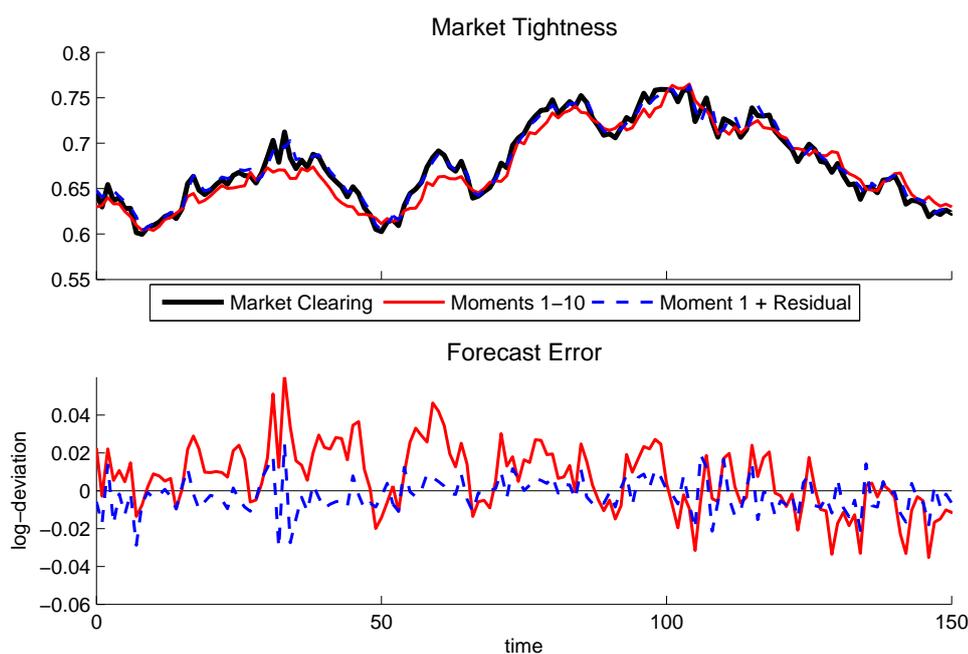


Figure 1.16: The upper part plots an example of market clearing values for market tightness in a simulation when $H = 5$. The dashed blue line draws the estimated values for market tightness using the residual method and the first moment of the distribution of positions. The red solid curve draws the estimated values for market tightness based on the first 10 moments of the distribution of positions without using the residual. The respective relative deviations are drawn in the lower subplot.

1.8 Conclusion

In this chapter, I have analysed the effects of uncertainty on labour market variables in a model with frictional labour markets and costs of creating positions. I demonstrated that in the absence of creation costs, uncertainty shocks hardly affect the job-finding rate. Then, higher uncertainty lowers employment solely by increasing the separation rate. In the data, however, uncertainty shocks reduce employment also by lowering the job-finding rate. When costs of creating positions are added to the model, the job-finding rate responds negatively to higher uncertainty, which brings the model closer in line with the data. I demonstrated that the model without uncertainty shocks cannot match the behavior of the US unemployment rate in the 2000s. In contrast, my full model with idiosyncratic volatility shocks and costs of creating positions can match the observed unemployment rate well. In addition, it can also account for a large share of the steep drop in the job-finding rate after 2007.

This chapter focuses on the demand for labour of heterogeneous firms, while it assumes that workers are risk neutral. If workers were risk averse and could not perfectly insure themselves against unemployment risk, it would become important whether changes in the unemployment rate are driven by changes in the separation rate or by changes in the job-finding rate. When workers are borrowing constrained, for example, it is likely that they prefer many short spells of unemployment over infrequent longer spells. In other words, their welfare is higher, when the same level of unemployment is caused by a high separation rate as opposed

Hence, the forecast error is the difference between the predicted residual \hat{x}_θ and the realized residual x_θ . Note that when market clearing is imposed using the residual method, firms know the exact value of market tightness in the current period, and they also base their decisions on predicted future values of market tightness according to equation (1.32), using the realized residual x_θ .

to a low job-finding rate. Hence, this chapter suggests that higher uncertainty is particularly bad for workers, because it decreases their chances of finding a job when they become unemployed. In chapter 2, we study the implications of longer unemployment spells on precautionary savings of risk averse agents. In the presence of nominal wage stickiness, an increased demand for safe assets can push up real wages and deepen recessions.

My model has also shown that similar shocks can have quantitatively very different responses depending on when the shock hits. For instance, an increase in uncertainty leads to a particularly large drop in employment, when the unemployment rate is already high. This suggests that policy makers should react differently to shocks depending on the state of the economy. Future work could explore the effectiveness of government policy in booms compared to recessions. The state-dependency found in this chapter suggests that macroeconomic variables will react stronger in response to policy interventions in times of higher unemployment.

Appendices to Chapter 1

1.A Empirical part

1.A.1 Data description

The following time series are used for the estimation of the VAR and when comparing the model to the US data.

Labour productivity. $\frac{Y}{L}$ is obtained from the series of seasonally adjusted nonfarm output (PRS85006043) and employment (PRS85006013) from the U.S. Department of Labor: Bureau of Labor Statistics (BLS).

Index of implied volatilities. Implied volatility data from OptionMetrics' Ivy DB database, available from 1996 until the second quarter of 2013, is used. Call options, whose underlying security is not an index, with a maturity of 30 days are used. Quarterly averages of log implied volatility are taken for each of the underlying securities. This gives between 1,758 and 3,836 observations per quarter. Then coefficients of dummy variables for each quarter are estimated, allowing for fixed effects for the securities. These dummy variables, which capture the average implied volatility, form the index of implied volatilities used in this chapter.

Unemployment rate, job-finding rate, and separation rate. They are constructed using data from the Current Population Survey (CPS). The seasonally adjusted unemployment rate is constructed by the BLS (LNS14000000). The job-finding rate and separation rate are calculated using the level of unemployed (LNS13000000) and the number of unemployed for less than 5 weeks (LNS13008396). I account for the redesignment of the unemployment duration question in the CPS, which led to a discontinuous drop of short term unemployment in 1994, as discussed in Shimer (2012). Therefore, the number of newly unemployed is multiplied by the constant 1.1 from 1994. I adjust for time aggregation as emphasized by Shimer (2012) to account for the possibility that workers experience more than one switch between unemployment and employment per month. The monthly job-finding probability is given by⁴²

$$F_t = 1 - \frac{u_{t+1} - u_{t+1}^S}{u_{t+1}}, \quad (1.35)$$

which gives the job-finding rate in a continuous time environment:

$$f_t = -\log(1 - F_t). \quad (1.36)$$

The separation rate in continuous time s_t is then implicitly given by

$$u_{t+1} = \frac{1 - \exp(-f_t - s_t) s_t}{f_t + s_t} l_t + \exp(-f_t - s_t) u_t. \quad (1.37)$$

Quarterly averages of the monthly unemployment rate, job-finding rate and separation rate are taken.

⁴²In what follows, u_t denotes unemployment, u_t^S denotes unemployment less than 5 weeks, and l_t denotes employment.

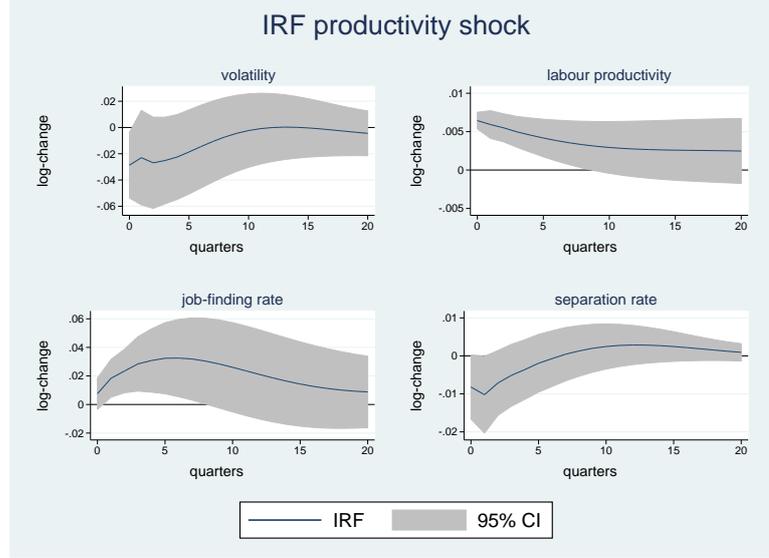


Figure 1.17: Impulse response functions after a shock to labour productivity in US data

1.A.2 Approximation of changes in unemployment

The law of motion of the unemployment rate is given by

$$u_{t+1} = s_t (1 - u_t) + (1 - f_t) u_t. \quad (1.38)$$

If the separation rate s_t and the job-finding rate f_t are constant, the unemployment rate converges to its steady state value

$$\bar{u} = \frac{\bar{s}}{\bar{s} + \bar{f}}. \quad (1.39)$$

The rate of convergence is determined by the sum of \bar{s} and \bar{f} :

$$u_{t+1} - \bar{u} = [1 - (\bar{s} + \bar{f})] (u_t - \bar{u}). \quad (1.40)$$

For example, when the monthly separation rate and job-finding rate sum to 50%, which is roughly the case for the US, any gap between the initial unemployment rate and the steady state value \bar{u} is reduced by 87.5% within a quarter. Therefore, the unemployment rate can be approximated well by only using the current separation rate and job-finding rate, as observed for example by Shimer (2012):

$$u_{t+1} \approx \frac{s_t}{s_t + f_t}. \quad (1.41)$$

Log-linearizing this relationship around \bar{u} , \bar{s} , and \bar{f} , yields

$$\begin{aligned} \log \frac{u_{t+1}}{\bar{u}} &\approx \log \frac{s_t}{\bar{s}} - \frac{1}{s_t + f_t} \left(s_t \log \frac{s_t}{\bar{s}} + f_t \log \frac{f_t}{\bar{f}} \right) \\ &= \frac{f_t}{s_t + f_t} \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right) \\ &= (1 - \bar{u}) \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right), \end{aligned} \quad (1.42)$$

where the last equality uses the definition of the steady state unemployment rate.

An analogous relationship can be derived for employment:

$$n_{t+1} \approx \frac{f_t}{s_t + f_t'} \quad (1.43)$$

$$\begin{aligned} \log \frac{n_{t+1}}{\bar{n}} &\approx (1 - \bar{n}) \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right) \\ &= \bar{u} \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right). \end{aligned} \quad (1.44)$$

This states that relative changes of employment are proportional to the unemployment rate, when the separation rate and the job-finding rate change by the same relative amount. In particular, same relative changes of market tightness have bigger effects on employment when the unemployment rate is high. This is due to the convexity of the quasi labour supply function given by the matching function, as emphasized by Michaillat (2014).

1.A.3 Robustness checks

This section shows that the impulse response function in figure 1.2 is robust to alternative identification, to an alternative volatility measure, and to alternative detrending of the variables. In figure 1.18, the volatility index is ordered first in the Cholesky decomposition, which implies that labour productivity can react instantaneously to volatility shocks. Figure 1.19 uses the volatility index VIX instead of the constructed index based on single stocks. Figures 1.20 and 1.21 use the Hodrick-Prescott filter to remove trends instead of removing a linear trend in the baseline. Figure 1.20 uses the smoothing parameter $\lambda = 100,000$ as preferred by Shimer (2005), whereas figure 1.21 uses the usual value of $\lambda = 1,600$ for quarterly data.

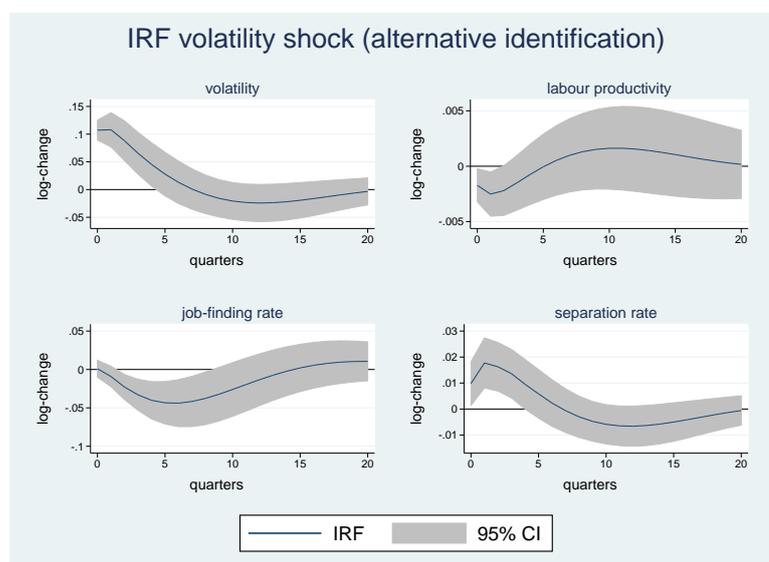


Figure 1.18: Impulse response functions after a volatility shock in US data using an alternative identification, in which volatility is ordered first in the Cholesky decomposition

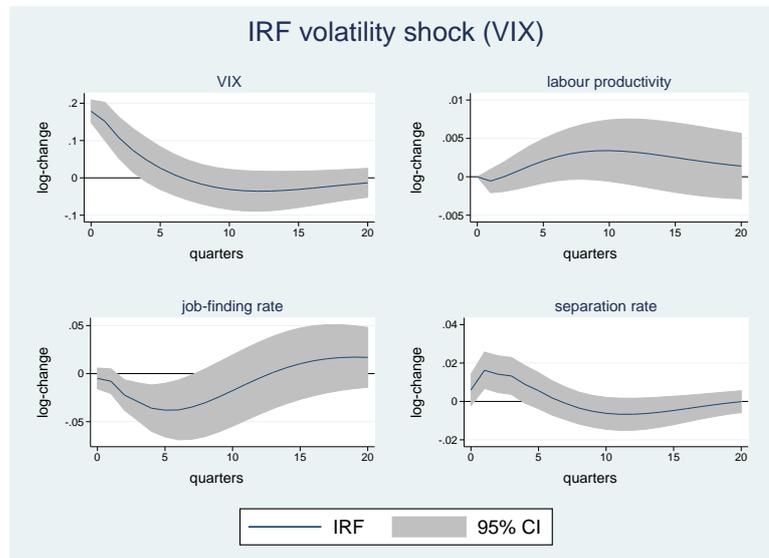


Figure 1.19: Impulse response functions after a volatility shock in US data. Instead of the volatility measure constructed from single stocks, the VIX is used as the volatility measure.

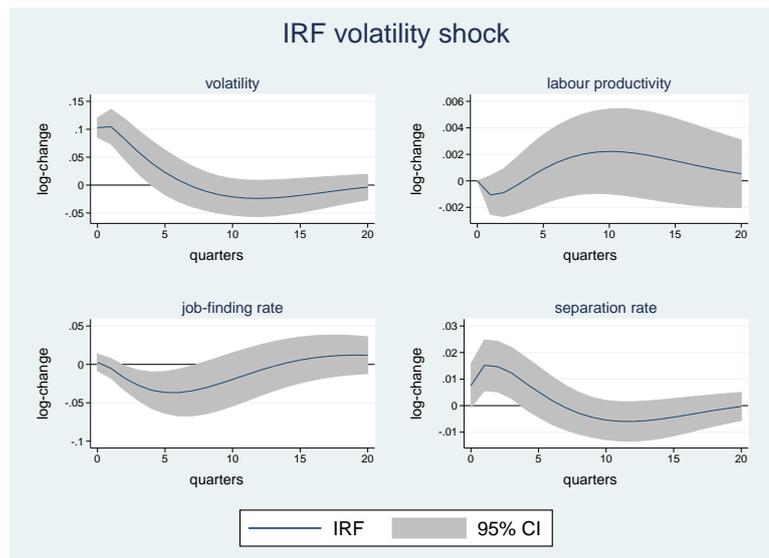


Figure 1.20: Impulse response functions after a volatility shock in US data. Instead of detrending using a linear trend, the series are detrended using the HP-filter with parameter $\lambda = 100,000$.

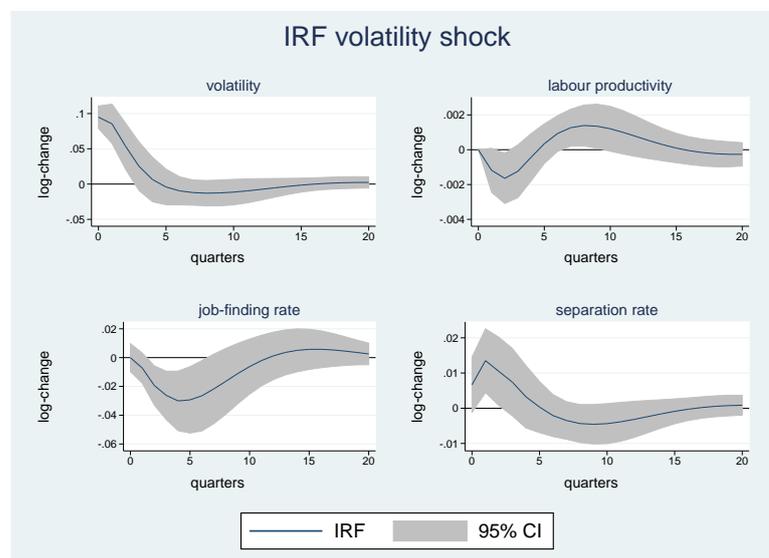


Figure 1.21: Impulse response functions after a volatility shock in US data. Instead of detrending using a linear trend, the series are detrended using the HP-filter with parameter $\lambda = 1,600$.

1.B Comparative statics

1.B.1 Sensitivity analysis

The purpose of this section is to show that the main result of this chapter is not sensitive to the choice of parameters. In the absence of creation costs ($H = 0$), the job-finding rate hardly reacts to changes in idiosyncratic volatility, and as the creation cost increases, the falling job-finding rate contributes increasingly to the fall in employment. This is shown in the lower right subplot of the comparative statics graphs in this section.

In the following exercises unless specified otherwise, the three parameters that are calibrated ($\bar{\sigma}_x$, ω_3 , and the mass of firms) are recalibrated such that the targets in steady state (unemployment rate, job-finding rate, filling cost relative to average wage) are met exactly.

Destruction cost instead of creation cost. In this exercise, there is no cost of creation, but a cost H has to be paid, whenever a position is shut down voluntarily. The results in figure 1.22 are almost identical to the cost paid upon creation. The reason is that firms are forward looking and take into account that they have to pay the destruction cost when they want to close the position in the future. As a result, the marginal value of a position shifts down and the three relevant values defining the cutoffs in equation (1.13) become $\left(-H, -H + (1 - \lambda) \frac{c}{q}, (1 - \lambda) \frac{c}{q}\right)$ instead of $\left(0, (1 - \lambda) \frac{c}{q}, H + (1 - \lambda) \frac{c}{q}\right)$.

Same absolute increase of idiosyncratic volatility. In this chapter, shocks to uncertainty are interpreted as relative changes of idiosyncratic volatility. The calibration strategy results in steady state volatility $\bar{\sigma}_x$ to be increasing in creation cost, because a higher creation cost makes firms more reluctant to fire. In order to meet the target for the separation rate, volatility has to be increased. This implies that a 10% increase of volatility increases volatility by a larger absolute amount when H is high, which explains the stronger response of employment. Figure 1.23 shows how the results differ when volatility is increased by 0.02, independent of H . Then the response of employment becomes almost flat, but the contribution of the job-finding rate to its change is hardly affected. Both the separation rate and the job-finding rate react less for large values of H .

Cyclical creation cost. Pissarides (2009) has shown that fluctuations of unemployment in the standard search and matching model are amplified if the total vacancy posting cost consists of a payment that is independent of market tightness. Without this assumption, higher market tightness in a boom makes it more costly for firms to fill their positions, which dampens the increase of vacancy posting. In my model, the cost of creating a position is assumed to be constant. This is not directly comparable with the assumption in Pissarides (2009) because the cost does not have to be paid when an empty position needs to be filled. The cost of finding workers in my model depends on market tightness as in a standard search and matching model. It is not clear whether the cost of creating (or destroying) positions should be higher when labour market tightness is high. If that were the case, higher costs of creating positions would dampen booms in addition to the higher filling cost. Note, however, that this does not affect the result that the job-finding rate does not move in response to uncertainty shocks when $H = 0$, because if market tightness does not react, it is irrelevant whether the cost of creating positions depends on market tightness. Figure 1.24 shows the results when the cost of creating positions is proportional to the filling cost $\frac{c}{q(\theta)}$.⁴³ The result is that for high values

⁴³Remember that in the initial steady state without aggregate uncertainty, $\frac{c}{q(\theta)}$ is normalized to be 1. Now, H

of H , the cost of creation falls relative to the baseline scenario. Consequently, firms become more willing to shut down positions, because it is cheaper to create them in the future and the separation rate increases. For the same reason, the job-finding rate is reduced because firms are less willing to fill their empty positions. As a result, the job-finding rate contributes less to the fall in employment, which becomes larger. Both lower employment and the reduced importance of the wait-and-see effect lead to an increase in measured labour productivity.

Filling cost. If it is more costly to fill positions, the wait-and-see effect becomes stronger even in the absence of a creation cost. Figure 1.25 shows that even doubling the target of the filling cost relative to average wages increases the contribution of the job-finding rate to less than 5% when $H = 0$. For larger values of H , the job-finding rate actually falls less than with the lower filling cost. The reason is that firms become more sensitive to changes in market tightness, because the filling costs are more important for them. This increases labour demand more in response to falling market tightness such that in general equilibrium the job-finding rate falls less.

Exogenous turnover. Figure 1.26 shows the solution for alternative values of the rate of obsolescence δ and quits λ . If both parameters are set to 0 as in Elsby and Michaels (2013), both the job-finding rate and the separation rate react more to higher uncertainty, and the job-finding rate contributes slightly less to changes in employment. In contrast, if λ was doubled, both labour market transition rates react less and the job-finding rate contributes more. Note that the calibration strategy ensures, that the separation rate is the same in steady state. This means that there are more *endogenous* separations when δ and λ are small. Then a shock that increases endogenous separations by the same relative amount, leads to a bigger increase of total separations.

Alternative wage equation. Figure 1.27 shows the results for alternative parameters in the wage equation (1.18). The red dashed solution halves the parameter ω_2 . This makes the wage less responsive to changes in market tightness. Therefore, market tightness and the job-finding rate need to fall more for high values of H to ensure market clearing. In the green dash-dotted solution, the parameter ω_1 is set to 1, which means that the wage does not depend on firm size. Increasing only ω_1 means that firms become more profitable as if they had a permanently higher productivity. This scales up their demand for labour and the average firm size. A corresponding decrease of the mass of firms could perfectly offset this, and the results would be identical to the baseline scenario. The only reason that they differ is that the wage coefficient ω_3 is recalibrated to meet the target for the filling cost relative to the wage. This has only minor consequences for the contribution of the job-finding rate to changes in employment.

Curvature of production function. In the literature of multi-worker firm models with decreasing returns to scale, a wide range of parameter values for the elasticity of the production or revenue function has been used. In the baseline scenario, I use an intermediate value of $\alpha = 0.75$. Figure 1.28 confirms that lowering it to 0.6 as in Elsby and Michaels (2013) or increasing it to 0.85 as in Schaal (2012) hardly affects the contribution of the job-finding rate to changes in employment. What α does affect is the response of employment, because it determines how much the marginal profit of the firm increases with fewer workers. When α is low, firms adjust their desired number of workers by less in response to shocks than for higher values of α .

Arrival rate of idiosyncratic shock. Figure 1.29 shows the results for different values of changes in proportion to changes in $\frac{c}{q(\theta)}$ in the new steady state.

the arrival rate of the idiosyncratic shock. Because the idiosyncratic productivity draws are uncorrelated in the baseline scenario, a higher arrival rate makes firms less responsive to their current level of productivity. Then they also react less to changes in idiosyncratic volatility. This primarily affects changes in employment, but it also reduces the contribution of the job-finding rate.

Distribution of idiosyncratic shock. Figure 1.30 shows that the contribution of the job-finding rate to changes in employment is not sensitive to the assumed distribution of idiosyncratic productivity. The red dashed curves show the result when a Pareto distribution instead of the log-normal distribution is used as in Elsby and Michaels (2013). The green dash-dotted curves assume that idiosyncratic productivity changes in each period, but it is autocorrelated as in Schaal (2012). In both cases, employment reacts less to uncertainty shocks, but the contribution of the job-finding rate increases slightly.

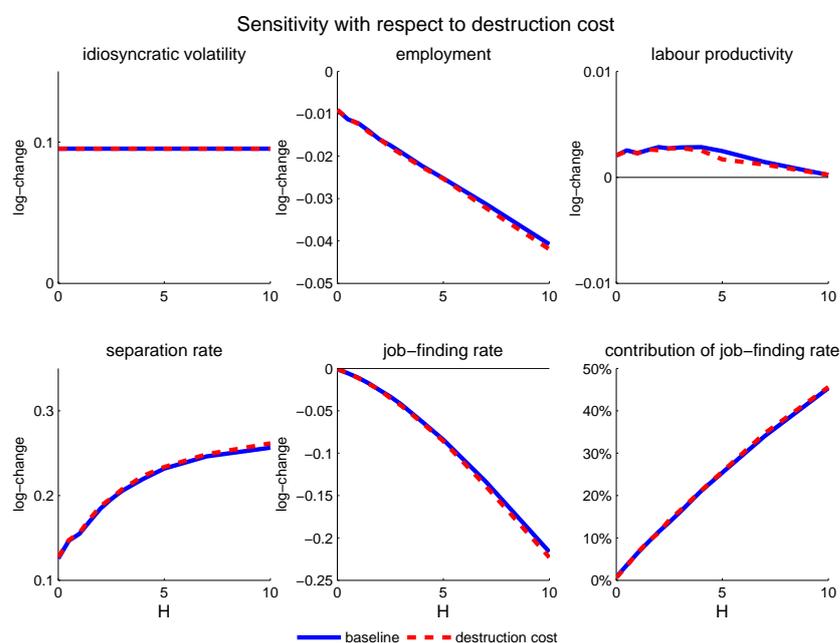


Figure 1.22: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the cost is not paid upon creation but upon voluntary destruction of a position.

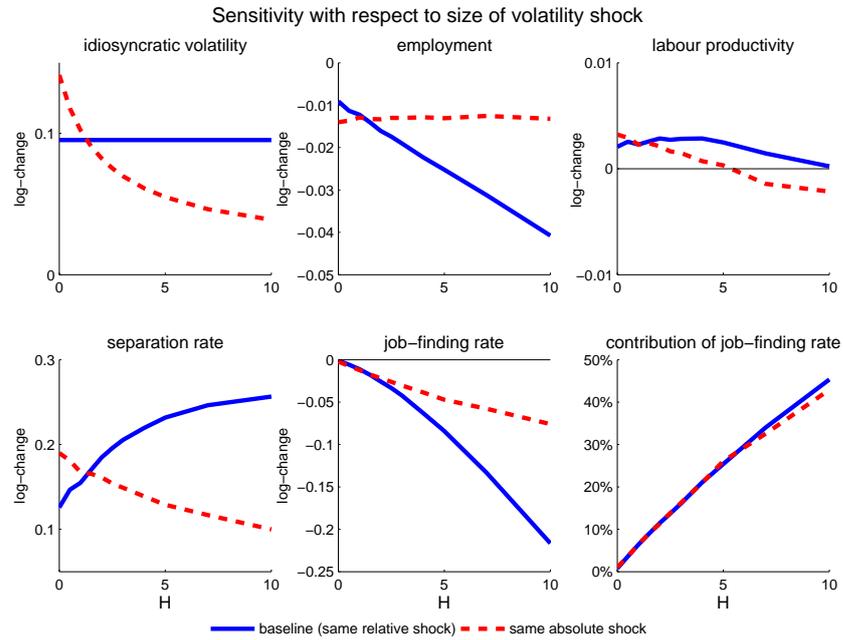


Figure 1.23: Responses after a permanent 10% increase compared with a 0.02 increase of idiosyncratic volatility for varying H .

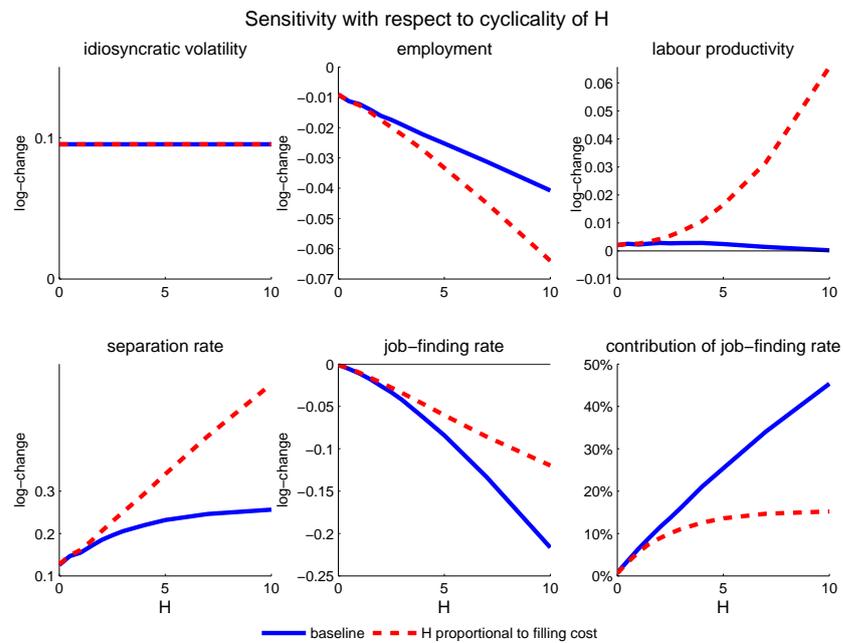


Figure 1.24: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the creation cost H is proportional to the filling cost $\frac{c}{q}$.

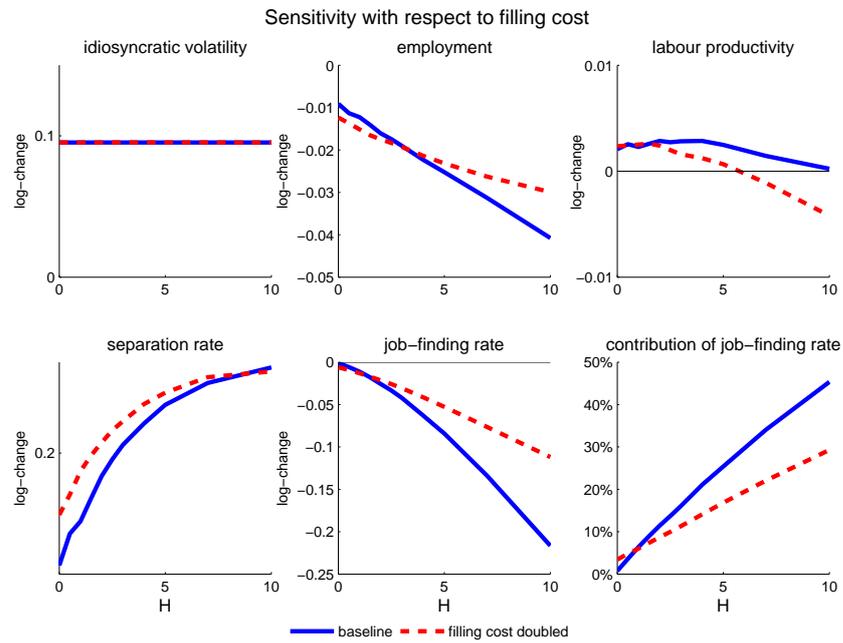


Figure 1.25: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the target for the filling cost is doubled relative to the baseline scenario. The creation cost remains unchanged, such that $H = 10$ now corresponds to the creation cost being 5 times the filling cost in steady state.

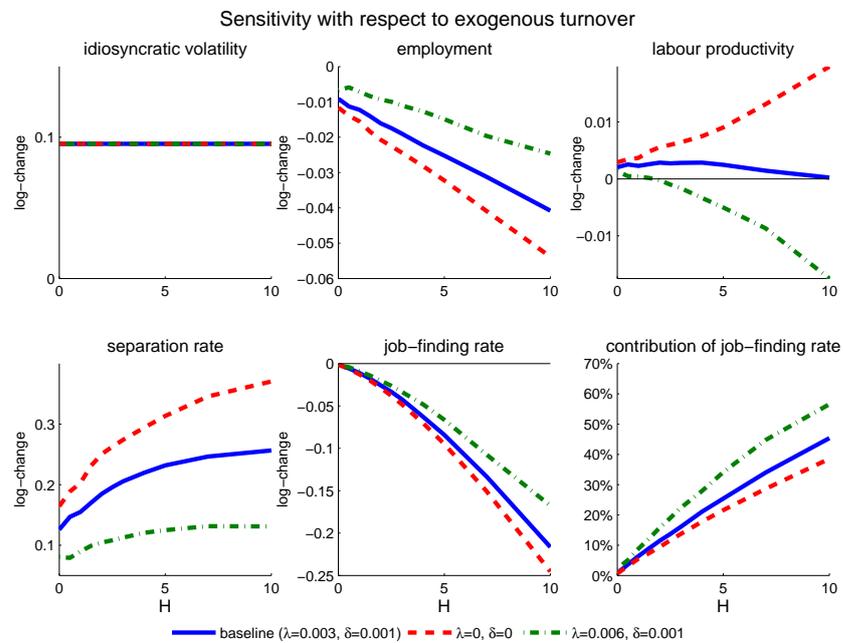


Figure 1.26: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution solves the model without depreciation of positions and without quits. The green dash-dotted solution solves the model for a doubled rate of quits.

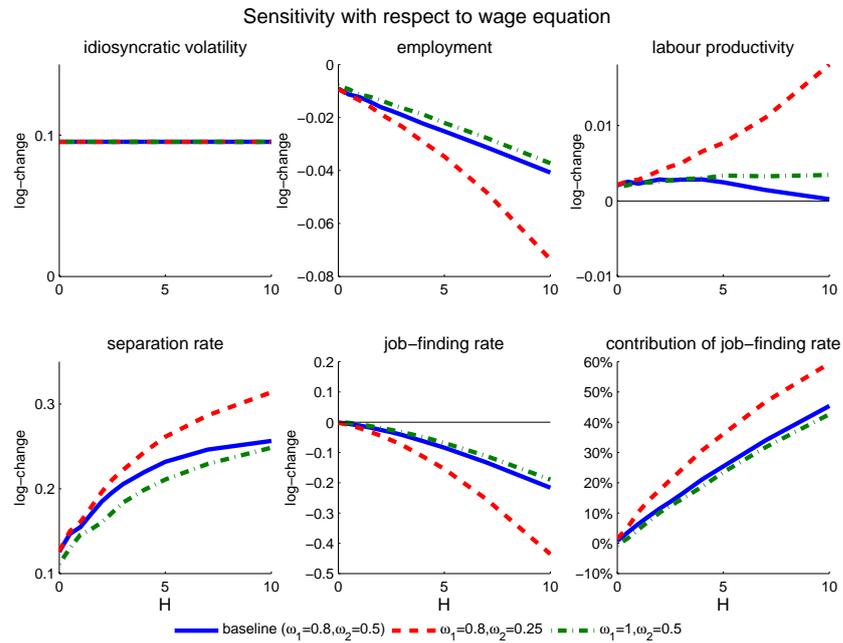


Figure 1.27: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution uses a wage equation that is less sensitive to changes in market tightness, and the green dash-dotted solution makes the wage independent of firm size.

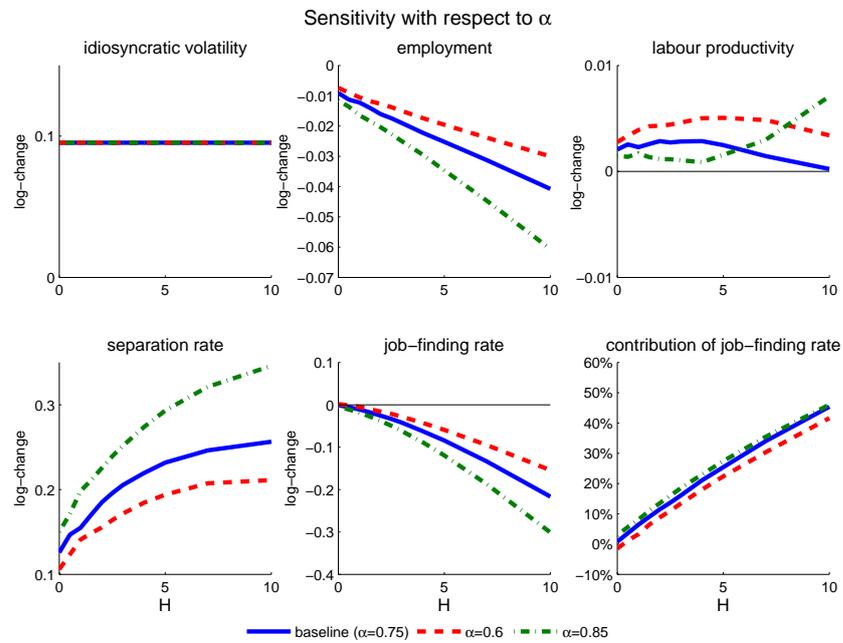


Figure 1.28: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The alternative solutions use different values of curvature in the production function.

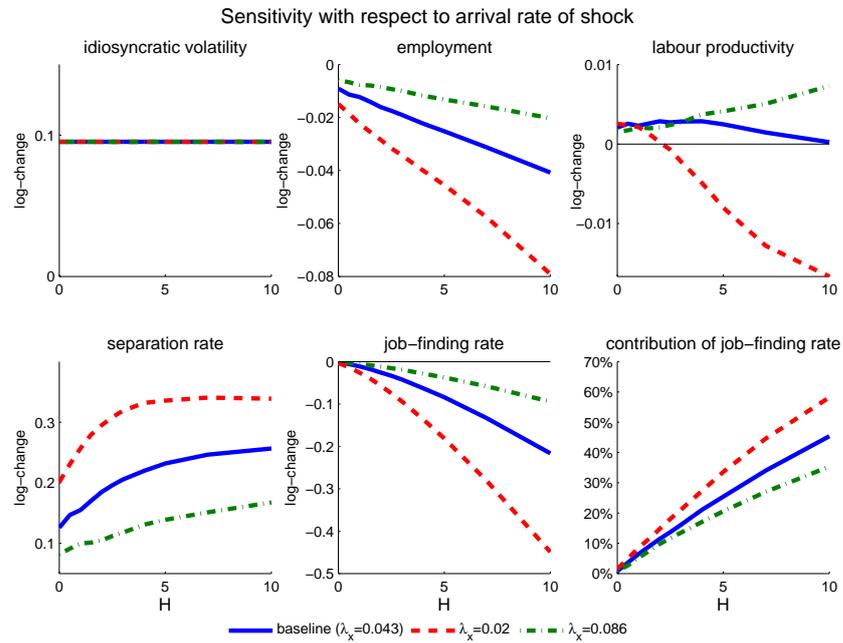


Figure 1.29: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solutions, the arrival rate of the idiosyncratic shock is varied.

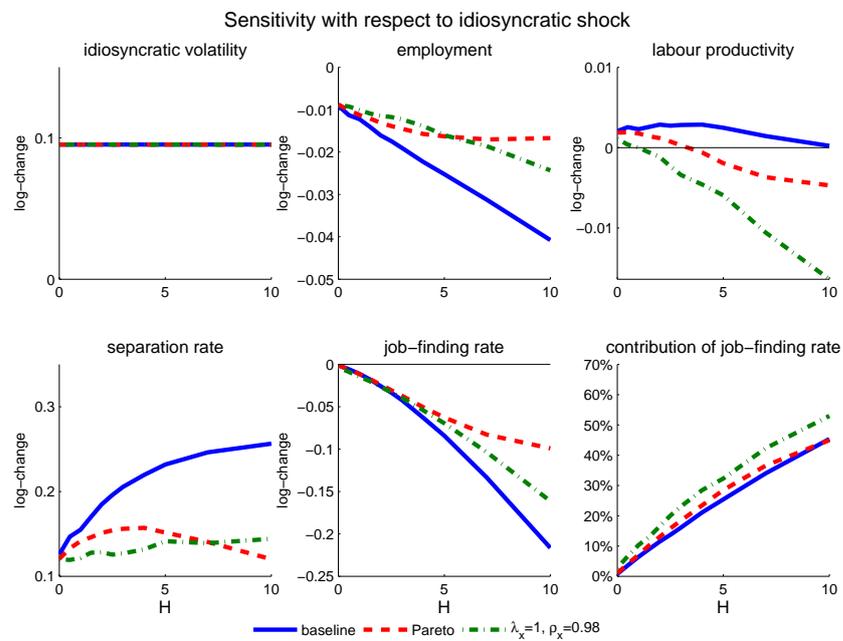


Figure 1.30: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution uses a Pareto distribution instead of the log-normal distribution for the idiosyncratic productivity draw. The green dash-dotted solution assumes that shocks arrive in each period but that productivity is autocorrelated.

1.C Decomposition exercise

1.C.1 Derivation of change in cutoffs for the wait-and-see effect

The goal is to compare how the cutoffs of firms at the same percentile of the idiosyncratic productivity distribution change when idiosyncratic volatility changes. It is not possible to just compare the change in cutoffs, because firms at the same percentile have a different value of x depending on the standard deviation of the cross-sectional distribution. In order to account for this mean effect, I will consider how the firm's cutoff changes relative to the "perfect insurance" solution. In this context, "perfect insurance" means the assumption that a firm can always sell an existing position or existing workers to other firms and thereby transfer the position or worker to this other firm. In equilibrium, other firms would be willing to pay exactly the cost that they would otherwise incur to create a position or to hire a worker, respectively. Hence, in terms of the marginal value of a position, firms are perfectly insured against idiosyncratic shocks. If they become less productive, they can sell some of their existing positions and workers to other firms until the marginal value of a position equals the price at which it could be sold. Then, equation (1.11) simplifies, and the number of positions is determined by

$$\omega_1 \alpha z x n^{\alpha-1} - (\omega_2 \theta + \omega_3) + \beta (1 - \delta) H + \beta (1 - \delta) (1 - \lambda) \mathbb{E} \left[\frac{c}{q(\theta(s'))} | s \right] = H + \frac{c}{q(\theta(s))}. \quad (1.45)$$

This states that the marginal product minus the marginal wage plus the discounted expected marginal value must be equal to the cost of creating and filling a position. The marginal value of a position created and filled today in the next period consists of the value of the position, $(1 - \delta) H$, taking into account the exogenous destruction of a share δ of positions, and the value of the worker filling it, $(1 - \delta) (1 - \lambda) \mathbb{E} \left[\frac{c}{q(\theta(s'))} | s \right]$, taking into account that a share λ of workers in not exogenously destroyed positions will leave the firm. Note that the only risk remaining for the firm is aggregate risk, which determines market tightness and hence the filling rate q' . For the remainder of this section, I will abstract from the effects expected changes in market tightness have on firms, and assume that $q' = q$. Then one can solve for the number of positions as a function of idiosyncratic productivity:

$$\log \tilde{n}_{LB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log \left\{ \omega_2 \theta + \omega_3 + [1 - \beta(1 - \delta)] H + [1 - \beta(1 - \delta)(1 - \lambda)] \frac{c}{q} \right\}}{1 - \alpha}. \quad (1.46)$$

This shows that the number of positions $\tilde{n}_{LB}(x)$ is log-linear in idiosyncratic productivity in the "perfect insurance" case. We can compare this solution to the creation cutoff $n_{LB}(x)$ in the model. In general, $n_{LB}(x) \leq \tilde{n}_{LB}(x)$ holds, because the continuation value in the model with idiosyncratic risk never exceeds the continuation value in the perfect insurance model. The difference between the creation cutoff and the perfect insurance number of positions can be interpreted as the amount of positions that a firm creates less in the presence of idiosyncratic risk, because it fears that the marginal value of the newly created positions falls in the future. One can show that $\lim_{x \rightarrow 0} (n_{LB}(x) - \tilde{n}_{LB}(x)) = 0$ holds, because for low values of idiosyncratic productivity it becomes increasingly unlikely that an even lower value of productivity is drawn in the next period. In particular, if there is a minimum value in the support of the distribution of idiosyncratic productivity, $n_{LB}(x_{\min}) = \tilde{n}_{LB}(x_{\min})$. Then, a firm knows for sure that productivity can't fall in the next period, and that the continuation value is equal to the creation and filling cost as in the perfect insurance economy. Figure 1.31 draws the perfect

insurance cutoff in addition to the creation cutoff for $H = 5$ and $H = 0$. Note that the difference between the two is larger for $H = 5$, because the continuation value could potentially fall by more. In the limiting case of no creation cost and no search frictions ($c = 0$), the two cutoffs would coincide as well.

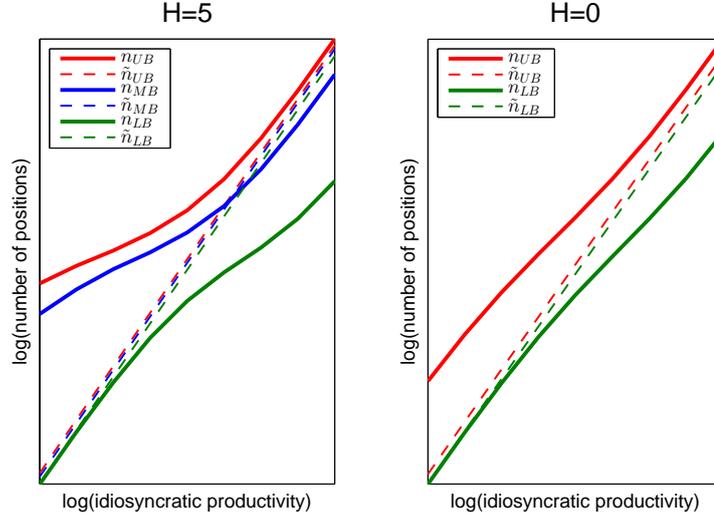


Figure 1.31: Optimal cutoffs (solid) and comparison cutoffs (dashed), that assume a fixed marginal continuation value.

Now, I can define analogous measures for the middle cutoff $n_{MB}(x)$, which determines whether a firm wants to fill empty positions, and for the upper cutoff $n_{UB}(x)$, which determines how many workers a firm wants to fire. According to equation (1.10), the middle cutoff is given by the condition

$$J(n_{MB}(x), x, s) = \omega_1 \alpha z x n_{MB}(x)^{\alpha-1} - (\omega_2 \theta + \omega_3) - \lambda \frac{c}{q(\theta(s))} + \dots \quad (1.47)$$

$$\begin{aligned} & \beta(1-\delta) \mathbb{E} [J((1-\delta)n_{MB}(x, s), x', s') | x, s] \\ &= (1-\lambda) \frac{c}{q(\theta(s))}. \end{aligned} \quad (1.48)$$

I want to compare it to a hypothetical cutoff, $\tilde{n}_{MB}(x)$, that assumes that the expected marginal continuation value is equal to the current period's marginal value:⁴⁴

$$\mathbb{E} [J((1-\delta)\tilde{n}_{MB}(x), x', s') | x, s] = J(\tilde{n}_{MB}(x), x, s). \quad (1.49)$$

Then the cutoff $\tilde{n}_{MB}(x)$ can be calculated:

$$\log \tilde{n}_{MB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log \left\{ \omega_2 \theta + \omega_3 + [1 - \beta(1-\delta)(1-\lambda)] \frac{c}{q(\theta(s))} \right\}}{1-\alpha}. \quad (1.50)$$

A comparison of this cutoff with the perfect insurance cutoff in equation (1.46) shows that it is obtained by a parallel upward shift in the $(\log x, \log n)$ -space. Note that the cutoff $n_{MB}(x)$ is above $\tilde{n}_{MB}(x)$ for low values of productivity and below it for high values of x . The reason is

⁴⁴Note that if both the aggregate and idiosyncratic state remain unchanged, $J((1-\delta)n_{MB}(x), x', s') > J(n_{MB}(x), x, s)$ holds because of depreciation ($\delta > 0$). The underlying assumption for the decomposition exercise is that this bias does not depend much on the size of idiosyncratic volatility.

that low-productivity firms are more likely to get a positive productivity shock, which leads to a continuation value above $(1 - \lambda) \frac{c}{q}$. Therefore, the expected continuation value for firms subject to idiosyncratic risk is higher than $(1 - \lambda) \frac{c}{q}$ in the definition of $\tilde{n}_{MB}(x)$. This implies that $n_{MB}(x) > \tilde{n}_{MB}(x)$ for low x . In contrast, high productivity firms are more likely to experience a negative productivity shock, that results in a continuation value below $(1 - \lambda) \frac{c}{q}$. Hence, the expected continuation value is below $(1 - \lambda) \frac{c}{q}$ for those firms and $n_{MB}(x) < \tilde{n}_{MB}(x)$. Finally, note that in the absence of creation costs the middle cutoffs coincide with the low cutoffs as shown in the right panel of figure 1.31.

The upper cutoff is implicitly defined by equation (1.6)

$$\begin{aligned} J(n_{UB}(x, s), x, s) &= (1 - \lambda) \left\{ \begin{aligned} &\omega_1 \alpha z x [(1 - \lambda) n_{UB}(x, s)]^{\alpha-1} - (\omega_2 \theta + \omega_3) + \\ &+ \beta (1 - \delta) \mathbb{E} [J((1 - \delta)(1 - \lambda) n_{UB}(x, s), x', s') | x, s] \end{aligned} \right\} \quad (1.51) \\ &= 0. \end{aligned}$$

Analogous to above, I define $\tilde{n}_{UB}(x)$ under the assumption that

$$\mathbb{E} [J((1 - \delta)(1 - \lambda) \tilde{n}_{UB}(x), x', s') | x, s] = 0. \quad (1.52)$$

Then the cutoff is given by

$$\log \tilde{n}_{UB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log(\omega_2 \theta + \omega_3)}{1 - \alpha} - \log(1 - \lambda). \quad (1.53)$$

This cutoff is parallel to $\tilde{n}_{MB}(x)$ and $\tilde{n}_{LB}(x)$ in the $(\log x, \log n)$ -space⁴⁵. In general, $n_{UB}(x) \geq \tilde{n}_{UB}(x)$, because the marginal continuation value that is relevant for $n_{UB}(x)$ is nonnegative, whereas it is 0 by assumption for $\tilde{n}_{UB}(x)$. The difference between the two cutoffs can be interpreted as the amount of workers that are not fired, and whose positions are not shut down, because the firm has a chance to get a positive productivity shock. The difference is getting smaller for high levels of productivity because it becomes less likely for the firm to become even more productive in the future⁴⁶.

Let $d_j^\sigma(p)$ be the log-difference between cutoffs n_j and \tilde{n}_j , evaluated at the $100p^{\text{th}}$ -percentile, when the standard deviation of idiosyncratic productivity is σ .⁴⁷ Denoting the inverse cumulative distribution function of the idiosyncratic shock by $F_{\sigma_x}^{-1}(\cdot)$, the three differences can be written as

$$d_{LB}^\sigma(p) = \log n_{LB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{LB}(F_{\sigma_x}^{-1}(p)) \leq 0, \quad (1.54)$$

$$d_{MB}^\sigma(p) = \log n_{MB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{MB}(F_{\sigma_x}^{-1}(p)), \quad (1.55)$$

$$d_{UB}^\sigma(p) = \log n_{UB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{UB}(F_{\sigma_x}^{-1}(p)) \geq 0. \quad (1.56)$$

⁴⁵In theory, $\tilde{n}_{UB}(x)$ could be smaller than $\tilde{n}_{MB}(x)$. In my calibration, λ is small enough relative $\frac{c}{q}$ such that $\tilde{n}_{UB}(x) > \tilde{n}_{MB}(x)$ holds.

⁴⁶In general it is not true that $\lim_{x \rightarrow \infty} (n_{UB}(x) - \tilde{n}_{UB}(x)) = 0$ holds. This is because in the presence of depreciation and attrition of workers, a firm's size shrinks over time. This increases marginal profits leading to a positive expected continuation value even if $J(n_{UB}(x_{\max}), x, s) = 0$ and $x' \leq x_{\max}$.

⁴⁷To keep notation simpler, I use subscript j , whenever an expression applies to all three cutoffs LB , MB , and UB .

1.C.2 Asymmetry of wait-and-see effect

The shifts of the cutoffs for the wait-and-see effect in figure 1.8 are not parallel. First, the effects are hump-shaped. They are smaller for p close to 0 or p close to 1, because then most new productivity shocks lead to the firm being at the lower or upper cutoff, respectively. For example, take a firm in a low percentile of current productivity. When a new shock arrives, it is likely to become significantly more productive. In that case, the firm is going to create positions after the good productivity draw even if it was currently (close to) firing workers. Thus the marginal continuation value of a position is likely to be at its upper bound after a new productivity draw. Increasing uncertainty still increases the expected continuation value by a bit, which leads to an increase of $d_{UB}^{\sigma}(p)$, but the effect is quantitatively small. Second, the change in cutoffs peaks at lower productivity values for the upper bound and at higher productivity values for the lower bound. The reason for this asymmetry is that $|d_{LB}^{\sigma}(p)|$ is larger for more productive firms as could be seen in figure 1.31. For those firms a future fall in productivity is more likely, which makes them more cautious in creating positions. Higher uncertainty exacerbates the potential losses and increases $|d_{LB}^{\sigma}(p)|$. In contrast, low-productivity firms are unlikely to get even less productive in the future. Therefore, higher uncertainty increases $|d_{LB}^{\sigma}(p)|$ by less. A similar argument can be applied for the upper cutoff: $d_{UB}^{\sigma}(p)$ is small for very productive firms, because they are unlikely to become even more productive. Therefore, changes in uncertainty affect $d_{UB}^{\sigma}(p)$ less. Note that this asymmetry increases the effects relative to a symmetric shift of the cutoffs, because there are more high-productivity firms close to the lower cutoff and more low-productivity firms close to the upper cutoff. Third, the effects are quantitatively stronger in the presence of creation cost. The larger the creation cost, the bigger becomes the difference between the upper and the lower bound of the continuation value, and the more likely it is for a firm to be in the interior region after a shock. This means that potential gains or losses from an increase in uncertainty become larger and firms adjust their cutoffs by more. As a result, the wait-and-see effect becomes more relevant when H is large.

1.D Comparison with US time series

This appendix provides additional graphs for the exercise matching the US time series in section 1.6.

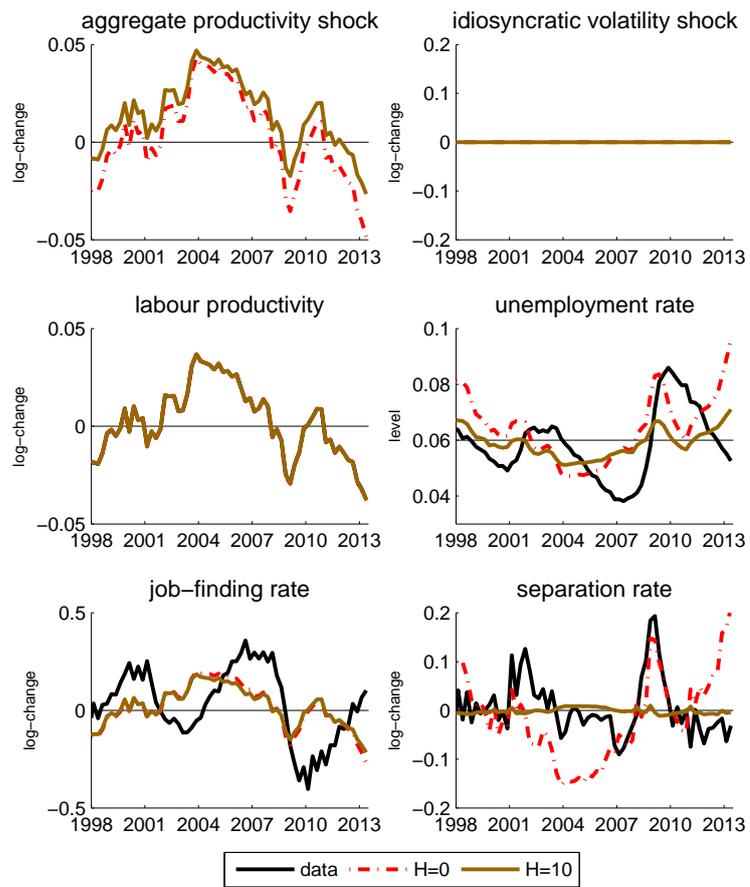


Figure 1.32: Model without uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data (linear trend removed). Relative to the baseline calibration, using $\omega_2 = 0.5$, a lower value of $\omega_2 = 0.2$ is used, which makes the wage less responsive to changes in market tightness.

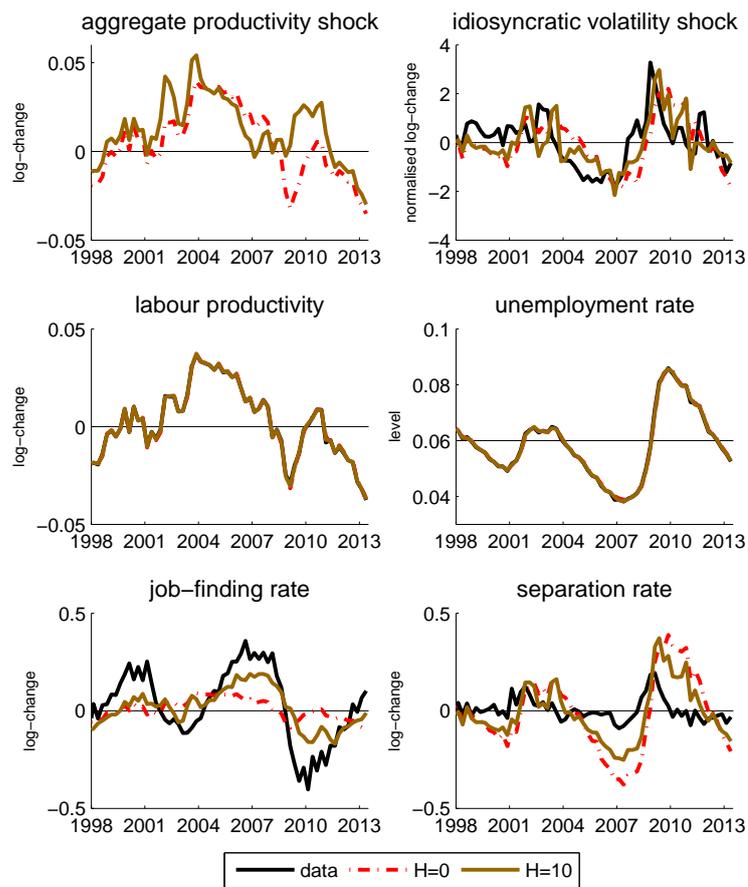


Figure 1.33: Model with aggregate productivity and uncertainty shocks: Aggregate productivity and idiosyncratic volatility are estimated in each quarter to match labour productivity and unemployment rate in the model to the data (linear trend removed). The top right panel compares the backed out volatility shocks with the empirical volatility index by normalising all series to have variance 1.

Chapter 2

Unemployment (Fears) and Deflationary Spirals: The Role of Sticky Nominal Wages

2.1 Introduction

The empirical literature documents that workers suffer substantial earnings losses and reductions in consumption levels during unemployment spells. For example, Kolsrud, Landais, Nilsson, and Spinnewijn (2015) document using Swedish data that consumption drops on average by 32% during the first year of an unemployment spell.¹ This observed inability to insure against unemployment spells has motivated several researchers to build business cycle models with incomplete markets. The hope (and expectation) has been that such models would not only generate more realistic behavior for individual variables, but also would be able to generate sufficiently volatile and prolonged business cycles without relying on large and persistent exogenous shocks hitting the economy. In these models, individual consumption is typically much more volatile than aggregate consumption, but the laws of motion for *aggregate* variables are often not that different from the outcomes in the corresponding representative-agent version. Exemplary papers are Krusell and Smith (1998) and Krusell, Mukoyama, and Sahin (2010). Moreover, McKay and Reis (2013) document that reductions in unemployment benefits actually *reduce* the volatility of aggregate consumption. The reason is that workers build up larger buffer stocks of assets when unemployment benefits are reduced, which makes the economy *as a whole* better equipped to smooth consumption.

We develop a model in which the inability to insure against unemployment risk generates business cycles which are much more volatile than the corresponding complete-markets (i.e., representative-agent) version in which all agents' consumption levels are identical in each and every period. This result is obtained by combining incomplete markets with incomplete adjustment of the nominal wage rate to changes in the price level.² The mechanism is strengthened by Diamond-Mortensen-Pissarides search frictions in the labor market.

¹See section 2.3.2 for a more detailed discussion of the empirical literature investigating the behavior of individual consumption during unemployment spells.

²Empirical motivation for this assumption is discussed in section 2.3.1.

Before explaining why the combination of incomplete markets and sticky nominal wages amplifies business cycles, we first explain why these features by themselves *dampen* business cycles in our model in which aggregate fluctuations are caused by productivity shocks. First, consider the model in which there are complete markets, but nominal wages do not respond one-for-one to price level changes. A negative productivity shock operates like a negative supply shock and increases prices, which—when nominal wages are sticky—would put downward pressure on real wages, which would *dampen* the reduction in profits caused by the direct negative effect of the reduction in productivity. Next, consider a model in which nominal wages are flexible, but workers cannot fully insure themselves against unemployment risk. Forward-looking agents understand that persistent negative productivity shocks increase their chances of being unemployed in the near future. If workers are not fully insured against this risk, then this would increase precautionary savings. In our general equilibrium model, this increased desire to save will lead to an increased demand for all assets including productive (and risky) assets such as firm ownership. This would *dampen* the reduction in the demand for productive assets triggered by the direct negative effect of reduced productivity levels.

Why does the combination of incomplete markets and sticky nominal wages lead to the opposite results? As mentioned above, the increase in the expected probability of being unemployed in the near future increases agents' desire to save more. Agents are particularly keen to increase savings by holding more of the liquid asset (money), which puts downward pressure on the price level, which in turn increases real wage costs and reduces profits. This effect dominates any positive effect that increased precautionary savings might have on the demand for productive investments. Once started, this channel will reinforce itself. That is, if precautionary savings lead—through downward (upward) pressure on prices (real wage costs)—to increased unemployment, then this in turn will lead to a further increase in precautionary savings. In our numerical examples, this downward spiral is (eventually) reversed by non-linearities in the matching market, that is, it becomes more attractive to invest in new firms when there are many unemployed workers searching for a job.

The property that workers are forward looking implies that concerns about *future* unemployment risk are preceded by increased real wage costs when nominal wages are sticky. The matching friction in the labor market magnifies this channel. The reason is that firms that foresee a future with higher real wage costs will reduce vacancy posting now and bring the concern about unemployment risk closer to the present.

We use our framework to study the advantages of alternative unemployment insurance (UI) policies and document that the effects of changes in unemployment benefits on the behavior of aggregate variables and on the wellbeing of workers differ from the effects in other models. For example, in the model of Krusell, Mukoyama, and Sahin (2010) most agents benefit from *reductions* in unemployment benefits even when benefits are reduced to very low levels. We consider an *increase* in unemployment benefits from the benchmark value of 50% to 55% of the prevailing wage rate and document that this increase in UI benefits improves the welfare of *all* agents when the switch occurs at the beginning of a recession and this is even true if wage rates adjust upwards to take into account the strengthened bargaining position of workers.³

There are quite a few factors that are important for this result. As a preview of the analysis, we mention here some factors that operate in our model and have not been emphasized in

³Wage increases have negative welfare consequences because they reduce job creation.

the literature. If unemployment benefits are increased during a recession, then this obviously benefits the unemployed directly. But the employed benefit too. Understandably, they benefit because they are better insured against future unemployment risk. Moreover, by reducing the negative downward spiral discussed above, the employed are less likely to be unemployed in the near future. Perhaps more surprisingly, they also benefit because the dampening of the downward spiral implies that the value of their equity holdings does not drop by that much relative to the case in which unemployment benefits are not increased.

As mentioned above, an increase in unemployment benefits has a negative effect on average investment, because agents save less if they are better insured. This leads to a reduction in the average employment rate. This channel is important in our model as well. In the version of our model *with* aggregate uncertainty, however, there are two quantitatively important factors that push average employment in the opposite direction and dominate the negative effect due to the reduction in precautionary savings. The first is that the demand for the productive asset increases, because asset prices are less volatile as increased unemployment benefits dampen the mechanism responsible for generating volatile business cycles. The second is that the non-linearity in the matching friction is such that increases in employment during booms are smaller than reductions during recessions. Consequently, a reduction in volatility would lead to an increase in average employment.

There is no interaction between the presence of sticky nominal wages and a *permanent* change in the level of unemployment benefits, except possibly during the transition phase. We document that there is a beneficial interaction when the change in the UI regime involves the introduction of a *countercyclical* component to UI benefits. An increase in UI benefits during recessions implies that precautionary saving does not increase by as much, which means that the price level does not drop by as much (and possibly even increases). Since this change in UI benefits and its induced effects on the price level are temporary, real wage costs are affected when nominal wages are sticky. In particular, there is less upward pressure on real wage costs, since prices drop by less. Another advantage of countercyclical benefits is that it does not increase workers' bargaining position by as much which also implies that there is less upward pressure on real wages.

In our model, precautionary savings can end up in the productive asset (firm ownership) and the unproductive asset (money). This complicates the analysis, because it means that the model simultaneously solves for agents' portfolio choices and equilibrium prices. Our numerical analysis ensures that the market for firm ownership (equity) is in equilibrium and all agents owning equity discount future equity returns with the correct, that is, their own individual-specific, marginal rate of substitution.⁴ By contrast, typical assumptions in the literature are that workers jointly own the productive asset at equal shares, that these shares cannot be sold, and that discounting of the returns of this asset occurs with some average marginal rate of substitution or a marginal rate of substitution based on aggregate consumption.^{5,6} Exceptions are Krusell, Mukoyama, and Sahin (2010) and Bayer, Lütticke, Pham-Dao, and Tjaden (2014), who—like us—allow trade in the productive asset and discount the agents' returns on this

⁴See section 2.2.7 for a detailed discussion.

⁵Examples are Shao and Silos (2007), Nakajima (2010), Gorneman, Kuester, and Nakajima (2012), Favilukis, Ludvigson, and Van Nieuwerburgh (2013), and Ravn and Sterk (2015).

⁶An alternative simplifying assumption is that the only agents who are allowed to invest in the productive asset are agents that are not affected by idiosyncratic risk (of any kind). Examples are Rudanko (2009), Bils, Chang, and Kim (2011), Challe, Matheron, Ragot, and Rubio-Ramirez (2014), and Challe and Ragot (2014).

asset with the correct marginal rate of substitution.⁷

In section 2.2, we describe the model. In section 2.3, we provide empirical motivation for the key assumptions underlying our model, sticky nominal wages and workers' inability to insure against unemployment risk. In this section, we also discuss the relationship between savings and idiosyncratic uncertainty. In section 2.4, we discuss the calibration of our model. In sections 2.5, and 2.6, we describe the behavior of individual and aggregate variables, respectively. In section 2.7, we discuss how business cycle behavior is affected by alternative UI policies.

2.2 Model

The economy consists of households, firms, and a government. The mass of firms is equal to q_t and all firms are identical. There is a unit mass of households and these differ in their employment status (employed or unemployed) and in their asset holdings. There are two types of assets, firm ownership (equity) and a liquid asset with a risk-free nominal value (money). In this section, we describe the behavior of the three different agents and the two asset markets.

Notation. Upper (lower) case variables denote nominal (real) variables. The subscript i refers to household i . Variables without a subscript i are either aggregate or variables that are identical across agents, such as prices.

2.2.1 Households

Each household consists of one worker who is either employed, $e_{i,t} = 1$, or unemployed, $e_{i,t} = 0$. The period- t budget constraint of household i is given by

$$\begin{aligned} P_t c_{i,t} + J_t (q_{i,t+1} - (1 - \delta) q_{i,t}) + M_{i,t+1} \\ = \\ (1 - \tau_t) W_t e_{i,t} + \mu (1 - \tau_t) W_t (1 - e_{i,t}) + D_t q_{i,t} + M_{i,t}, \end{aligned} \quad (2.1)$$

where $c_{i,t}$ denotes the consumption of household i , P_t the price of the consumption good, $M_{i,t}$ the amount of the liquid asset held at the *beginning* of period t (and chosen in period $t - 1$), τ_t the tax rate on wage income, W_t , the nominal wage rate, and μW_t the level of unemployment benefits. The variable $q_{i,t}$ is the amount of equity held at the beginning of period t . These ownership shares pay out nominal dividends D_t . Each period, a fraction δ of all firms go out of business and this leads to a corresponding loss in equity.⁸ When $q_{i,t+1} - (1 - \delta) q_{i,t}$ is larger (smaller) than zero, then the worker is buying (selling) equity and the nominal value of this transaction is equal to $J_t (q_{i,t+1} - (1 - \delta) q_{i,t})$, where J_t is the nominal price of equity.

Households are not allowed to take short positions in equity, that is

$$q_{i,t+1} \geq 0. \quad (2.2)$$

⁷The procedure in Krusell, Mukoyama, and Sahin (2010) is exact if the aggregate shock can take on as many realizations as there are assets and no agents are at the constraint. Our procedure does not require such restrictions, which is important, because the fraction of agents at the constraint is nontrivial in our model.

⁸In our benchmark model, we assume that households hold a diversified portfolio of equity, which means that all investors face the same loss, δ .

The household maximizes the following objective function:⁹

$$\mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j \left(\left(\frac{c_{i,t+j}^{1-\gamma} - 1}{1-\gamma} \right) + \chi \frac{\left(\frac{M_{i,t+1+j}}{P_{t+j}} \right)^{1-\zeta} - 1}{1-\zeta} \right) \right].$$

The first-order conditions are given by the budget constraint, the short-sale constraint and

$$c_{i,t}^{-\gamma} = \beta \mathbb{E}_t \left[c_{i,t+1}^{-\gamma} \frac{P_t}{P_{t+1}} \right] + \chi \left(\frac{M_{i,t+1}}{P_t} \right)^{-\zeta}, \quad (2.3)$$

$$c_{i,t}^{-\gamma} \geq \beta \mathbb{E}_t \left[c_{i,t+1}^{-\gamma} \left(\frac{D_{t+1} + (1-\delta)J_{t+1}}{J_t} \right) \frac{P_t}{P_{t+1}} \right], \quad (2.4)$$

$$0 = q_{i,t+1} \left(c_{i,t}^{-\gamma} - \beta \mathbb{E}_t \left[c_{i,t+1}^{-\gamma} \left(\frac{D_{t+1} + (1-\delta)J_{t+1}}{J_t} \right) \frac{P_t}{P_{t+1}} \right] \right). \quad (2.5)$$

In this model, agents hold money to facilitating transactions, but also to insure themselves against unemployment risk. Telyukova (2013) documents that households hold on average 50% more liquidity than they spend on average each month. The utility specification implies that agents will always choose positive real money balances. That is, there is no borrowing. Short positions in the liquid asset would become possible if the argument of the utility function is equal to $(M_{i,t} + \Phi)/P_t$ with $\Phi > 0$ instead of $M_{i,t}/P_t$. At higher values of Φ , agents can take larger short positions in money and are, thus, better insured against unemployment risk. Increases in χ —while keeping Φ equal to zero—have similar implications, since higher values of χ imply higher levels of financial assets.

2.2.2 Firms

An existing firm produces z_t , which is identical across firms. The value of z_t follows a first-order Markov process with a low (recession) and a high (boom) value. The partition into a recession and a boom regime simplifies characterizing model properties.¹⁰

There is one worker attached to each firm. Thus, the number of firms, q_t , is equal to the economy-wide employment rate. The nominal wage rate, W_t , is the only cost to the firm. Consequently, nominal firm profits, D_t , are given by

$$D_t = P_t z_t - W_t. \quad (2.6)$$

The nominal wage rate is set according to

$$W_t = \omega_0 \left(\frac{z_t}{\bar{z}} \right)^{\omega_z} \bar{z} \left(\frac{P_t}{\bar{P}} \right)^{\omega_p} \bar{P}, \quad (2.7)$$

where \bar{z} is the average productivity level, P_t is the price level, and \bar{P} is the average price level. Since the focus of this paper is on the responsiveness of nominal wages, W_t , to nominal prices, P_t , we need a wage setting rule with which we can vary this responsiveness. If $\omega_p = 1$

⁹If money and consumption enter the utility function additively, then money does not enter the Euler equation of other assets directly. This is consistent with the empirical results in Ireland (2004) who finds that the coefficient of money in the IS curve is very small (the standard error is quite sizable, however, so other specifications cannot be ruled out).

¹⁰Although we can solve the model for richer processes, this simple specification for z_t clearly helps in keeping the computational burden manageable.

($\omega_P < 1$), then wages adjust fully (partially) to changes in P_t . The coefficient ω_0 indicates the fraction of output that goes to the worker when z_t and P_t take on their average values and the coefficient ω_z indicates the sensitivity of the wage rate to changes in productivity. The value of ω_z controls how wages vary with business cycle conditions. This sensitivity is a key question in the labor search literature. In particular, Hall and Milgrom (2008) argue that the popular Nash bargaining framework makes wages cyclical by making the relevant reference point the value of being unemployed.¹¹

2.2.3 Government

The government taxes wages to finance unemployment benefits. Since the level of unemployment benefits is equal to a fixed fraction of the wage rate and taxes are proportional to wage income, the government's budget constraint can be written as

$$\tau_t q_t W_t = (1 - q_t) \mu (1 - \tau_t) W_t. \quad (2.8)$$

From this equation we get an expression for taxes, τ_t , which only depend on the employment rate. That is,

$$\tau_t = \mu \frac{1 - q_t}{q_t + \mu (1 - q_t)}. \quad (2.9)$$

An increase in q_t means that there is an increase in the tax base and a reduction in the number of unemployed. Both lead to a reduction in the tax rate.

2.2.4 Firm creation and equity market¹²

Agents that would like to increase their equity position in firm ownership, i.e., agents for whom $q_{i,t+1} - (1 - \delta) q_{i,t} > 0$, can buy equity at price J_t from agents that would like to sell equity, i.e., from agents for whom $q_{i,t+1} - (1 - \delta) q_{i,t} < 0$. Alternatively, agents who would like to buy additional equity can also acquire new firms by creating them. The creation of new firms requires making an investment $v_{i,t}$. How many new firms are created with a certain investment level depends on the number of unemployed workers, u_t , and the *aggregate* amount invested, v_t . In particular, the total number of new firms created is equal to

$$h_t \equiv q_t - (1 - \delta) q_{t-1} = \psi v_t^\eta u_t^{1-\eta} \quad (2.10)$$

and an individual investment of $v_{i,t}$ results in $(h_t/v_t)v_{i,t}$ new firms. In equilibrium, the cost of creating one new firm, v_t/h_t , has to be equal to the real market price, J_t/P_t , since new firms

¹¹Under Nash bargaining, workers' wages vary with their individual wealth level, which would increase the computational burden. One could question whether this is an empirically relevant feature. Moreover, the results in Krusell, Mukoyama, and Sahin (2010) indicate that this complication may only have a substantial effect on the wages of the poorest agents.

¹²Our representation of the matching market looks somewhat different than usual. As documented in appendix 2.C, however, it is identical to the standard setup. Our way of "telling the story" has three advantages. First, in our model, there is only one type of investor, namely the household. That is, we do not have zombie entrepreneurs who fulfill a crucial role in the standard setup, but do not get any positive net benefits out of this. Second, all agents in our economy have access to the same two assets, namely firm ownership and money. By contrast, households and entrepreneurs have different investment opportunities in the standard setup. Third, we have one parameter less and avoid the feature of the standard setup that different combinations of the vacancy posting cost and the scalings coefficient of the matching function lead to identical results for all variables except the level of vacancies, for which the data does not give good guidance anyway.

are identical to existing firms. Setting v_t/h_t equal to J_t/P_t and using equation (2.10) gives

$$v_t = \left(\psi \frac{J_t}{P_t} \right)^{1/(1-\eta)} u_t, \quad (2.11)$$

that is investment in new firms/jobs is increasing in J_t and increasing in the mass of workers looking for a job.

Equilibrium in the equity market requires that the supply of equity is equal to the demand of equity, that is,

$$\begin{aligned} h_t + \int \int \int_{e_i, q_i, M_i} ((1-\delta)q_i - q(e_i, q_i, M_i; s_t))_+ dF_t(e_i, q_i, M_i) \\ = \\ \int \int \int_{e_i, q_i, M_i} (q(e_i, q_i, M_i; s_t) - (1-\delta)q_i)_+ dF_t(e_i, q_i, M_i), \end{aligned} \quad (2.12)$$

where $(x)_+$ equals x when $x \geq 0$ and equals 0 when $x < 0$. Also, $F_t(e_i, q_i, M_i)$ is the period- t cumulative distribution function of the cross-sectional distribution of the three individual state variables: the employment state, e_i , money holdings, M_i , and equity holdings, q_i . The variable s_t denotes the set of aggregate state variables and its elements are discussed in section 2.2.6.

Combining the last three equations gives

$$\psi^{1/(1-\eta)} \left(\frac{J_t}{P_t} \right)^{\eta/(1-\eta)} u_t = \int \int \int_{e_i, q_i, M_i} (q(e_i, q_i, M_i; s_t) - (1-\delta)q_i) dF_t(e_i, q_i, M_i). \quad (2.13)$$

The supply of new equity (the left-hand side) is increasing in J_t and the net demand for equity (right-hand side) is decreasing in J_t .

2.2.5 Money market

Similarly, equilibrium in the market for money holdings requires that the net demand of households wanting to increase their money holdings is equal to the net supply of households wanting to decrease their money holdings. That is,

$$\begin{aligned} \int \int \int_{e_i, q_i, M_i} (M(e_i, q_i, M_i; s_t) - M_i)_+ dF_t(e_i, q_i, M_i) \\ = \\ \int \int \int_{e_i, q_i, M_i} (M_i - M(e_i, q_i, M_i; s_t))_+ dF_t(e_i, q_i, M_i). \end{aligned} \quad (2.14)$$

In section 2.7.2, we describe how liquidity injections would affect model outcomes and whether central banks are likely to pursue such policies.

2.2.6 Equilibrium and model solution

In equilibrium, the following conditions hold: (i) asset demand is determined by the households' optimality conditions, (ii) the cost of creating a new firm equals the market price of an existing firm, (iii) the demand for equity from households that want to buy equity equals the creation of new firms plus the supply of equity from households that want to sell equity, (iv) the demand for the liquid assets from households that want to increase their holdings is equal to the supply from households that want to reduce their holdings and (v) the government's budget constraint is satisfied.

The state variables for agent i are his asset holdings, his employment status, and the aggregate state variables. The latter consist of the aggregate productivity level, z_t , and the cross-sectional joint distribution of employment status and asset holdings. We use an algorithm similar to the one used in Krusell and Smith (1998) to solve for the laws of motion of aggregate variables. Details on the numerical procedure are given in appendix 2.B.1.

2.2.7 Discounting firm profits correctly with heterogeneous ownership

With incomplete markets and heterogeneous firm ownership, the question arises how to discount future firm profits. In our model, each and every firm owner discounts firm profits correctly. The reason is that agents can buy and sell equity. This means that the Euler equation for equity holds with equality for all investors holding equity, which implies that all firm owners discount the proceeds of the equity investment with the correct, i.e., their own individual-specific, marginal rate of substitution.¹³ In our model (and in our numerical algorithm) the market price and quantities are determined simultaneously such that both the equilibrium condition and each agent's Euler equation are satisfied. In our model, agents can choose to invest in the liquid, less risky, and unproductive asset and in the riskier and productive asset. Models with heterogeneous agents often assume that agents can only trade in the unproductive asset and assume that there is some form of communal ownership of the productive asset with ownership shares that remain fixed through time.¹⁴ Investment decisions in the productive asset are then based on an Euler equation using a marginal rate of substitution based on aggregate consumption, an average of the marginal rate of substitution of all agents, or risk neutral discounting. Another approach is to assume that investments in the productive asset are made by another type of agents who is not affected by uncertainty.¹⁵ Both approaches simplify the analysis a lot, but both direct any possible consequences of precautionary savings induced by idiosyncratic risk towards the unproductive asset only.

A long outstanding and unresolved debate in corporate finance deals with firm decision making when owners are heterogeneous and markets are incomplete. This is not an issue here, because firms do not make decisions with intertemporal consequences.¹⁶ If firms had to

¹³Krusell, Mukoyama, and Sahin (2010) also describe a procedure to discount firm profits (almost) correctly. They assume that the number of assets is equal to the number of realizations of the aggregate shock. Firm profits can then be discounted with the prices of the two corresponding contingent claims and this would be exactly correct if borrowing or short-sell constraints are not binding for any investor. Our procedure allows investors to be constrained and the number of realizations of the aggregate shock can exceed the number of assets.

¹⁴Examples are Shao and Silos (2007), Nakajima (2010), Gorneman, Kuester, and Nakajima (2012), Favilukis, Ludvigson, and Van Nieuwerburgh (2013), and Ravn and Sterk (2015).

¹⁵Examples are Rudanko (2009), Bills, Chang, and Kim (2011), Challe, Matheron, Ragot, and Rubio-Ramirez (2014), and Challe and Ragot (2014).

¹⁶Note that firm creation is a static decision and all agents in the economy would compare the cost of creating one firm, v_t/h_t , and its market value, J_t/P_t , in the same way.

make such decisions, we would have to deal with this challenging issue and specify how firm decisions are made. Conditional on this specification, however, our approach can still be used and firm owners would still discount firm revenues correctly.¹⁷

2.3 Empirical motivation

In this section, we discuss some key empirical observations that motivate our analysis. First, we discuss the evidence in favor of sticky nominal wages and whether that has or has not affected wage costs during the recent economic downturn. Second, we discuss the inability of individuals to insure themselves against unemployment spells. Third, we discuss whether savings respond to an increase in idiosyncratic uncertainty. The discussion mainly highlights the behavior of key Eurozone variables during the recent financial crisis, although we will also discuss evidence from other periods and countries outside the Eurozone. Details on the data sources are given in appendix 2.A.

2.3.1 Deflationary pressure and sticky nominal wages

In our heterogeneous-agent model, precautionary savings motives put upward pressure on the demand for money, which in turn puts downward pressure on prices. If nominal wages do not fully respond to changes in prices, then this puts upward pressure on real unit wage costs during recessions.

There are four elements in this story. First, there is downward pressure on prices.¹⁸ Second, nominal wages do not fully adjust for inflation. Third, real unit wage costs increase, that is, upward pressure on real wages is not offset by increases in labor productivity.¹⁹ Fourth, the increase in wage costs is also relevant for new jobs. These four elements are discussed next.

Deflationary pressure. Our paper focuses on recessions during which households' inability to fully insure themselves against increased idiosyncratic risk increases households' desire to save, which puts downward pressure on prices. The top panel of figure 2.1 plots the GDP deflator for the Eurozone (18 countries as of 2014) together with the time path that the deflators would have followed if inflation during the period following the fourth quarter of 2007 had been equal to the average inflation rate over the five preceding years. The figure shows that the growth in the price level slowed considerably during the crisis.²⁰

¹⁷This remains a tricky problem, even if the firms' objective function is given and all firms have the same objective. For example, suppose that all firms maximize current market value. Identical firms could then very well end up making different decisions. To see why, suppose that all firms make the same intertemporal decision. By deviating and providing different future payoff realizations, a firm can create value by "completing the market". As discussed in Ekern and Wilson (1974), if firms decisions do not alter the set of returns available to the whole economy, then investors can "undo" the effects of firm decisions on the payoffs of their individual portfolio. Consequently, investors would agree on what choices the firm should make. Carceles-Poveda and Coen-Pirani (2009) show that this happens in their model in which firms have constant return to scale technology and there are no binding borrowing constraints.

¹⁸Our story does not require prices to be procyclical. That is, the channel we identify is also present when the precautionary motive only dampens countercyclical behavior.

¹⁹Our model has ambiguous predictions for the cyclicity of real wages. If nominal wages respond little to lower inflation *and* little to lower productivity, then it is possible that the real wage rate increases during a recession. In our benchmark model, real wages initially increase following a negative productivity shock and then start to decrease.

²⁰Remarkable deflationary pressure is also visible in the US consumer price index (CPI). It dropped by 3.4% during the period from September 2008 to December 2008.

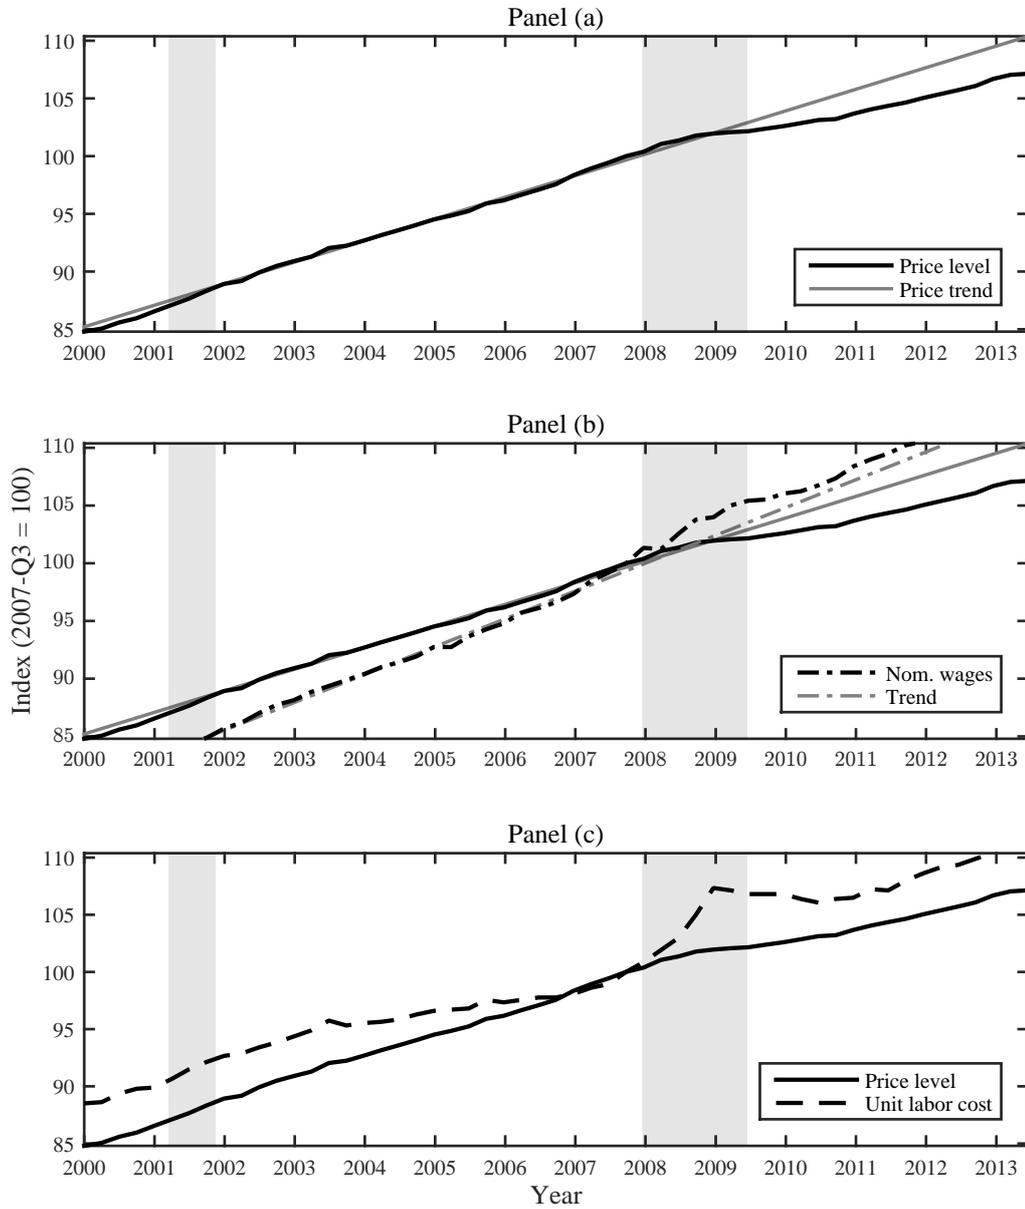


Figure 2.1: Key Eurozone variables before and after the financial crisis.

Notes. Panel (a) illustrates the GDP deflator for the Eurozone (18 countries as of 2014) together with its pre-crisis trend. Panel (b) illustrates nominal hourly earnings, the GDP deflator, and their associated pre-crisis trends. Panel (c) illustrates nominal unit labor costs together with the GDP deflator.

Source: OECD.

Nominal wage stickiness and inflation. There are many papers that document that nominal wages are sticky.²¹ Important for our paper is the question to what extent nominal wages adjust to aggregate shocks and in particular to changes in the aggregate price level. Well suited for our purpose is Druant, Fabiani, Kezdi, Lamo, Martins, and Sabbatini (2009) which provides survey evidence for a sample of European firms with a focus on the wages of the firms' main occupational groups; these would not change for reasons such as promotion. Another attractive feature of this study is that it explicitly investigates whether nominal wages adjust to inflation. In their survey, only 29.7% of Eurozone firms indicate that they have an internal policy of taking inflation into account when setting wages and only half of these firms do so using automatic indexation. Moreover, most firms that take inflation into account are backward looking. Both findings imply that real wages increase (or decrease by less) when inflation rates fall. Even though there are not that many firms that have an automatic policy in place, inflation is the prevalent factor triggering wage adjustment. In particular, around 50% of Eurozone firms adjust wages to inflation yearly and around 6% do this more frequently.²²

Papers that document nominal wage rigidity typically highlight the importance of *downward* nominal wage rigidity. Suppose there is downward, but no upward nominal wage rigidity. Does this imply that all nominal wages respond fully to changes in aggregate prices as long as aggregate prices increase? The answer is no. The reason is that firms are heterogeneous and a fraction of firms can still be constrained by the inability to adjust nominal wages downward. In fact, downward nominal wage rigidity is supported by the empirical finding that the histogram of firms' nominal wage changes has a large peak at zero.²³ The fraction of firms that is affected by the constraint that nominal wages cannot be adjusted downward would be increasing if the aggregate price level increases by less. In fact, Daly, Hobijn, and Lucking (2012) document that the fraction of US workers with a constant nominal wage increased from 11.2% in 2007 to 16% in 2011, whereas the fraction of workers facing a reduction in nominal wages was roughly unchanged.²⁴ This indicates that there is upward pressure on real wages when the inflation rate falls even if it remains positive and nominal wages are only rigid downward.

To investigate whether nominal wages followed the slowdown in inflation, we plot in the second panel of figure 2.1 nominal hourly earnings together with the GDP deflator. The figure also plots the realizations of both variables if they would have grown at a rate equal to the average observed in the five years before the crisis. We find that nominal wages continued to grow at pre-crisis rates, despite a substantial reduction in inflation rates. This means that real wages actually *increased* relative to trend.²⁵

Real wage costs. The observed increases in real wages are not necessarily due to a combination of low inflation and downward nominal wage rigidity. It is possible that solid real

²¹See, for example, Barattieri, Basu, and Gottschalk (2010), Dickens, Goette, Groshen, Holden, Messina, Schweitzer, Turunen, and Ward (2007), Daly, Hobijn, and Lucking (2012), Daly and Hobijn (2013), and Druant, Fabiani, Kezdi, Lamo, Martins, and Sabbatini (2009).

²²The actual number may be a bit higher, since roughly 30% of firms select as their answer the residual category, which includes "don't know" in addition to "never".

²³See Barattieri, Basu, and Gottschalk (2010), Dickens, Goette, Groshen, Holden, Messina, Schweitzer, Turunen, and Ward (2007), Daly, Hobijn, and Lucking (2012), and Daly and Hobijn (2013).

²⁴Similarly, at <http://nadaesgratis.es/?p=39350>, Marcel Jansen documents that from 2008 to 2013 there was a massive increase in the fraction of Spanish workers with no change in the nominal wage. There is some increase in the fraction of workers with a decrease in the nominal wage, but this increase is small relative to the increase in the spike of the histogram at constant nominal wages.

²⁵Similarly, Daly, Hobijn, and Lucking (2012), Daly and Hobijn (2013), and Rendahl (2012) document that real wages increased during the recent recession in the US.

wage growth reflects an increase in labor productivity, for example, because workers that are not laid off are more productive than those workers that left employment. To shed light on this possibility, we compare the nominal unit wage cost with the price level.²⁶ The results are shown in the bottom panel of figure 2.1. The figure shows that nominal unit labor costs have grown faster than prices in the private sector since the onset of the crisis, whereas the opposite was true before the crisis. This indicates that real labor costs increased during the crisis even if one corrects for productivity.²⁷

The observations are consistent with the hypothesis that the combination of deflationary pressure and nominal wage stickiness increased wage costs. In principle, it is still possible that nominal wages in the Eurozone did respond fully to prices. However, in that case, it must be true that the reduction in employment is mainly due to an outflow of workers that earn low wages *and* could produce at low real unit labor cost, since both real wages and real unit labor costs increased. That is, it must be the case that the workers who left employment were the ones who had a wage that was low relative to their productivity. That does not seem a very plausible possibility.

Wages of new and existing relationships. What matters in labor market matching models is the flexibility of wages of newly hired workers. Haefke, Sonntag, and van Rens (2013) argue that wages of new hires respond almost one-to-one to changes in labor productivity. Gertler, Huckfeldt, and Trigari (2014) argue that this result reflects changes in the composition of new hires and that—after correction for such composition effects—the wages of new hires are not more cyclical than wages of existing workers. More importantly, what matters for our paper is whether nominal wages respond to changes in the price level and this question is not addressed in either paper. As mentioned above, Druant, Fabiani, Kezdi, Lamo, Martins, and Sabbatini (2009) find that many firms do not adjust wages to inflation. Since their survey evidence focuses on firms' main occupational groups one would think that their results apply to new as well as old matches.

2.3.2 Inability to insure against unemployment risk

An important feature of our model is that workers cannot fully insure against unemployment risk. They are limited to do so because there are only two assets, neither of which has individual-specific payoffs. Moreover, short positions are not allowed. Consequently, workers' consumption levels drop when agents become unemployed.

Using Swedish data, Kolsrud, Landais, Nilsson, and Spinnewijn (2015) document that expenditures on consumption goods drop sharply during the first year of an unemployment spell after which they settle down at a level that is 34% lower than the pre-displacement level. This sharp drop is remarkable given that Sweden has quite generous unemployment benefits. As discussed in section 2.4, one reason is that the amount of assets workers hold at the start of an unemployment spell is low. Another reason is that average borrowing actually *decreases* during observed unemployment spells.

²⁶The nominal unit wage cost is defined as the cost of producing one unit of output, i.e., the nominal wage rate divided by labor productivity. The price index used as comparison is the price index used in defining labor productivity.

²⁷The observation that real unit labor costs are not constant over the business cycle is interesting in itself. If the real wage rate is equal to the marginal product of capital and the marginal product is proportional to average labor productivity—properties that hold in several business cycle models—then real unit labor costs would be constant.

Stephens (2004) and Saporta-Eksten (2014) provide empirical support using US data. Using the four 1992-1996 waves of the Health and Retirement Study (HRS), Stephens (2004) finds that annual food consumption is 16% lower when a worker reports that he is no longer working for the employer of the previous wave either because of a layoff, business closure, or business relocation, that is, the worker was displaced between two waves. Stephens (2004) finds similar results using the Panel Study of Income Dynamics (PSID).²⁸ Using the 1999-2009 biannual waves of the PSID, Saporta-Eksten (2014) finds that job loss leads to a drop in total consumption of 17%. In particular, a drop of about 9% occurs before job loss and a drop of about 8% around job loss. The drop before job loss suggests that either the worker anticipated the layoff or labor income was already under pressure. Moreover, this drop in consumption is very persistent and is only slightly less than 17% six years after displacement.

2.3.3 Savings and idiosyncratic uncertainty

The idea that idiosyncratic uncertainty plays an important role in the savings decisions of individuals has a rich history in the economics literature. The theoretical literature shows that idiosyncratic uncertainty increases savings when the third-order derivative of the utility function with respect to consumption is positive and/or the agent faces borrowing constraints.²⁹ Moreover, idiosyncratic uncertainty regarding unemployment is more important for recessions—like the recent recession—which are characterized by a prolonged downturn and an increase in the average duration of unemployment spells. Krueger, Cramer, and Cho (2014) document that the number of long-term unemployed increased during the recent recession in all countries considered, except in Germany. The results are particularly striking for the US. During the recent recession, the amount of workers who were out of work for more than half a year relative to all unemployed workers reached a peak of 45%, whereas the highest peak observed in previous recessions was about 25%.

Several papers have provided empirical support for the hypothesis that increases in idiosyncratic uncertainty increases savings. Using 1992-98 data from the British Household Panel Survey (BHPS), Benito (2004) finds that an individual whose level of idiosyncratic uncertainty would move from the bottom to the top of the cross-sectional distribution reduces consumption by 11% *ceteris paribus*. An interesting aspect of this study is that the results hold both for a measure of idiosyncratic uncertainty based on individuals' own perceptions as well as on an econometric specification.³⁰ Empirical evidence for this relationship during the recent downturn can be found in Alan, Crossley, and Low (2012) who argue that the observed sharp rise in the savings ratio of the UK private sector is driven by increases in uncertainty, rather than other explanations such as tightening of credit standards. In line with the mechanism emphasized in this paper, Carroll (1992) argues that employment uncertainty is especially important because unemployment spells are the reason for the most drastic fluctuations in household income. He provides empirical evidence that the fear of unemployment leads to an increased desire to save even when controlling for expected income growth.

²⁸Stephens (2004) also reports that the amount of the drop does not depend much on whether the job loss was expected or not. This result raises the question whether savings always do significantly increase in anticipation of an increased probability of job loss.

²⁹See Kimball (1992).

³⁰Although the sign is correct, the results based on individuals' own perceptions are not significant.

2.4 Calibration

In this section, we motivate the parameter values used to generate the results. The section starts with a discussion of the parameter values that play a key role in generating the results followed by a discussion of the remaining parameter values. The model period is one quarter. Targets are constructed using Eurozone data from 1980 to 2012.³¹

Key parameter values. Regarding the choice of key parameter values, our strategy is to show that our main results can be generated with conservative choices. For example, we set the coefficient of relative risk aversion equal to 2. Even though risk aversion is not that high, the differences between our heterogeneous-agent model and the representative-agent version are substantial.

The incidence and duration of unemployment spells are obviously important. The probability of job destruction, δ , and the parameter characterizing efficiency in the matching market, ψ , are chosen to ensure that the unemployment rate and the expected duration of an unemployment spell in the economy without aggregate risk match their observed counterparts, which are equal to 10.7% and 3.57 quarters.³² The latter corresponds to a worker matching probability of 28%. These numbers imply a 3.36% job separation rate and a value for ψ equal to 0.574.

The generosity of the unemployment insurance regime is a key variable affecting the severity of unemployment spells. They vary a lot across countries in Europe. Esser, Ferrarini, Nelson, Palme, and Sjöberg (2013) report that net replacement rates for insured workers vary from 20% in Malta to just above 90% in Portugal. Most countries have net replacement rates between 50% and 70% with an average duration of around one year. Coverage ratios vary from about 50% in Italy to 100% in Finland, Ireland, and Greece. Net replacement rates for workers that are not covered are much lower. In most countries, these are less than 40% and in some countries substantially so. In the model, unemployment benefits are set equal to 50% and—for computational convenience—are assumed to last for the duration of the unemployment spell no matter how long. The 50% used in the model is possibly a bit less than the average observed, but this is compensated by the longer duration of unemployment benefits in the model.

The inability to fully insure against unemployment risk plays a key role in our model. It is, therefore, important that the model generates a realistic drop in consumption during an unemployment spell. The best empirical counterpart for us is provided in Kolsrud, Landais, Nilsson, and Spinnewijn (2015) who use Swedish data. They find that consumption drops on average by 34% during the first year of an unemployment spell. A key parameter to target this number is the scale parameter, $\chi = 4 \cdot 10^{-5}$, which characterizes the liquidity benefits of money. This parameter affects the level of financial assets held and, thus, the ability of the agent to insure against unemployment spells. The literature also provides some evidence on pre-displacement wealth levels. For example, Gruber (2001) provides evidence for the US. In

³¹ Average unemployment duration data are based on all of Europe, since no Eurozone data is available for this time period. Details about data sources are given in appendix 2.A.

³²We use the model *without* aggregate risk for this part of the calibration, because finding the right parameter values requires solving the model numerous times, which is computer intensive for the model *with* aggregate risk. With aggregate risk, the average unemployment rate is 11.7% and the average expected duration of an unemployment spell is 4.03 quarters.

section 2.5, we will show that our calibration is conservative. That is, we generate the targeted consumption drop without making agents unrealistically poor.

The main focus of this paper is on the interaction between sticky nominal wages and deflationary pressure induced by uncertain job prospects. Consequently, a key role is played by ω_p , the parameter that indicates how responsive nominal wages are to changes in the price level. If ω_p is equal to 0, then nominal wages do not respond to changes in the price level at all and if ω_p is equal to 1, then nominal wages respond fully so that real wages are not affected by inflation. We set ω_p equal to 0.70, but we will report results for other parameter values as well. This value is conservative for a quarterly model. As mentioned above, Druant, Fabiani, Kezdi, Lamo, Martins, and Sabbatini (2009) report that only 6% of European firms adjust wages (of their main occupational groups) more than once a year to inflation and only 50% do so once a year.³³

Finally, the curvature parameter in the utility component for liquidity services, ζ , plays an important role, because it directly affects the impact that changes in future job security have on the demand for the liquid asset. With more curvature, the demand for the liquid asset is less sensitive and increased concerns about future job prospects will generate less deflationary pressure. We set ζ equal to 2, which corresponds to a money demand elasticity with respect to the nominal interest rate equal to -0.5 , which is the preferred value in Lucas (2000).³⁴ The elasticity of money demand with respect to transactions by the household, consumption, is equal to γ/ζ , which equals 1 for our choices for the coefficient of relative risk aversion, γ , and ζ .

Other parameter values. Based on the empirical estimates in Petrongolo and Pissarides (2001), the elasticity of the matching probability with respect to tightness, η , is set equal to 0.5. The average share of the surplus received by workers, ω_0 , and the elasticity of the wage rate with respect to changes in aggregate productivity, ω_z , are set such that the standard deviation of employment relative to the standard deviation of output are in line with their empirical counterpart. Our mechanism creates additional volatility, but not enough to get enough volatility in employment for a wide range of values for ω_0 and ω_z .³⁵ If ω_z is increased and wages become more flexible, then the same level of employment volatility can be achieved by increasing the value of ω_0 . In our benchmark calibration, we use $\omega_0 = 0.97$ and $\omega_z = 0.3$.

In our model, the presence of idiosyncratic risk lowers average real rates of return. If we would set β equal to 0.99, then average return would be unrealistically low.³⁶ Therefore, we

³³Moreover, even if firms adjust for inflation they typically do so using backward looking measures of inflation, which reduces the responsiveness to changes in inflationary pressure.

³⁴The first-order condition for a bond with a risk-free nominal interest rate is given by

$$1 = \beta (1 + R_t) \beta \mathbb{E}_t \left[\left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\gamma} P_t / P_{t+1} \right].$$

Using $(1 + R_t)^{-1} \approx 1 - R_t$, we get

$$\ln (M_{i,t+1} / P_t) \approx -\zeta^{-1} \ln R_t + \zeta^{-1} (\ln \chi + \gamma \ln c_{i,t}).$$

³⁵One would think that very large employment fluctuations are possible in a model with deflationary spirals. As a researcher, however, one faces limits when one explores possibilities to increase the magnitude of such a mechanism. One limitation is computational. As fluctuations increase, it becomes more difficult to obtain an accurate numerical solution. Another limitation is that at some point the true model solution could very well be non-stationary. This would of course be interesting, but would require development of a new set of numerical solution algorithms.

³⁶In particular, it is important that we avoid the situation in which the average rate at which agents are willing

set β equal to 0.985.

Aggregate productivity z_t is equal to 0.978 in a recession and 1.023 in a boom. The probability of switching from boom to recession and vice versa is equal to 2.5%.³⁷

Without loss of generality, we use aggregate money holdings $\bar{M} = 0.13$ to normalize $P_t = 1$ in the economy without aggregate risk.

Parameters values in the representative-agent model. We will compare the results of our model with those generated by the corresponding representative-agent economy. Parameter values in the representative-agent model are identical to those in the heterogeneous-agent model, except for β . In the representative-agent economy, the value of β is set equal to $\beta \mathbb{E}_t \left[(c_{i,t+1}/c_{i,t})^{-\gamma} \right]$ for agents holding equity when there are no aggregate shocks.³⁸ Without this adjustment, the agent in the representative-agent economy would have a more short-sighted investment horizon and average employment would be lower.

2.5 Agents' consumption, investment and portfolio decisions

In section 2.5.1, we describe key aspects of individual consumption and in particular its behavior during an unemployment spell. In section 2.5.2, we focus on the individual's investment decisions.

2.5.1 Post-displacement consumption

In this section, we discuss the reduction of individual consumption following displacement and the key factors that are behind the substantial drop generated by the model. Two key factors are the stance of the business cycle (boom versus recession) and the level of pre-displacement asset holdings.

Magnitude of the post-displacement drop in consumption. Figure 2.2 plots changes in post-displacement consumption relative to consumption in the last period of employment. The model's parameters are calibrated such that the one-year drop equals its empirical equivalent, that is 34%. Although not targeted, the model predicts a proportional decrease over the first year similar to what is observed in the data.³⁹ The model's predictions differ from the data after the first year as the average consumption drop only settles down after two years, whereas the data indicates that this happens after one year. However, there are not that many agents who are unemployed for more than one year. This is documented in figure 2.3, which plots the distribution of the duration of unemployment spells.

There are several reasons why the drop in consumption is of such a nontrivial magnitude. One reason is that unemployment benefits are only half as big as labor income. The second key factor affecting the magnitude of the drop is the average level of wealth at the beginning of the unemployment spell. Using US data, Gruber (2001) finds that the median agent holds enough

to hold risk-free real bonds is negative. This is not only realistic, but also problematic for our model in which the marginal utility of money is always positive and there is no inflation. These two features put an upper bound on $\mathbb{E}_t \left[(c_{i,t+1}/c_{i,t})^{-\gamma} \right]$ and a lower bound on average real interest rates.

³⁷The process is chosen such that $\mathbb{E} [\ln z_t] = 0$, $\mathbb{E}_t [\ln z_{t+1}] = 0.95 \ln z_t$, and $\text{std} (\ln z_{t+1} - 0.95 \ln z_t) = 0.007$ hold.

³⁸Without aggregate shocks, this value is the same across all agents, because equity is risk-free in that case.

³⁹See Kolsrud, Landais, Nilsson, and Spinnewijn (2015).

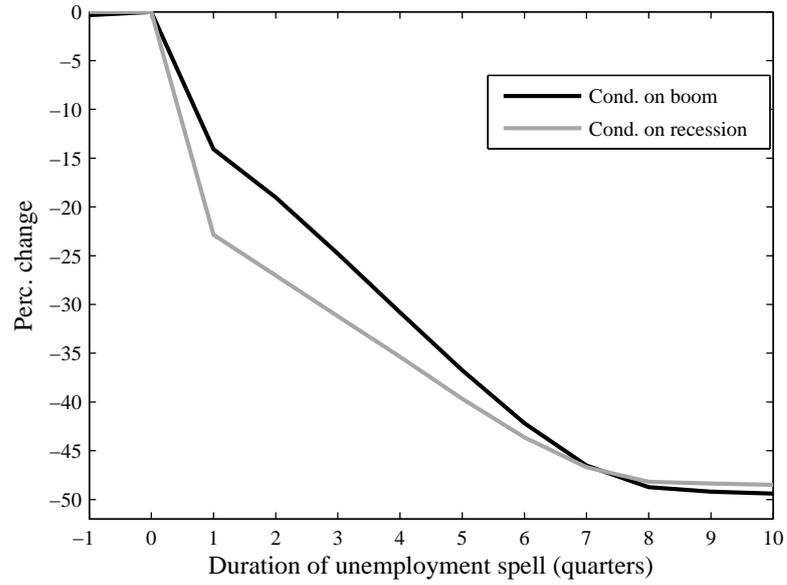


Figure 2.2: Evolution of consumption drop over the unemployment spell.

Notes. The black line illustrates the consumption drop of an individual that becomes unemployed in period 0 conditional on a boom at the time of displacement. The grey line illustrates the equivalent path conditional on a recession.

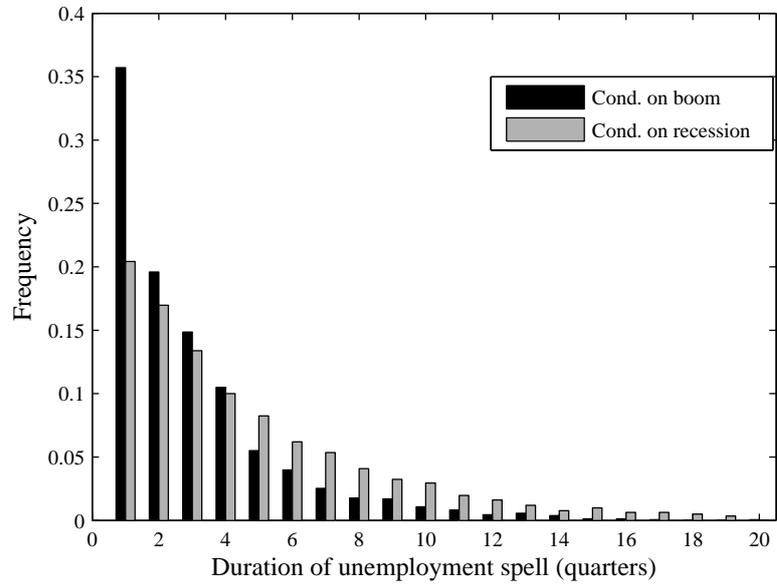


Figure 2.3: Distribution of the unemployed.

Notes. The black bars measure the fraction of unemployed at various durations conditional on a boom at the time of displacement. The grey bars provide the corresponding measure conditional on a recession.

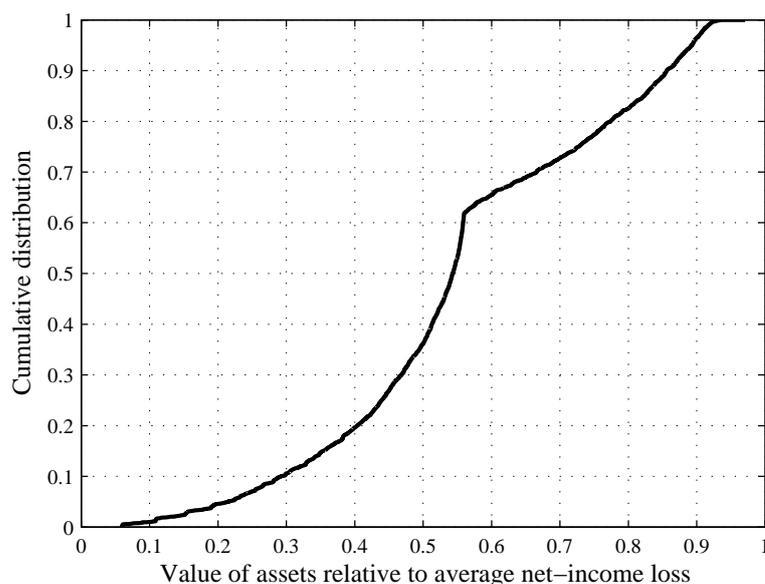


Figure 2.4: Financial assets at the beginning of an unemployment spell.

gross financial assets to cover 73% of the average net-income loss during an unemployment spell. Moreover, in terms of *net* financial assets, the median agent does not even have enough to cover 10% of the average net-income loss.⁴⁰ One would expect that agents accumulate less savings in Europe where unemployment benefits are higher.⁴¹ In our model, the median agent's asset holdings are equal to 54% of the average net-income loss during unemployment spells. This is true for both gross and net asset holdings, since there is no debt in our model. Thus, relative to these observed levels, the median agent is less wealthy in terms of gross financial assets and a lot wealthier in terms of net financial assets. Figure 2.4 plots the complete cumulative distribution function of wealth levels at the beginning of unemployment spells as a function of the average net-income loss. Agents in the bottom of the wealth distribution are substantially richer than their real world counterparts, even if we focus on gross assets. For example, Gruber (2001) documents that 38% of all workers do not have enough assets to cover 25% of the average net-income loss and that fraction is only 7% in our model. Thus, it is not the case that one needs unrealistically low wealth levels to generate the nontrivial post-displacement drops in consumption that are observed in the data.

Another aspect affecting consumption during unemployment spells is the ability to borrow. In our model, agents cannot go short in any asset and agents would presumably hold less financial assets if they had the option to borrow. Kolsrud, Landais, Nilsson, and Spinnewijn (2015) report, however, that the amount of consumption that is financed out of an increase

⁴⁰Gross financial assets would be the relevant measure if debt can be rolled over during an unemployment spell, whereas net financial assets would be the relevant measure if that is not the case. Kolsrud, Landais, Nilsson, and Spinnewijn (2015) find that average debt *decreases* during unemployment spells, which means that the observed gross measure overestimates the amount of funds agents have to cover income losses.

⁴¹In contrast to Gruber (2001), Kolsrud, Landais, Nilsson, and Spinnewijn (2015) do not provide pre-displacement wealth levels as a function of expected earnings losses. But some information is available. In particular, using an average unemployment spell duration of 4 months, the median Swedish agent's level of gross financial assets is equal to roughly 13% of average net-income loss. In our model, calibrated to an average level of European unemployment benefits, net income drops by more (by half as opposed to one third in Sweden), but agents that become unemployed are wealthier.

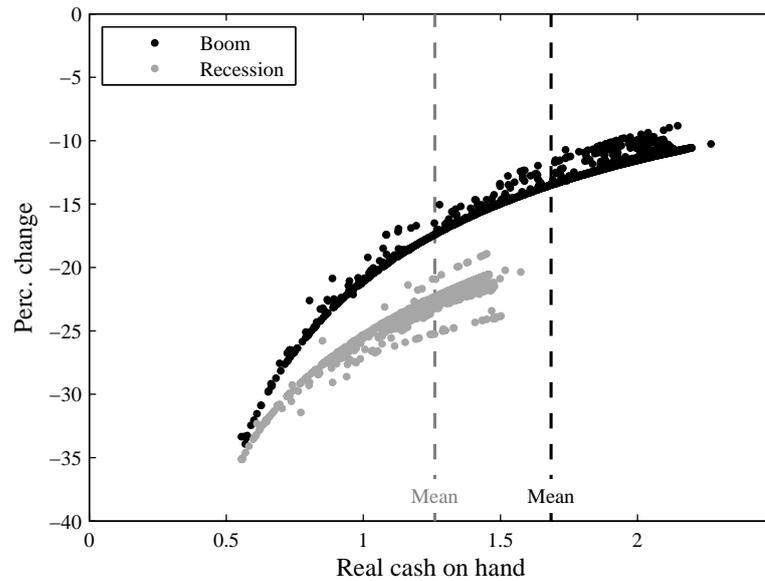


Figure 2.5: Consumption drop upon unemployment.

in debt actually *decreases* following a displacement. More importantly, we think that the key feature to capture is the level of the consumption drop, not whether this is accomplished by no borrowing or some borrowing and a lower level of financial assets.

State dependence of consumption drop. Figure 2.5 presents a scatter plot of the reduction in consumption and beginning-of-period cash on hand, where both are measured in the period when the agent becomes unemployed.⁴² There are two distinct scatter plots, one for booms and one for recessions.⁴³

Consistent with figure 2.2, figure 2.5 documents that the consumption drop is *much* more severe if the unemployment spell starts in a recession. The drop in consumption varies from 18.9% for the richest agent to 35.1% for the poorest agent during recessions. The range increases during a boom: The richest agent faces a consumption drop of "only" 8.8% whereas the drop for the poorest agent is very similar, namely 33.9%. There are several reasons why consumption drops by more during recessions. First, job finding rates are lower during recessions, which means that agents anticipate a longer unemployment spell. For a given amount of cash on hand, they will therefore reduce consumption by more during recessions. A second factor is that the amount of cash on hand is substantially lower during recessions, because the value of equity holdings is substantially lower during recessions. The value of cash on hand held by a newly unemployed agent is on average equal to 1.26 in a recession and on average equal to 1.68 in a boom.

In reality, workers may not face such a large drop in the value of their equity portfolio when the economy enters a recession. After all, quite a few workers do not own equity. We think

⁴²Cash on hand is equal to the sum of non-asset income (here unemployment benefits), money balances, dividends, and the value of equity holdings.

⁴³The level of employment is also important for the observed consumption drop, which explains the scatter of observations. In particular, the consumption drop is usually lower at the beginning of a boom and steeper at the beginning of a recession.

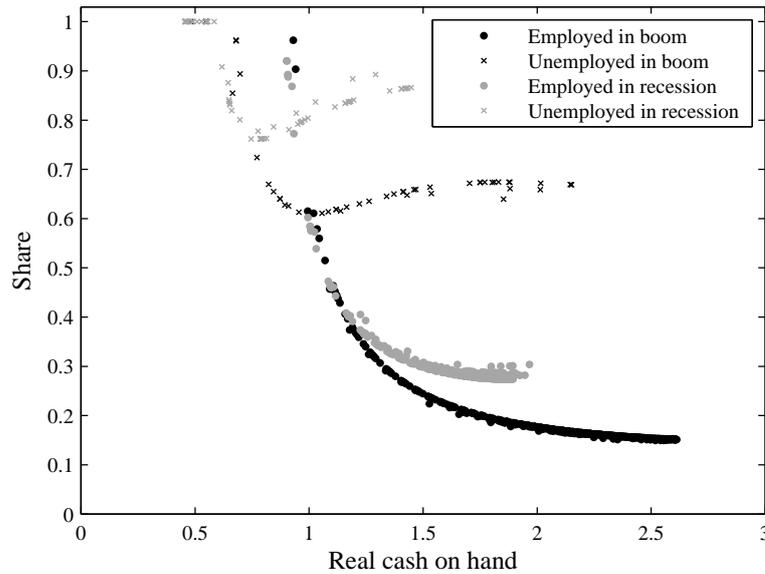


Figure 2.6: Portfolio shares in liquid asset.

that the cyclical nature of the consumption drop associated with unemployment spells driven by the cyclical behavior of equity prices in our model do capture real world phenomena. First, although clearly not all workers hold equity, many hold assets such as housing that also have volatile and cyclical prices. Second, unemployed workers may receive handouts and or loans from affluent family members, friends, or financial intermediaries whose ability and willingness to help may be affected by the value of their assets, which could very well be cyclical.

Figure 2.5 also underscores the non-trivial role of agents' wealth levels for the post-displacement consumption drop when they become unemployed.

2.5.2 Investment decisions

In this section, we first discuss how portfolio shares vary with agents' wealth levels and employment status. Next, we discuss agents' holdings of money and equity during unemployment spells.

Portfolio composition and cash-on-hand levels. Figure 2.6 presents a scatter plot of the share of the liquid asset in the agent's investment portfolio (y-axis) and beginning-of-period cash-on-hand levels (x-axis). The graph documents that there are three distinct patterns depending on the stance of the business cycle (boom or recession) and the agent's employment status. Although, the model predicts intricate and non-monotonic relationships, the model basically predicts that the share of the portfolio invested in the liquid asset (i) is higher during recessions (at the same cash-on-hand level), (ii) is higher for unemployed than for employed agents (again at the same cash-on-hand level), and (iii) is higher when the agent's cash-on-hand level is lower.

More precisely, the following four observations can be made. The first observation is that the fraction invested in the liquid asset is higher during recessions for almost all beginning-of-

period cash-on-hand levels for both employed and unemployed workers. The reason is that agents save less during recessions and diminishing returns of money's transactions services imply that agents that invest less would invest a larger fraction in money. The second observation is that the poorest employed agents with lower cash-on-hand levels hold a larger fraction in the liquid asset than unemployed agents with the same cash-on-hand levels. This can also be explained by diminishing returns. The reason is that these poor employed agents carry few financial assets into the next period. High marginal transaction benefits at low levels of money holdings induces these agents to hold most of their financial assets as money. In fact, some are constrained by the short-sell constraint and only hold money. An unemployed agent with the same amount of cash on hand as an employed agent saves more and would—according to the same argument—hold a smaller fraction as money.

The third observation is that the last observation is only true for the lowest cash-on-hand levels that we observe for employed agents: At higher cash-on-hand levels, unemployed agents hold a larger fraction in liquid assets than employed agents. Moreover, the difference is increasing with the agent's wealth level and the difference becomes huge. For example, at the highest cash-on-hand levels observed for an unemployed agent, i.e., around 2.15, the unemployed agent holds 67% of his portfolio in the liquid asset, whereas the employed agent only holds 17%. The fourth observation is that the relationship between the share invested in the liquid asset and beginning-of-period cash-on-hand levels is not monotonic.

These last two observations cannot be explained with diminishing transaction benefits to holding money. These observations can be explained by another factor that works in the opposite direction, that is, there is a reason why poor agents want to invest *more* in the risky asset.⁴⁴ That other factor is risk. Keeping beginning-of-period resources the same, an unemployed agent invests more (in money and equity) than an employed agent, since expected labor income is lower for the unemployed. As asset income increases relative to non-asset income, the agent's total income becomes riskier. The unemployed agent invests more in the safer asset to reduce this risk, which explains the third observation. Risk also explains the observed non-monotonic behavior. For simplicity, suppose that the safer asset does not have any transaction benefits. If the amount invested is low relative to non-asset income, then investment risk is of second-order importance whereas the higher expected return on the risky asset is of first-order importance. This latter effect is dominated by diminishing returns of money at low investment levels, but this is not the case as the cash-on-hand level of the unemployed agent becomes bigger and the agent starts investing more.

The discussion does not give a complete picture, because it ignores general equilibrium effects such as changes in prices, which will dampen the effects discussed and are important for the quantitative outcomes.

Money demand and cash on hand. This paper investigates the possibility whether concerns about future unemployment affect the demand for money, which in turn would affect the price level, which—in the presence of sticky nominal wages—would affect real wage costs, which in turn would affect unemployment. Consequently, it is key to fully understand the demand for money and in particular how the aggregate demand for money varies over the business cycle. Figure 2.7 presents a scatter plot of the demand for real money balances and beginning-of-period cash-on-hand levels. There are four distinct patterns depending on the stance of the

⁴⁴Surely, this channel is affecting the other observations as well, but is then dominated by the diminishing returns.



Figure 2.7: Money demand (real).

business cycle (boom or recession) and the agent's employment status. As discussed above, almost all unemployed workers hold larger shares of their portfolio in the liquid asset than employed workers and both employed and unemployed workers typically hold a larger share in the liquid asset during recession than during booms. Figure 2.7 shows that both properties are also true when we consider the amount of real money balances as opposed to the share of money in the portfolio.⁴⁵ The figure also illustrates that—everything else equal—the demand for real money balances increases with beginning-of-period cash on hand.

This paper argues that the interaction between sticky nominal wages and the inability to insure against unemployment risk deepens recessions. Key in understanding the underlying mechanism is the cyclical behavior of the demand for real money balances. When the economy enters a recession, then aggregate money demand is pushed in opposite directions by different factors. During recessions aggregate cash on hand falls. This would reduce aggregate demand for real money balances. Figure 2.7 documents, however, that there are two reasons why aggregate demand for money *increases* in our economy with incomplete markets during recessions. First, all agents demand more money for given cash-on-hand levels during recessions. Second, unemployed agents demand more money (for given cash-on-hand levels) and there are more unemployed during recessions. In the next section, it will become clear that the last two effects dominate and aggregate money demand *increases* during recessions, whereas aggregate money demand *decreases* during recessions in the representative-agent version of our economy.

To see that this is a remarkable result, consider the partial equilibrium version of our model in which the price level and the equity price are fixed. Also assume that there is no short-sale constraint. Markets are still incomplete because the agent cannot insure against unemployment risk. Now consider a decrease in the job finding probability. Could this lead to an in-

⁴⁵Whereas the observations are typically true when the *share* invested in the liquid asset is considered, the observations are always true when the *level* of money demand is considered.

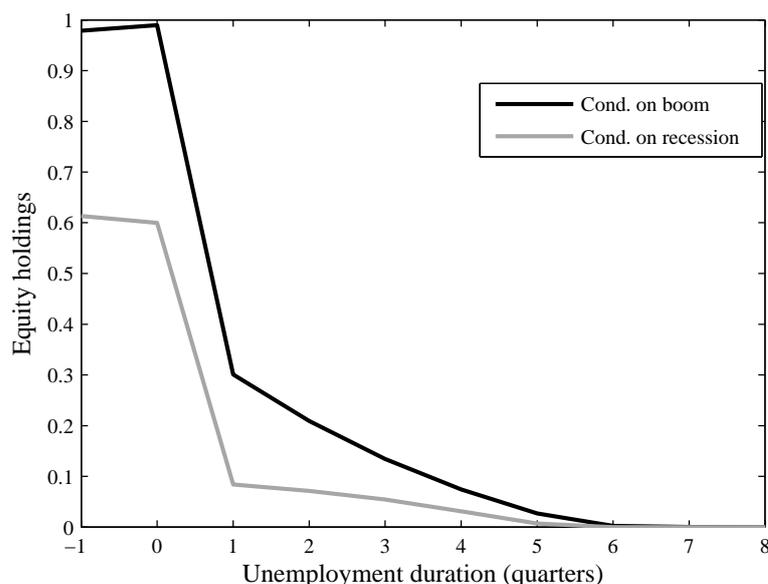


Figure 2.8: Post displacement equity holdings.

Notes. The black line illustrates the path for equity holdings of an individual that becomes unemployed in period 0 conditional on a boom at the time of displacement. The grey line illustrates the equivalent path conditional on a recession.

crease in the demand for real money balances? The answer is no. In this economy, demand for real money balances, $M_{i,t+1}/P_t$, and consumption, $c_{i,t}$ always move in the same direction. Since agents will lower $c_{i,t}$ in respond to a *decrease* in the job finding probability, money demand will decrease as well. The reason is the following. Equation (2.4)—which now holds with equality—implies that $\mathbb{E}_t \left[(c_{i,t+1}/c_{i,t})^{-\gamma} \right]$ is not affected. Equation (2.3), then directly implies that $c_{i,t}$ and $M_{i,t+1}/P_t$ move in the same direction. By contrast, in our model—in which prices adjust to clear markets and the short-sale constraint is always binding for some agents—aggregate money demand and aggregate consumption do move in different directions.

Financial assets during unemployment spells. Consumers dampen the drop in consumption following displacement by selling financial assets. Figures 2.8 and 2.9 document what this means for equity and money holdings, respectively. Although the total amount of financial assets and the amount invested in equity sharply decrease, the amount held in the liquid asset actually *increases* during the first two periods of an unemployment spell. The loss of wage income means that workers' cash-on-hand levels drop when they become unemployed. This reduces the demand for real money balances. For a given cash-on-hand level, however, the unemployed actually hold more money. Figure 2.9 documents that the last effect dominates in the beginning of an unemployment spell.

2.6 Business cycle properties

In the previous section, we showed that the inability of agents to insure against unemployment risk meant that workers face a sharp drop in consumption when they become unemployed.

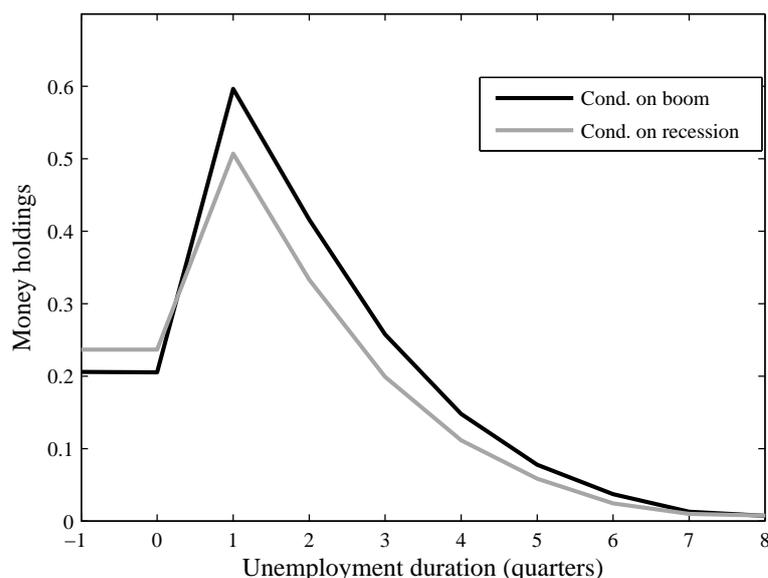


Figure 2.9: Post displacement money holdings.

Notes. The black line illustrates the path for real money holdings of an individual that becomes unemployed in period 0 conditional on a boom at the time of displacement. The grey line illustrates the equivalent path conditional on a recession.

We also discussed how this imperfect insurance affects money demand in ways that are not present in an economy with complete markets. In this section, we discuss what this means for business cycles. In particular, we document and explain why the interactions between sticky nominal wages, gloomy outlooks regarding future employment prospects, and the inability to insure against unemployment risk deepen recessions. We first discuss business cycle properties of the benchmark economy and compare those with the analogues in an economy with full risk sharing. Next, we compare economies with and without sticky nominal wages.

2.6.1 Benchmark: Business cycles with sticky nominal wages

Figure 2.10 plots the responses of key aggregate variables to a negative productivity shock for our benchmark economy and for the corresponding representative-agent economy.⁴⁶ The two panels in the top row of the figure display the responses for output and employment. These two panels document that the economy with incomplete risk sharing faces a much deeper recession than the economy with complete risk sharing. In particular, output drops by 7.2% in the heterogeneous-agent economy and by only 4.3% in the representative-agent economy.

The key in understanding this large difference in the depth of the recession is the behavior of the price level. In the representative-agent economy, the reduction in real activity decreases the demand for money and increases the price level. In our benchmark economy, $\omega_P = 0.7$, that is, a 1% increase in the price level leads to only a 0.7% increase in nominal wages and thus

⁴⁶In our benchmark calibration, productivity takes on only two values. The IRFs are calculated as follows. The starting point is period τ when productivity takes on its boom value and employment is equal to its mean value conditional on being in a boom. We then calculate the following two time paths for each variable. The “no-shock” time path is the *expected* time path when productivity takes on the high value in period $\tau + 1$. The “shock” time path is the *expected* time path when the productivity switches to the low value in period $\tau + 1$. The IRF is the difference between these two time paths.

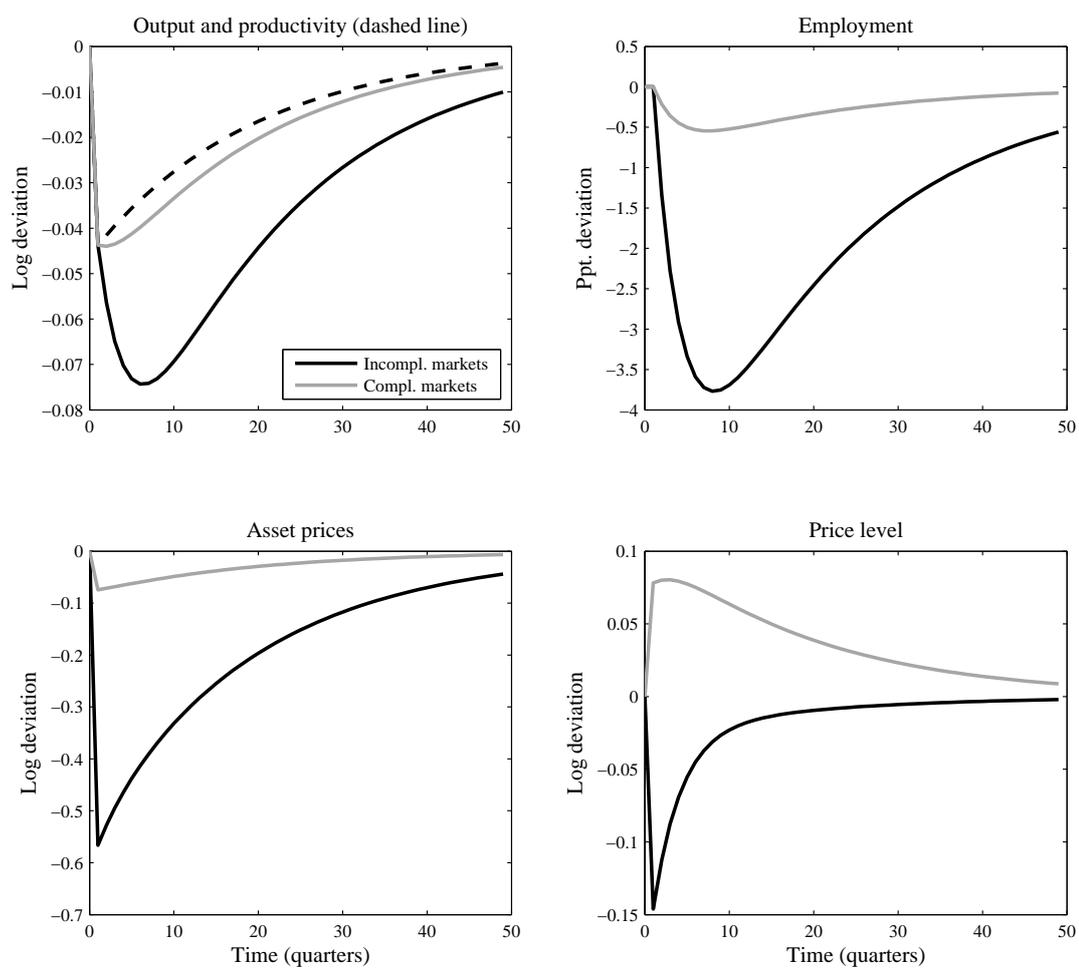


Figure 2.10: Impulse responses with sticky nominal wages.

Notes. These graphs plot the expected time paths of the indicated variable when the economy enters a recession (low-productivity regime) in period 1 relative to the expected time paths when the economy remains in a boom. $\omega_p = 0.7$, i.e., nominal wages increase with 0.7% when prices increase with 1%.

a 0.3% decrease in real wages. Thus, the direct effect of the reduction of productivity, z_t , on real wages is strengthened because nominal wages do not fully respond to changes in the price level. That is, our starting point is an economy in which the sluggish response of nominal wages to changes in prices actually *dampens* the economic downturn.

By contrast, the price level falls in the heterogeneous-agent economy, which is caused by an *increase* in the aggregate demand for the safer asset, i.e., money. To understand this different outcome, consider again figure 2.7, which plots the relationship between the demand for money as a function of beginning-of-period cash-on-hand levels during booms and recessions for both employed and unemployed agents. The reduction in real activity lowers cash-on-hand levels which reduces the demand for money by both employed and unemployed agents. The drop in cash-on-hand levels is substantial because the value of equity drops sharply. Nevertheless, aggregate money demand increases in the heterogeneous-agent economy, because there are strong forces pushing aggregate money demand up. As discussed above, both employed and unemployed agents hold more money during recessions for the same cash-on-hand level. Second, there are more unemployed agents during recessions and unemployed agents have larger money holdings, again for the same cash-on-hand level.

Whereas sticky nominal wages reduce the depth of recessions in the representative-agent economy, they worsen recessions in the heterogeneous-agent economy. Moreover, this is a quantitatively important effect, because a reduction in the price level (for any reason) starts the following self-reinforcing process that deepens recessions: The reduction in the price level puts upward pressure on real wages, which reduces profits, which in turn reduces investment in new jobs, which in turn reduces employment.⁴⁷ Since this reduction in employment is persistent it worsens employment prospects which leads to an increase in the demand for money when agents cannot insure themselves against unemployment risk. The impulse responses show that this mechanism is powerful enough to completely overturn the dampening effect that sticky nominal wages have in an economy with complete risk sharing.

Although it is a powerful mechanism, there is a counterforce. That is, as unemployment increases, the probability a firm finds a worker increases. For the results reported here, this counterforce is strong enough to ensure stability. For some parameter values, the numerical algorithm does not converge. This happens even if we use the homotopy approach to move very slowly to these parameter values (and every other trick that we could think of). It is possibly that the algorithm could not find a solution, because the model does not have a non-explosive solution at the parameter values used. Perturbation methods impose that aggregate shocks will not destabilize the economy as long as really small shocks do not do so.⁴⁸ Our experience suggests that this may impose stability where there is none.⁴⁹

2.6.2 Business cycles with flexible nominal wages

In this section, we discuss business cycle properties when changes in the price level leave *real* wages unaffected. Figure 2.11 plots the IRFs for the heterogeneous-agent economy and the

⁴⁷The negative productivity shock still has a direct negative effect on real wages. Which effect is stronger depends on parameter values. In our benchmark economy, the impulse response function of real wages is positive for the first two years and then turns negative.

⁴⁸For example, the technique developed in Reiter (2009) to solve models with heterogeneous agents relies on a perturbation solution for changes in the aggregate shock, which implies that the solution is *imposed to be stable* for the shocks considered as long as the Blanchard-Kahn conditions are satisfied.

⁴⁹We cannot prove this statement, since the inability to find a stable solution may also be due to the fact that one should try (even) harder or develop a better algorithm.

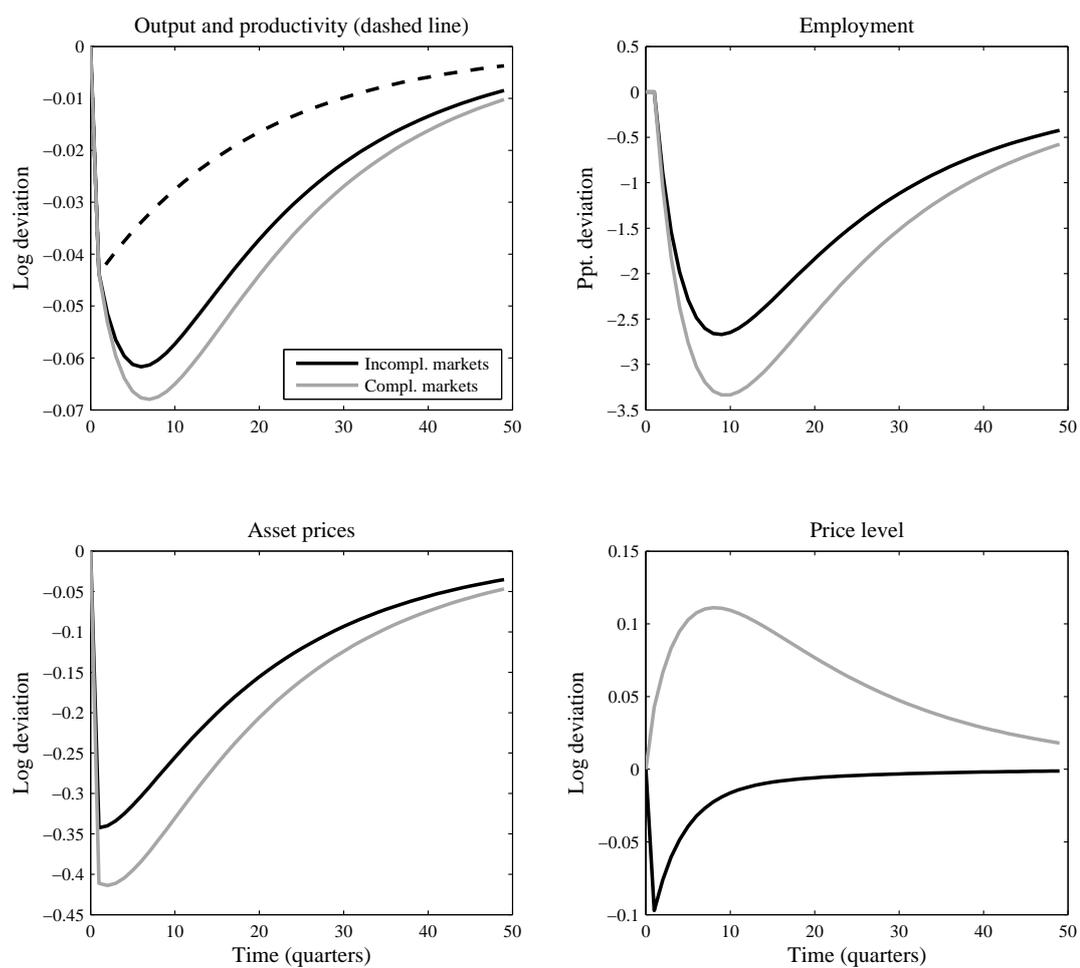


Figure 2.11: Impulse responses with flexible nominal wages.

Notes. These graphs plot the expected time paths of the indicated variable when the economy enters a recession (low-productivity regime) in period 1 relative to the expected time paths when the economy remains in a boom. $\omega_p = 1$, that is, nominal wages respond 1-for-1 to price changes.

IRFs for the corresponding representative-agent economy. There are several similarities with our benchmark results, but also one essential difference. We start with the similarities.

A negative productivity shock still has a direct negative effect on profits, which leads to a reduced demand for equity (firm ownership), which in turn means that fewer jobs are created. Also, increased concerns about employment prospects still induces agents in the heterogeneous-agents economy to increase their demand for money holdings, which again is strong enough to push the price level down, while it increases in the representative-agent economy.

There is also a striking difference. In the economy with flexible nominal wages, recessions are *less* severe when agents cannot insure themselves against unemployment risk. The reason is the following. In principle, a desire to save more because of precautionary motives could also increase the demand for equity. If this effect was present, then increased precautionary savings would *dampen* the reduction in the demand for equity induced by the direct negative effect of the productivity shock on profits. The IRFs document that this happens when $\omega_p = 1$, that is when nominal wages respond one-for-one to changes in the price level. The magnitude of the dampening effect is nontrivial. Whereas the biggest drop in employment is 3.3 percentage points in the representative-agent economy, it is equal to 2.7 percentage points in the heterogeneous-agent economy. These results make clear that a researcher would bias model predictions if this dampening aspect of precautionary savings during recessions is not allowed to operate. Increased uncertainty about the future increases the expected value of the marginal rate of substitution. Thus, increased uncertainty affects the first-order condition of the liquid asset *as well as* the first-order condition of the productive investment and both investments are therefore valued more. We argue that it is important to use the correct individual-specific marginal rate of substitution to discount all agents' future revenues and to let the model determine how increased idiosyncratic uncertainty affects the economy.

In our benchmark economy, we do allow this channel to operate, but it is dominated by the interaction between sticky nominal wages and uninsured unemployment risk. Increased uncertainty may increase the demand for equity, but it will also increase the demand for money. The latter decreases the price level, which increases real wages when nominal wages are sticky. The latter reduces profits, which in turn lower the demand for equity. The latter dominates any positive effect that precautionary savings may have on the demand for equity.

Robustness of the dampening effect. In all cases considered, we find that recessions are less severe in the heterogeneous-agent economy than in the representative-agent economy *if* nominal wages respond one-for-one to change in the price level. That is, this dampening effect is very robust. During the nineties, several papers argued that an increase in idiosyncratic risk could lead to a reduction in the demand for the risky investment when investors can save through a risky and a risk-free investment even though it would increase total savings.⁵⁰ This effect is referred to as temperance. We find that this result is quite fragile for several reasons.

The first reason is that it is a partial equilibrium result. In general equilibrium prices would adjust. This is important. Suppose that the economy *as a whole* can increase savings through the risky investment, but not through the risk-free investment, then the price of the risk-free asset would increase making the riskier asset more attractive. This plays a role in our economy, because the only way the economy as a whole can in the current period take action to get more

⁵⁰See Kimball (1990), Kimball (1992), Gollier and Pratt (1996), and Elmendorf and Kimball (2000).

goods in the future is by investing more in the productive asset, that is, in the risky asset. There are several other features, typically present in macroeconomic models, that make temperance less likely. One such reason is that the temperance result relies on idiosyncratic risk to be completely independent of investment risk. In macroeconomic models, that is not the case. The amount of idiosyncratic risk depends on the level of the wage rate.⁵¹ But the level of the wage rate is clearly correlated with the return of the risky asset, since both are affected by the same shocks.⁵² Another feature that works against the temperance result is the short-sale constraint on equity, which directly prevents a reduction in the demand for equity, at least for some agents. In our model, diminishing returns on the transactions aspect of money also work against temperance. This makes increased investment in the risk-free asset less attractive relative to a framework in which the return remains fixed. Finally, in our model, money is not nearly as risky as equity, but it is also not completely risk-free.

It may be the case that temperance can be generated in models with different utility functions, for example, if the utility function is such that the price of risk increases during recessions.⁵³ We leave this for future research.

2.7 Government policy

In this section, we discuss the two components of government policy in this model: unemployment insurance and monetary policy.

2.7.1 Unemployment-insurance (UI) policies

In this section, we analyze the impact of alternative unemployment-insurance policies. In our model, changes in such policies affect the economy quite differently than in many other models. For example, in the standard labor search model with a representative agent and flexible nominal wages an increase in unemployment benefits results in more volatile business cycles. By contrast, it would dampen business fluctuations in our environment. In the labor-market matching model with incomplete markets of Krusell, Mukoyama, and Sahin (2010), almost all agents (everyone except for the very poorest unemployment workers, who are borrowing constraint) benefit from a 25% reduction in unemployment benefits relative to the benchmark value. In fact, even an almost complete elimination of unemployment benefits (to only 1% of the benchmark value) is preferred by 92% of all workers. Our model is also characterized by a labor-market search friction and incomplete markets. Our model differs from the one in Krusell, Mukoyama, and Sahin (2010), however, in that the inability to insure against unemployment risk interacts with sticky nominal wages in such a way that it does not only affect the behavior of individual variables, but also makes aggregate variables behave quite differently than they do when there are complete markets. The consequence is that we draw very different conclusions regarding the desirability of changing unemployment benefits.

We start with a discussion of the impact of an increase in the level of unemployment-insurance when we all other parameter values remain constant, including those of the wage setting rule. In section 2.7.1.1, we discuss the case for our benchmark economy in which

⁵¹In the extreme case when the wage rate is zero or equal to the value of home production, there is no unemployment risk.

⁵²In the model, considered here they are both directly affected by z_t .

⁵³We considered models with different degrees of risk aversion, but this does not seem to matter for this issue.

$\omega_p = 0.7$, i.e., an $x\%$ change in the price level leads to an $0.7x\%$ ($0.3x\%$) increase (decrease) in the nominal (real) wage rate. In section 2.7.1.2, we discuss the results when there is no nominal wage stickiness, that is, $\omega_p = 1$. In section 2.7.1.3, we consider the results when the increase in the level of unemployment benefits does affect wage setting. In these first sections, we compare separate economies with different unemployment benefits. In section 2.7.1.4, we address the question how the economy would respond when μ is increases and we also address the question whether workers would prefer this switch taking into account transition dynamics. Finally, we document that the introduction of only a small countercyclical component in the level of unemployment benefits can achieve similar welfare improvements as much larger increases in the average level.

2.7.1.1 Higher UI when nominal wages are sticky, $\omega_p = 0.7$.

In this section, we consider an increase in the level of unemployment benefits, μ , from 0.5 to 0.55 in the benchmark economy when nominal wages only partially respond to changes in the price level. The increase in μ leads to a substantial decrease in the volatility of individual consumption. In particular, the standard deviation of individual log consumption drops by 15.3%.

The increase in μ also has a big impact on the behavior of aggregate variables. For example, the standard deviation of aggregate employment drops sharply by 49.7%. Moreover, the increase in μ leads to an *increase* in the average employment rate of 0.31ppt. By contrast, we find that the increase in μ leads to a *reduction* in average employment of 0.52ppt in the version of our model *without* aggregate uncertainty. Such comparative statics typically result in similar answers for economies with and without aggregate uncertainty, because aggregate uncertainty is relatively small. Volatility of the only aggregate random variable, productivity, is indeed modest in our model. Nevertheless, the induced volatility in asset prices and the nonlinearity of the matching function are important enough to get these two opposite implications for the change in average employment when μ increases.

In the economy with aggregate uncertainty, there are two effects associated with the increase in μ that increase the demand for equity and, thus, increase job creation. The first effect is that an increase in μ reduces the risk of holding equity, because the increase in μ not only reduces the volatility of real activity, it also leads to a substantial reduction in the volatility of stock prices. In fact, the standard deviation of the log of real equity prices drops with 49.8%. This reduction in risk leads to more job creation and an increase in average employment of 0.42ppt.⁵⁴ The second effect that pushes average employment up is related to the nonlinearity of the matching function, that is, increases in firm value have a smaller effect on employment than decreases in firm value. This means that decreases in the volatility of firm value increase average employment. This effect increases average employment with 0.41ppt.⁵⁵ Combining

⁵⁴We calculate this as follows. If there is no aggregate uncertainty, then the increase in μ leads to a decrease in employment of 0.52ppt and a decrease in real equity value of 5%. If there is aggregate uncertainty, then the same change leads to a decrease in the average real equity value by only 1%. This lower drop in equity value is due to the fact that there also is a decrease in aggregate uncertainty. Assuming that these effects are linear, the difference between the 5% and the 1% drop corresponds to an increase in average employment of 0.416ppt (= $4/5 \times 0.52$ ppt).

⁵⁵We calculate this as follows. When $\mu = 0.5$, then the introduction of aggregate uncertainty leads to a reduction in employment of 1.01ppt of which 46ppt can be explained by the reduction in the *average* equity price. The remainder of 0.55ppt is, thus, due to the nonlinearity of the matching function. When $\mu = 0.55$, then this nonlinearity effect is only 0.14ppt. Thus, when μ increases from 0.5 to 0.55 in the economy with aggregate uncertainty, then there is a reduction of the impact of the nonlinearity on average employment of 0.55ppt–0.14ppt=0.41ppt.

the two positive effects on average employment related to the decrease in aggregate volatility with the negative effect related to the reduction in savings because of better individual insurance, we get an increase in average employment of 0.31ppt.

The direct effect of an increase in μ is an increase in the tax rate. This direct effect is dampened by the increase in the tax base induced by the increase in employment.⁵⁶ The indirect effect is not strong enough to decrease *average* tax rates, but it is strong enough to do so during recessions. The reason is that the smaller business cycle fluctuations and the higher average employment rate imply that the tax base at the higher level of μ is especially higher during recessions. If tax rates would be distortionary—which they are not in our model—then lower tax rates during recessions could lead to a further dampening of business cycle fluctuations.

2.7.1.2 Higher UI when nominal wages are flexible, $\omega_p = 1$

The consequences of an increase in μ are quite different when nominal wages respond one-for-one to changes in the price level, i.e., when $\omega_p = 1$. Individual consumption becomes again less volatile, but the reduction in the standard deviation of individual log consumption is only 8.4%, whereas the drop is equal to 15.3% in our benchmark economy. Even bigger differences are observed for the impact of the increase in μ on aggregate variables. Whereas the standard deviation of the aggregate employment rate dropped by almost 50% in the benchmark economy, the standard deviation *increases* with 9.3% when $\omega_p = 1$. As discussed in section 2.6.2, increased precautionary savings *dampens* business cycles when $\omega_p = 1$. At higher values of μ , agents are better insured and precautionary savings increase by less during recessions. Consequently, there is also less dampening and business cycle fluctuations become more volatile as μ increases.

2.7.1.3 Higher UI when average real wages adjust

The discussion above considered an increase in unemployment insurance while keeping the wage setting rule the same. This is not unreasonable given that several empirical papers find that UI benefits do *not* have a significant effect on wages.⁵⁷ Schmieder, von Wachter, and Bender (2014) find that UI benefits have a significant *negative* effect on wages. This could happen if higher UI benefits prolong unemployment spells and increases skill loss. Nekoei and Weber (2015) find that UI benefits have a positive effect on re-employment wages. This could happen because an increase in UI benefits increases workers' outside option or because it allows workers to find a better match. If it is the former, then higher UI benefits would decrease the surplus of the match and the share that accrues to firm owners, which in turn would negatively affect job creation.

Even though the empirical evidence is inconclusive, it is interesting to see how results change if wages do adjust following an increase in UI benefits. In our next exercise, ω_0 —and thus average wages—increase when μ increases according to the following mechanism. For the firm owner, the value of being in a match with a worker is equal to that period's dividend plus the market value of the firm, adjusted for the probability that the firm may be in its last period; his outside option is zero. For both employed and unemployed agents,

⁵⁶In our model, taxes are only used to finance unemployment benefits and are, thus, very low. When average tax rates were higher, the increase in revenues because of an increase in the tax base would be higher.

⁵⁷See, for example, Card, Chetty, and Weber (2007), Lalive (2007), van Ours and Vodopivec (2008), and Le Barbanchon (2012).

we can calculate expected utility as a function of their financial wealth. Consequently, we can also calculate the cash equivalent of the difference between the expected utility of being employed and being unemployed, $a_{i,t}$. Worker i 's implied Nash bargaining weight is then equal to $f_{i,t} = a_{i,t} / (a_{i,t} + D_t/P_t + (1 - \delta) J_t/P_t)$. Since all workers receive the same wage, but have different wealth levels, the value of $f_{i,t}$ is worker specific. An increase in μ leads to a decrease in $a_{i,t}$ and a decrease in $f_{i,t}$, if the wage setting rule remains the same. Here we adjust ω_0 such that the average implied bargaining weight is not affected by the increase in μ . As pointed out in Hall and Milgrom (2008), $f_{i,t}$ may overstate the importance of fluctuations in the value of unemployment, because the worker's threat in bargaining is typically not leaving the relationship and becoming unemployed, but prolonging the negotiations. Consequently, our procedure may overstate the upward pressure on wages following an increase in μ . By considering the case when wages do not respond at all as well as a case when wages respond probably too much, we can bound likely outcomes of the increase in μ .

When the increase in μ goes together with an increase in ω_0 , then the standard deviation of individual consumption drops by 8.9% instead of 15.3% and the standard deviation of aggregate employment drops by 40% instead of 49.7%.⁵⁸ Volatility does not drop by as much because an increase in ω_0 lowers average profits and makes profits more sensitive to changes in productivity, which in turn makes job creation and employment more volatile. The decrease in average profits, induced by the increase in ω_0 , would lower average employment. In fact, whereas employment increased with 0.3ppt when ω_0 remains constants, it decreases with 0.56ppt when ω_0 changes.

2.7.1.4 Transition dynamics and desirability of switching to higher UI levels

In the previous sections, we compared different economies that are identical except for the unemployment regime. In this section, we analyze how an economy responds to a one-time increase in μ that is completely unexpected and is believed to be a permanent change. The long-run results were discussed above. Here we discuss the transition paths. In addition, we ask the question which workers prefer this change in UI policy.

We focus on a change in UI policy when the economy has just entered a recession and we consider the increase in μ when wage setting does and does not adjust. Figure 2.12 plots the time paths for employment when the economy moves from a boom to a recession and back to a boom. It plots the series when the UI does and when it does not change. The results above made clear that an increase in μ leads to smaller fluctuations and a higher average employment level if ω_0 remains the same. Consequently, employment should drop by less if μ is increased at the start of a recession. The same turns out to be true if ω_0 does increase. That is, the negative effect of the induced increase in ω_0 on average employment is smaller than the dampening effect of this change in parameter values on business cycle fluctuations. When the economy gets out of the recession, however, the recovery is dampened by the higher unemployment benefits, both when ω_0 does and when ω_0 does not adjust. When μ remains equal to 0.5, then the employment level exceeds the employment level when μ and ω_0 both increase in the first quarter of the recovery and exceed the employment level corresponding to the case when only μ increases after two quarters. The result that higher unemployment benefits can damage economic activity is consistent with Hagedorn, Karahan, Manovskii, and

⁵⁸The value of ω_0 increases to 0.973 which means that average firm profits decrease by 10%.

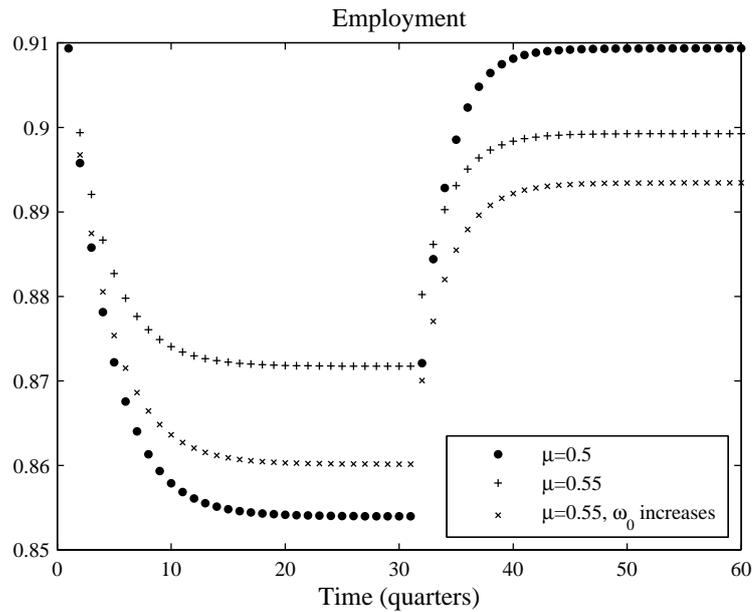


Figure 2.12: Switch to higher unemployment insurance at the start of the recession.

Notes. This graph compares the benchmark time path of employment with the time path when μ increases (unexpectedly and permanently) to 0.55, both when ω_0 does and when ω_0 does not adjust upwards.

Mitman (2015) who argue that the extension of unemployment benefits in the US increased unemployment in 2011—when the US recovery had started—by 2.5 percentage points.⁵⁹

We now turn to the question whether agents prefer the increase in μ . It is clear that agents are affected quite differently by the policy change. Unemployed workers benefit immediately from the increase in unemployment benefits. Employed workers benefit from the increase in μ because (i) the dampening of the downturn increases the value of their equity holdings, (ii) the higher μ increases their income if they become unemployed, and when ω_0 does not increase (iii) the long term average employment increases which means that all workers can expect to be less affected by unemployment. Moreover, although average tax rates increase, they are lower during this initial and future recessions when ω_0 does not increase.

To evaluate whether agents like the increase in μ , we calculate the expected utility when μ changes and when it does not change. To make the utility changes comparable across agents, we calculate the cash equivalent.⁶⁰ Figure 2.13 displays the cash equivalent of the proposed change as a function of the agent's beginning-of-period cash-on-hand level.⁶¹ The figure documents that *all* unemployed and *all* employed agents prefer the switch to the higher level of unemployment benefits, both when ω_0 does and when ω_0 does not adjust. For the same cash-on-hand levels, an unemployed worker benefits more than an employed worker. This is not surprising given that an unemployed worker benefits directly from the higher unemployment benefits. Rich agents benefit more than poor agents. One reason is that they hold more equity and, thus, benefit from the fact that stock prices drop by less when μ is increased. All

⁵⁹ Amaral and Ice (2014) argue that the extension only had a minor impact and part of the increase in the unemployment rate was due to a reduction in the number of unemployed leaving the labor force (and thus unemployment).

⁶⁰ To be able to do this, we calculate expected utility as a function of beginning-of-period cash-on-hand levels.

⁶¹ Cash on hand is measured at the point when it is known that the economy has entered a recession, but before it is known that μ has changed

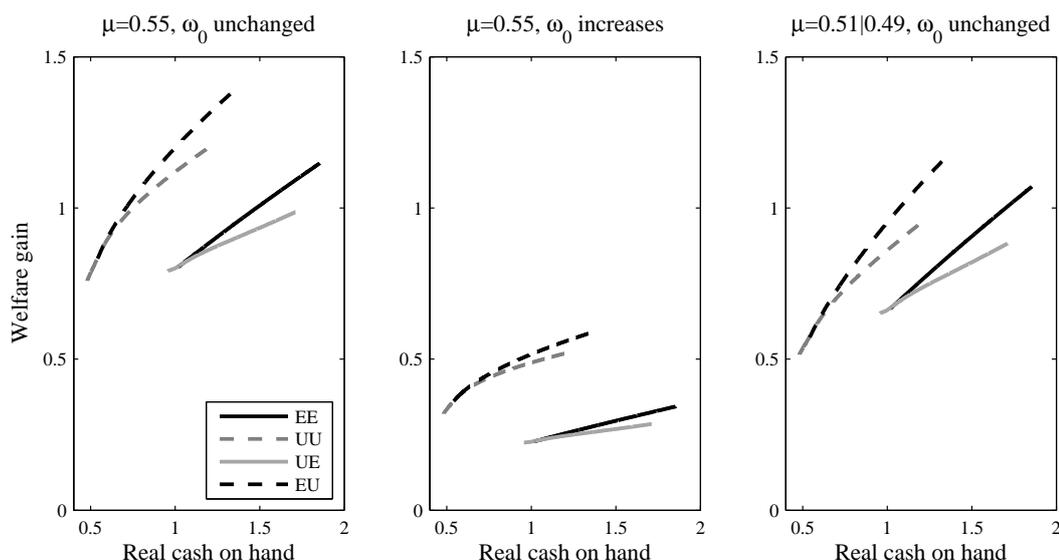


Figure 2.13: Welfare gains from changing unemployment insurance in the first period of a recession.

Notes. Welfare gains are measured as cash-on-hand equivalents. The average welfare gain with ω_0 unchanged equals 1.04. The average welfare gain with $\omega_0 = 0.973$ equals 0.32. The average welfare gain when unemployment insurance is made cyclical equals 0.93. Welfare gains are drawn for the four possible combinations of employment status in the previous and current period, because an agent's portfolio of assets in the beginning of the current period depends on his employment status in the last period.

agents benefit less for the case when both μ and ω_0 increase. This is even true for employed workers who hold no equity. The simultaneous increase in ω_0 implies that the downturns in real activity and stock prices are not dampened by as much. The value of poor worker is not affected by changes in stock prices since he does not hold equity. He does benefit from the higher wage brought about by the increase in ω_0 . But negative aspects of the increase in ω_0 , such as worsened future employment prospects, weigh more heavily.⁶²

If the change in μ occurs at the beginning of a boom, then the calculated cash equivalents are lower for all workers. The reason is that an increase in μ not only dampens recessions, it also dampens booms since the upward pressure on prices induced by a reduction in precautionary savings is smaller at higher levels of μ . It is still the case that all workers prefer the change when ω_0 does not increase. When ω_0 does increase, however, the richer employed and the richer unemployed workers do not prefer the increase in μ .

2.7.1.5 Higher UI benefits and unemployment duration

As discussed above, it is not clear from empirical studies whether changes in UI benefits affect wages. There is much more empirical support for the hypothesis that more generous UI benefits increase unemployment duration.⁶³ Our framework can explain increased unemployment duration even though search intensity is fixed and, thus, not affected by the level of UI benefits both when wages do and when wages do not depend on the level of UI benefits.

If wages depend on the level of UI benefits, then average employment and the average

⁶²At our calibration, employment is below the socially optimal level.

⁶³A long list of papers is given in Le Barbanchon (2012).

job finding rate are lower with higher UI benefits. Higher UI benefits do still lead to a higher employment level and a higher job finding rate during recessions.

If wages do not depend on the level of UI benefits, then higher UI benefits can negatively affect employment and the job finding rate during recoveries. For our benchmark calibration, they do not negatively affect employment and the job finding rate *on average*. But this outcome was the results of different aspects pushing job creation in different directions. For parameter values that are such that the negative effect, which is caused by workers saving less if they are better insured, is stronger, then higher UI benefits could have a negative effect on average employment and the average duration of unemployment spells even if wages do not depend on UI benefits.

2.7.1.6 Cyclical UI

Here we consider a cyclical unemployment insurance regime under which benefits increase during recessions and decrease during booms. In particular, the value of μ is equal to 0.51 in a recession and equal to 0.49 in a boom. Quantitatively, this is a smaller change than the increase in μ from 0.5 to 0.55 considered above. Moreover, the average value of μ stays the same.

This more modest change in the unemployment insurance policy reduces business cycle fluctuations by almost the same and does better in terms of its effect on average employment.

As discussed above, there are factors which push up average employment when μ increases. There is also a reason why average employment decreases: An increase in the value of μ implies that workers are better insured against unemployment risk, which in turn implies that they invest less (in creating new jobs), which reduces average employment.⁶⁴ By keeping the average level of μ constant we reduce this downward effect on employment. That is, average employment when μ equals 0.55 is 0.31ppt above average employment when μ equals 0.5. In the economy with the cyclical UI policy average employment is 0.74ppt higher.

The cyclical UI policy is almost as effective in dampening business cycles as the UI policy with a constant μ equal to 0.55. In particular, whereas the standard deviation of the employment rate drops from 0.026 to 0.013 when μ is increased from 0.5 to 0.55, it drops to 0.015 for the economy with the cyclical UI policy.

The combination of a higher average employment level and an only slightly higher level of cyclical employment fluctuations means that the employment level in the economy with the cyclical UI policy is higher than the employment level in the economy in which μ equals 0.55 in *every* period. Figure 2.14 compares the time paths for employment when the economy switches in the first period of the recession to either the policy with μ equal to 0.55 or to the cyclical UI policy. In both cases, starting from an economy in which μ is constant and equal to 0.5. The graph clearly documents that the modest but cyclical change in UI policy reduces business cycle fluctuations by almost as much as the larger acyclical change. The graphs also shows that employment levels are always higher in the economy with the cyclical UI policy.⁶⁵ The third panel of figure 2.13 shows that all agents prefer a switch to cyclical unemployment insurance.

⁶⁴Above, we documented that average employment increases when μ increases (and ω_0 does not), because an increase in μ reduces the risk of investment and the non-linearity of the matching function reduces average employment by less if business cycles are dampened.

⁶⁵This is even true in the first period, when the economy has not yet fully benefitted from the fact that agents invest more when the average value of unemployment benefits is lower.

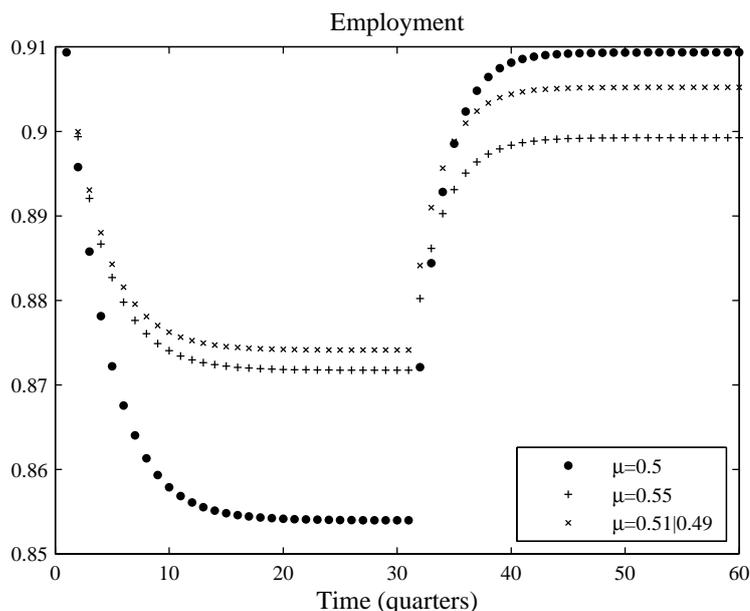


Figure 2.14: Switch to cyclical UI policy at the start of the recession.

Notes. This graph compares the benchmark time path of employment with the time path when μ increases (unexpectedly and permanently) to 0.55 and with the time path when the unemployment insurance scheme (unexpectedly and permanently) changes to a cyclical scheme in which $\mu = 0.51$ in recessions and $\mu = 0.49$ in booms. The value of ω_0 is kept the same.

Relative to the economy in which μ equals 0.55, the economy with the cyclical UI policy does also better in that tax rates are lower in every period. The reason is that unemployment benefits are lower in every period and the employment rate is higher.

2.7.2 Monetary policy

Could the severity of recessions that occur in our model possibly be alleviated by monetary policy?⁶⁶ If the central bank could respond to changes in z_t instantaneously and if the central bank could increase the money supply by “helicopter drops”, then the central bank could prevent deflationary pressure on the price level and upward pressure on real wages when nominal wages are sticky. In practice, however, there are several reasons why it may not be that easy for central banks to offset the harmful effects caused by the interaction of precautionary savings and sticky nominal wages.

In our heterogeneous-agent model, this interaction leads to a drop in the price level followed by a period of sustained inflation. Suppose that the central bank cannot prevent this initial drop, because it cannot respond instantaneously to the drop in productivity. To undo the harmful effects of nominal wage stickiness, the central bank would have to increase infla-

⁶⁶Several papers with heterogeneous-agent models adopt the cashless-economy approach. In a representative-agent economy, there are different ways to motivate a cashless economy. One is to assume that the level of real money balances does not interact with the real economy and that the central bank performs open market operations to ensure that the quantity of money outstanding is such that the desired interest rate is achieved. In other words, one could model the demand for money and equilibrium on the market for money, but this block would be independent of the other equations. At least, this motivation for the cashless economy is unlikely to carry over to a model with heterogeneous agents, since asset holdings, including money are essential for modelling precautionary savings and, thus, for understanding the role of money holdings for real activity. In this paper, there definitely is no such separation between demand for liquid assets and real activity.

tion when expected inflation is already higher than normal. That is, the central bank would have to adopt a regime of price-level targeting instead of a regime of inflation targeting. But price-level targeting may not be the best policy when the economy faces other problems. So the first reason why monetary policy may not eliminate the channel identified in this paper is that they may not be willing to adopt a price-level targeting regime.

Another aspect is, of course, that helicopter drops of money are not part of the usual set of central bank instruments. The typical way for a central bank to increase liquidity in the economy is to purchase government bonds from banks. This increases the liquidity position of banks. If the extra liquidity induces banks to issue more loans, then bank deposits will increase. That is, money holdings of the private sector will increase. Note, however, that the liability of the private sector to the financial sector has increased with the same amount. It is possible that this combined increase of liquid assets and debt eases workers concerns about future unemployment, for example, because the loans are (perceived to be) long-term loans. If workers care about their net-liquidity position, however, then this monetary stimulus would not undo workers desire to hold more money balances and there still would be downward pressure on the price level during recessions. This latter case would be especially relevant if bank loans cannot be rolled over if the worker becomes unemployed.

Monetary policy that undoes changes in the price level to offset nominal wage stickiness will have distributional consequences. In our model, all workers are *ex ante* identical. But suppose that there is another group of workers, who also have nominal sticky wages, but who are never unemployed. These workers would not like the central bank to decrease real wages by pushing up the price level.

Finally, another factor that is likely to make it difficult to figure out the right monetary response is a lack of information about the state of the economy and a lack of information about the "true" model. In our economy, changes in the price-level deepen the recession when nominal wages are sticky and there is an inability to insure against unemployment risk. When these two features are not present or when they are not strong enough, then changes in the price-level *dampen* recessions. Since the strength of different channels may change over time, it wouldn't be clear whether it always would be good if the central bank tried to undo any changes in the price level.

Appendices to Chapter 2

2.A Data Sources

- Eurozone GDP implicit price deflators are from the Federal Reserve Economic Data (FRED). Data are seasonally adjusted.
- Eurozone private sector hourly earnings are from OECD.STATEextracts (MEI). The target series for hourly earnings correspond to seasonally adjusted average total earnings paid per employed person per hour, including overtime pay and regularly recurring cash supplements. Data are seasonally adjusted.
- Unit labor costs are from OECD.STATEextracts. Data are for the total economy (employment based). Unit labour costs are calculated as the ratio of total labour costs to real output. Data are seasonally adjusted.
- Average unemployment rate: Average unemployment rate for the four large Eurozone economies, France, Germany, Italy, and Spain. Data is from OECD.STATEextracts (ALFS).
- Average unemployment duration: Average unemployment duration for Europe from OECD.StatExtracts. This is annual data. The data series for Europe is used because no data for the Eurozone is available, nor data for the big Eurozone countries. Starting in 1992, separate data is given for Europe, the European Union with 21 countries, and the European Union with 28 countries, and the series are quite similar over this sample period.

2.B Solution algorithm

2.B.1 Solution algorithm for heterogeneous agent model

In appendix 2.B.1.2, we document how we solve the individual problem taking as given perceived laws of motion for prices and aggregate state variables. In appendix 2.B.1.1, we document how to generate time series for the variables of this economy, including the complete cross-sectional distribution, taking the individual policy rules as given. The simulation is needed to update the laws of motion for the aggregate variables and to characterize the properties of the model. We make a particularly strong effort in ensuring that markets clear *exactly* such that there is no “leakage” during the simulation. This is important since simulations play a key role in finding the numerical solution and in characterizing model properties.⁶⁷

⁶⁷If equilibrium does not hold exactly, then the extent to which there is disequilibrium is likely to accumulate over time, unless the inaccuracy would happen to be *exactly* zero on average. Such accumulation is problematic, since long time series are needed to obtain accurate representation of model properties.

2.B.1.1 Simulation and solving for laws of motion of key aggregate variables

The perceived laws of motion for the real stock price, \tilde{J}/\tilde{P} and the price level, \tilde{P} , are given by the following two polynomials (using a total of 12 coefficients):

$$\ln \tilde{J}/\tilde{P} = a_0(z) + a_1(z) \ln q + a_2(z) (\ln q)^2, \quad (2.15)$$

$$\ln \tilde{P} = b_0(z) + b_1(z) \ln q + b_2(z) (\ln q)^2. \quad (2.16)$$

Note that q is not only the level of employment, but also the number of firms, and the aggregate amount of equity shares held. We only use the first moment of the distribution of equity holdings, as in Krusell and Smith (1997), but we use a nonlinear function.⁶⁸ To update the coefficients of this law of motion, we run a regression using simulated data. In this appendix, we describe how to simulate this economy taking the policy rules of the individual agents as given. We start by describing the general idea and then turn to the particulars.

General idea of the simulation part of the algorithm. Policy functions are typically functions of the state variables, that is, functions of *predetermined* endogenous variables and *exogenous* random variables. These functions incorporate the effect that prices have on agents' choices, but this formulation does not allow for prices to adjust if equilibrium does not hold *exactly* when choices of the individuals are aggregated. If used the true policy functions, then equilibrium would hold exactly by definition. Unfortunately, this will not be true for numerical approximations, not even for very accurate ones. Since long simulations are needed, errors accumulate, driving supply and demand further apart, unless these errors happen to be exactly zero on average. Our simulation procedure is such that equilibrium does hold exactly. The cost of achieving this is that actual prices, J and P , will be different from perceived prices, \tilde{J} and \tilde{P} and some of the actual individual choices will be different from those according to the original policy functions.⁶⁹ These are errors too, but there is no reason that these will accumulate.

Preliminaries. To simulate this economy, we need laws of motions for perceived prices, $\tilde{J}(q, z)$ and $\tilde{P}(q, z)$, as well as individual policy functions, q_i and M'_i , which are calculated as described in appendix 2.B.1.2. At the beginning of each period, we would also need the joint distribution of employment status, e_i , and cash on hand, x_i . This distribution is given by $\psi(\tilde{x}_i, e_i)$, where the tilde indicates that cash on hand is evaluated at perceived prices. The distribution is such that,

$$\int_{e_i} \int_{\tilde{x}_i} \tilde{x}_i d\psi_i = zq + (1 - \delta)q \frac{\tilde{J}}{\tilde{P}} + \frac{\bar{M}}{\tilde{P}}, \quad (2.17)$$

where the dependence of prices on the aggregate state variables has been suppressed. Below, we discuss how we construct a histogram for the cross-sectional distribution each period and show that this property is satisfied. We do not specify a joint distribution of equity and money holdings. As discussed below, we do know what level of beginning-of-period equity holdings, q_i , corresponds with what level of beginning-of-period cash on hand, \tilde{x}_i , in each period. In

⁶⁸Note that the first-moment of money holdings is constant, since money supply is constant.

⁶⁹Throughout this appendix, perceived variables have a tilde and actual outcomes do not.

particular, the distribution satisfies

$$\int_{e_i} \int_{\tilde{x}_i} q_i d\psi_i = q. \quad (2.18)$$

A household's cash on hand is given by

$$\tilde{x}_i = e_i(1 - \tau) \frac{\tilde{W}}{\tilde{P}} + (1 - e_i)\mu(1 - \tau) \frac{\tilde{W}}{\tilde{P}} + q_i \left(\frac{\tilde{D}}{\tilde{P}} + (1 - \delta) \frac{\tilde{J}}{\tilde{P}} \right) + \frac{M_i}{\tilde{P}}, \quad (2.19)$$

and the household can spend this on consumption and asset purchases, that is,

$$\tilde{x}_i = c_i + q'_i \frac{\tilde{J}}{\tilde{P}} + \frac{M'_i}{\tilde{P}}. \quad (2.20)$$

The government has a balanced budget each period, that is,

$$\tau = \mu \frac{1 - q}{q + \mu(1 - q)}. \quad (2.21)$$

Even if the numerical solutions for q'_i , M'_i , \tilde{J} , and \tilde{P} are very accurate, it is unlikely that equilibrium is *exactly* satisfied if we aggregate q'_i and M'_i across agents. To impose equilibrium exactly, we modify the numerical approximations for equity and money holdings such that they are no longer completely pinned down by exogenous random variables and predetermined variables, but instead depend directly—to at least some extent—on prices.⁷⁰ In the remainder of this section, we explain how we do this and how we solve for equilibrium prices.

Modification and imposing equilibrium. To impose equilibrium we adjust q'_i , M'_i , \tilde{J} , and \tilde{P} . The equilibrium outcomes are denoted by $q_{i,+1}$, $M_{i,+1}$, J , and P . The individual's demand for assets is modified as follows:

$$q_{i,+1} = \frac{\tilde{J}/\tilde{P}}{J/P} q'_i, \quad (2.22)$$

$$M_{i,+1} = \frac{P}{\tilde{P}} M'_i. \quad (2.23)$$

We will first discuss how equilibrium prices are determined and then discuss why this is a sensible modification. An important accuracy criterion is that this modification of the policy functions is small, that is, actual and perceived laws of motions are very similar.⁷¹

We solve for the actual law of motion for employment, q_{+1} , the number of new firms created, h , the amount spent on creating new firms in real terms, $v = hJ/P$, the market clearing

⁷⁰The policy functions $q(\tilde{x}_i, e_i, q, z)$ and $M(\tilde{x}_i, e_i, q, z)$ do depend on prices, but this dependence is captured by the aggregate state variables.

⁷¹As explained above, it is important to do a modification like this to ensure that equilibrium holds *exactly*, even if the solution is very accurate and the modification small.

asset price, J , and the market clearing price level, P , from the following equations:⁷²

$$q_{+1} = (1 - \delta)q + h, \quad (2.24)$$

$$h = \psi v^\eta (1 - q)^{1-\eta}, \quad (2.25)$$

$$v = hJ/P, \quad (2.26)$$

$$\begin{aligned} h &= \int_{e_i} \int_{\tilde{x}_i} (q_{+1}(\tilde{x}_i, e_i, q, z) - (1 - \delta)q_i) d\psi_i \\ &= \int_{e_i} \int_{\tilde{x}_i} \left(\frac{\tilde{J}}{J/P} \tilde{P} q'(\tilde{x}_i, e_i, q, z) - (1 - \delta)q_i \right) d\psi_i, \end{aligned} \quad (2.27)$$

$$\bar{M} = \int_{e_i} \int_{\tilde{x}_i} M_{+1}(\tilde{x}_i, e_i, q, z) d\psi_i = \int_{e_i} \int_{\tilde{x}_i} \frac{P}{\tilde{P}} M(\tilde{x}_i, e_i, q, z) d\psi_i. \quad (2.28)$$

Logic behind the modification. Recall that $q(\tilde{x}_i, e_i, q, z)$ and $m(\tilde{x}_i, e_i, q, z)$ are derived using perceived prices, $\tilde{J}(q, z)$ and $\tilde{P}(q, z)$. Now suppose that—in a particular period—aggregation of $q(\tilde{x}_i, e_i, q, z)$ indicates that demand for equity exceeds supply for equity. This indicates that $\tilde{J}(q, z)$ is too low in that period. By exactly imposing equilibrium, we increase the asset price and lower the demand for equity. Note that our modification is such that any possible misperception on prices does not affect the real amount each agent spends, but only the number of assets bought.

Throughout this section, the value of cash on hand that is used as the argument of the policy functions is constructed using *perceived* prices. In principle, the equilibrium prices that have been obtained could be used to update the definition of cash on hand and one could iterate on this until convergence. This would make the simulation more expensive. Moreover, our converged solutions are such that perceived and actual prices are close to each other, which means that this iterative procedure would not add much.

Equilibrium in the goods market. It remains to show that our modification is such that the goods market is in equilibrium as well. That is, Walras' law is not wrecked by our modification.

From the budget constraint we get that actual resources of agent i are equal to

$$x_i = e_i(1 - \tau) \frac{W}{P} + (1 - e_i)\mu(1 - \tau) \frac{W}{P} + \left(\frac{D}{P} + (1 - \delta) \frac{J}{P} \right) q_i + \frac{M_i}{P} \quad (2.29)$$

and actual expenditures are equal to

$$x_i = c_i + \frac{J}{P} q_{i,+1} + \frac{M_{i,+1}}{P}. \quad (2.30)$$

The value of c_i adjusts to ensure this equation holds. Aggregation gives

$$x = zq + \frac{J}{P} (1 - \delta) q + \frac{\bar{M}}{P} \quad (2.31)$$

⁷²Recall that we define variables slightly different and v is not the number of vacancies, but the amount spent on creating new firms.

and

$$x = c + \frac{J}{P} \int \int_{e_i \tilde{x}_i} q_{i,+1} d\psi_i + \frac{\int \int_{e_i \tilde{x}_i} M_{i,+1}}{P} = c + \frac{J}{P} q_{+1} + \frac{\bar{M}}{P}. \quad (2.32)$$

Equation (2.31) uses the definition of dividends and equation (2.18). Equation (2.32) follows from the construction of J and p .

Since

$$\frac{J}{P} q' - \frac{J}{P} (1 - \delta) q = v, \quad (2.33)$$

we get

$$zq = c + v, \quad (2.34)$$

which means that we have goods market clearing in each and every time period.

Implementation. To simulate the economy, we use the “non-stochastic simulation method” developed in Young (2010). This procedure characterizes the cross-sectional distribution of agents’ characteristics with a histogram. This procedure would be computer intensive if we characterized the cross-sectional distribution of both equity and bond holdings. Instead, we just characterize the cross-sectional distribution of cash on hand for the employed and unemployed. Let $\psi(\tilde{x}_{i,-1}, e_{i,-1})$ denote last period’s cross-sectional distribution of cash on hand and employment status. The objective is to calculate $\psi(\tilde{x}_i, e_i)$.

1. As discussed above, given $\psi(\tilde{x}_{i,-1}, e_{i,-1})$ and the policy functions, we can calculate last period’s equilibrium outcome for the total number of firms (jobs) carried into the current period, q , job-finding rate, $h_{-1}/(1 - q_{-1})$, last period’s prices, J_{-1} and P_{-1} , and for each individual the equilibrium asset holdings brought into the current period, q_i and M_i .
2. Current employment, q , together with the current technology shock, z , allows us to calculate perceived prices \tilde{J} and \tilde{P} .
3. Using the perceived prices together with asset holdings, q_i and M_i , we calculate perceived cash on hand conditional on last-period’s cash on hand and both the *past* and the *present* employment status. That is,

$$\begin{aligned} \tilde{x}(e_i, \tilde{x}_{i,-1}, e_{i,-1}) &= e_i(1 - \tau) \frac{\tilde{W}}{\tilde{P}} + (1 - e_i)\mu(1 - \tau) \frac{\tilde{W}}{\tilde{P}} \\ &\quad + q_i(\tilde{x}_{i,-1}, e_{i,-1}) \left(\frac{\tilde{D}}{\tilde{P}} + (1 - \delta) \frac{\tilde{J}}{\tilde{P}} \right) + \frac{M_i(\tilde{x}_{i,-1}, e_{i,-1})}{\tilde{P}}. \end{aligned}$$

4. Using last period’s distribution $\psi(\tilde{x}_{i,-1}, e_{i,-1})$ together with last-period’s transition probabilities, we can calculate the joint distribution of current perceived cash on hand, \tilde{x}_i , past employment status, and present employment status, $\hat{\psi}(\tilde{x}_i, e_i, e_{i,-1})$.
5. Next, we retrieve the current period’s distribution as

$$\psi(\tilde{x}_i, e_i) = \hat{\psi}(\tilde{x}_i, e_i, 1) + \hat{\psi}(\tilde{x}_i, e_i, 0). \quad (2.35)$$

6. Even though we never explicitly calculate a multi-dimensional histogram, in each period we do have information on the joint cross-sectional distribution of cash on hand at

perceived prices and asset holdings.

2.B.1.2 Solving for individual policy functions

When solving for the individual policy functions, aggregate laws of motion as specified in appendix 2.B.1.1 are taken as given. Individual policy functions for equity, $q'_i = q(\tilde{x}_i, e_i, q, z)$, and money, $M'_i = M(\tilde{x}_i, e_i, q, z)$, are obtained by iteration:

1. Using initial guesses for q'_i and M'_i , a policy function for consumption can be calculated from the agent's budget constraint:

$$c(\tilde{x}_i, e_i, q, z) = \tilde{x}_i - \frac{q'_i \tilde{J} + M'_i}{\tilde{P}}. \quad (2.36)$$

2. Conditional on the realizations of the aggregate shock and the agent's employment state, cash on hand and consumption in the next period can be calculated:

$$\tilde{x}'(e'_i, z') = e'_i(1 - \tau') \frac{\tilde{W}'}{\tilde{P}'} + (1 - e'_i)\mu(1 - \tau') \frac{\tilde{W}'}{\tilde{P}'} + q'_i \left(\frac{\tilde{D}'}{\tilde{P}'} + (1 - \delta) \frac{\tilde{J}'}{\tilde{P}'} \right) + \frac{M'_i}{\tilde{P}'}, \quad (2.37)$$

$$c'(e'_i, z') = c(\tilde{x}'(e'_i, z'), e'_i, q', z'). \quad (2.38)$$

3. Using the individual and aggregate transition probabilities, the expectations $\mathbb{E} \left[c'^{-\gamma} \frac{\tilde{P}}{\tilde{P}'} \right]$ and $\mathbb{E} \left[c'^{-\gamma} \frac{\tilde{D}' + (1 - \delta)\tilde{J}'}{\tilde{J}} \frac{\tilde{P}}{\tilde{P}'} \right]$, in the first-order conditions 2.3 and 2.4 can be calculated. Then, the first-order condition for equity holdings gives an updated guess for consumption of agents holding positive amounts of equity:

$$c^{new}(\tilde{x}_i, e_i, q, z) = \left(\beta \mathbb{E} \left[c'^{-\gamma} \frac{\tilde{D}' + (1 - \delta)\tilde{J}'}{\tilde{J}} \frac{\tilde{P}}{\tilde{P}'} \right] \right)^{-\frac{1}{\gamma}}. \quad (2.39)$$

The first-order condition for money gives an updated policy function for money:

$$M^{new}(\tilde{x}_i, e_i, q, z) = \tilde{P} \chi^{\frac{1}{\zeta}} \left(c^{new}(\tilde{x}_i, e_i, q, z)^{-\gamma} - \beta \mathbb{E} \left[c'^{-\gamma} \frac{\tilde{P}}{\tilde{P}'} \right] \right)^{-\frac{1}{\zeta}}. \quad (2.40)$$

The budget constraint in the current period gives the updated policy function for equity:

$$q^{new}(\tilde{x}_i, e_i, q, z) = \max \left(0, \frac{\tilde{x}_i \tilde{P} - c^{new}(\tilde{x}_i, e_i, q, z) \tilde{P} - M^{new}(\tilde{x}_i, e_i, q, z)}{\tilde{J}} \right). \quad (2.41)$$

For agents with a binding short-sale constraint, updated policy functions for consumption and money are instead calculated using only the first-order condition for money and the budget constraint:

$$c^{new,constraint}(\tilde{x}_i, e_i, q, z) = \left(\beta \mathbb{E} \left[c'^{-\gamma} \frac{\tilde{P}}{\tilde{P}'} \right] + \chi \left(\frac{M'_i}{\tilde{P}} \right)^{-\zeta} \right)^{-\frac{1}{\gamma}}, \quad (2.42)$$

$$M^{new,constraint}(\tilde{x}_i, e_i, q, z) = \tilde{x}_i \tilde{P} - c^{new,constraint}(\tilde{x}_i, e_i, q, z) \tilde{P}. \quad (2.43)$$

4. A weighted average of the initial guesses and the new policy functions is used to update the initial guesses. The procedure is repeated from step 1 until the differences between initial and updated policy functions become sufficiently small.

2.C Equivalence with standard matching framework

In the standard matching framework, new firms are created by “entrepreneurs” who post vacancies, \tilde{v}_t , at a cost equal to κ per vacancy. The number of vacancies is pinned down by a free-entry condition. In the description of the model above, such additional agents are not introduced. Instead, creation of new firms is carried out by investors wanting to increase their equity holdings.

Although, the “story” we tell is somewhat different, our equations can be shown to be identical to those of the standard matching model. The free-entry condition in the standard matching model is given by

$$\kappa = \frac{\tilde{h}_t J_t}{\tilde{v}_t P_t}, \quad (2.44)$$

where

$$\tilde{h}_t = \tilde{\psi} \tilde{v}_t^\eta u_t^{1-\eta}. \quad (2.45)$$

Each vacancy leads to the creation of $\tilde{h}_t / \tilde{v}_t$ new firms, which can be sold to households at price J_t .

Equilibrium in the equity market requires that the *net* demand for equity by households is equal to the supply of *new* equity by entrepreneurs, that is

$$\begin{aligned} & \int \int \int_{e_i, q_i, M_i} (q(e_i, q_i, M_i; s_t) - (1 - \delta) q_i) dF_t(e_i, q_i, M_i) \\ &= \tilde{\psi} \tilde{v}_t^\eta u_t^{1-\eta}. \end{aligned} \quad (2.46)$$

Using equations (2.44) and (2.45), this equation can be rewritten as

$$\begin{aligned} & \int \int \int_{e_i, q_i, M_i} (q(e_i, q_i, M_i; s_t) - (1 - \delta) q_i) dF_t(e_i, q_i, M_i) \\ &= \tilde{\psi}^{1/(1-\eta)} \left(\frac{J_t}{\kappa P_t} \right)^{\eta/(1-\eta)} u_t. \end{aligned} \quad (2.47)$$

This is equivalent to equation (2.13) if

$$\tilde{\psi} = \psi \kappa^\eta. \quad (2.48)$$

It only remains to establish that the number of new jobs created is the same in the two setups, that is,

$$h_t = \tilde{h}_t \quad (2.49)$$

or

$$\psi v_t^\eta u_t^{1-\eta} = \tilde{\psi} \tilde{v}_t^\eta u_t^{1-\eta}. \quad (2.50)$$

From equations (2.44) and (2.45), we get that

$$\tilde{v}_t = \left(\frac{\tilde{\psi} J_t}{\kappa P_t} \right)^{1/(1-\eta)} u_t. \quad (2.51)$$

Substituting this expression for \tilde{v}_t and the expression from equation (2.11) for v_t into equation (2.50) gives indeed that $h_t = \tilde{h}_t$. Moreover, the total amount spent on creating new firms in our representation, v_t , is equal to the number of vacancies times the posting cost in the traditional representation, $\kappa \tilde{v}_t$.

The focus of this paper is on the effect of negative shocks on the savings and investment behavior of agents in the economy when markets are incomplete. We think that our way of telling the story behind the equations has the following two advantages. First, there is only one type of investor, namely, the household and there are no additional investors such as zombie entrepreneurs (poor souls who get no positive benefits out of fulfilling a crucial role in the economy).⁷³ Second, all agents have access to investment in the same two assets, namely equity and the liquid asset, whereas in the standard labor market model there are households and entrepreneurs and they have different investment opportunities.

⁷³One could argue that entrepreneurs are part of the household, but with heterogeneous households the question arises which households they belong to.

Chapter 3

Efficiency of On-the-Job Search in a Search and Matching Model with Endogenous Job Destruction

3.1 Introduction

This chapter analyses the efficiency of on-the-job (OTJ) search in a search and matching model. OTJ search has been increasingly incorporated into search models to account for the large direct flows of workers from one job to another. For instance, Fallick and Fleischman (2004) find that nearly two fifths of new jobs in the U.S. are taken by previously employed workers.

I use a search model that allows for OTJ search and endogenous job destruction, which is close to the model outlined in Pissarides (2000). The main assumptions are that search is random and commitment in terms of wages and OTJ search is limited. Random search means that there is only one job market for firms and workers. Job seekers, however, differ in terms of their current employment status: unemployed workers benefit from finding a job more than employed job seekers, who are looking for a better job. If it were possible to separate these different groups of job seekers in a directed search model like Menzio and Shi (2011), then different submarkets would be created. It would be efficient to have relatively more vacancies for the unemployed to increase their chances to find a job. Without these separate markets, the efficient level of market tightness takes the potential gains of both the unemployed and the employed job-seekers into account. In section 3.2, the efficient level of market tightness is analysed in a simple model with OTJ search. In the spirit of the condition in Hosios (1990), I derive the optimal level of the worker's bargaining power. It is shown that, when OTJ search is present, it exceeds the level indicated by the Hosios rule. This is because the efficient level of market tightness is lower due to the trade-off described above that translates to the necessity of a higher worker's bargaining power.

The assumption of limited commitment creates another possible source of inefficiency: it is assumed that workers cannot commit to stay with their current firms but they can secretly perform OTJ search. Thereby, they impose an externality: the firm suffers a loss if its worker leaves for a better job. In contrast, a certain amount of OTJ search is efficient when the gain from a better job outweighs the loss of destroying the current job. This chapter finds that in

general there is too much OTJ search taking place. Only if firms can partially commit to the wage they pay, the efficient level can be reached. This is possible, because by committing to a higher wage, the firm can make it less attractive for the worker to perform OTJ search. A result of this mechanism, as pointed out by Shimer (2006), is that the bargaining set becomes non-convex; when the firm pays a higher wage, the probability of the worker quitting is reduced, which increases the surplus of a match. A suitable bargaining game is used to account for this non-convexity. The outcome reduces inefficient OTJ search and subsequently it is possible for the efficient level to be obtained.

Pissarides (1994) introduced OTJ search into a search model. My model differs from the one in Pissarides (2000) insofar as I do not allow for a joint OTJ search decision and I introduce partial commitment in terms of wages. The efficiency of OTJ search has been discussed in Stevens (2004). In her model, firms are allowed to post wage contracts that are more or less restricted. In the unrestricted case, workers pay a fee at the beginning of a match and then they get the whole rent. To allow for such contracts would yield an efficient OTJ search decision in my model, because employers do not care when a worker leaves. Hence, workers bear the full cost of destroying a match and choose the efficient threshold for performing OTJ search. Yet, Stevens (2004) concludes that turnover is too low, because workers do not incorporate the benefits a firm has from being matched. This is because she considers a partial equilibrium, in which firms and workers are matched at a constant rate. In general equilibrium with free entry of firms, however, these benefits are equal to the expected job creation cost and cancel out in an analysis of social welfare.

The result that there are too many vacancies is also obtained by Gautier, Teulings, and Van Vuuren (2010), who consider a circular model of heterogeneity. When posting vacancies, firms do not take into account the loss of other firms whose employees are to be potentially hired.

In a different context, Moen and Rosen (2011) have adapted the Hosios rule, when workers have private information. A higher worker bargaining power of workers is desirable as it induces them to exert more effort. In this sense, they provide an alternative explanation for the result in proposition 1.

There are two approaches for wage bargaining in models including OTJ search: Cahuc, Postel-Vinay, and Robin (2006) allow for counter-offers when employed workers get another job offer. Workers can then benefit from the competition between the two possible employers. In this chapter, I follow Shimer (2006), for instance, and exclude such competition.

Another approach for modelling OTJ search is the class of directed search models like in Menzies and Shi (2011) and Menzies and Shi (2010). The main advantage of the so called block recursive equilibrium they obtain, is the tractability of the solution. Introducing different sub-markets for job seekers then yields an efficient level of market tightness for each group.

The chapter is structured as follows: section 3.2 discusses a simple version of the model that abstracts from the decisions regarding OTJ search and job destruction and focuses on the efficient number of vacancies in the presence of OTJ search. The adaptation of the Hosios rule needed to obtain the efficient level of market tightness is derived. The full model with endogenous job destruction and OTJ search decision is analysed in section 3.3. In section 3.4, partial commitment is introduced, and the bargaining process is adapted to account for the non-convex bargaining set. Finally, section 5 concludes.

3.2 Simple model

3.2.1 Setup

The simple model I consider first is a special case of the endogenous job destruction model in Pissarides (2000) that allows for OTJ search. It is a continuous time search and matching model. The mass of workers is normalised to 1. Firms and workers are risk-neutral, having a discount rate r . Every firm-worker pair can have two possible levels of productivity: all jobs are created at productivity p . At rate λ , shocks arrive that reduce the productivity to px . Jobs of either productivity are destroyed at the exogenous rate δ .

Vacancies and jobseekers are matched according to a constant returns to scale matching function $m(u + e, v)$, where unemployed and employed job seekers are equally likely to be matched. From the matching function it follows that job seekers find a new job at rate $a(\theta) = m(1, \theta)$, where $\theta \equiv \frac{v}{u+e}$ denotes labour market tightness. This is the ratio of vacancies to the total number of job seekers, who consist of the unemployed u and the OTJ seekers e . The function $a(\theta)$ is increasing and concave in θ and the Inada conditions shall hold.¹ The arrival rate of employees to a vacancy is then given by the rate $q(\theta) = \frac{a(\theta)}{\theta}$.

The flow cost of maintaining a vacancy for a firm is pc . Unemployed workers receive a flow payoff of b . When employed workers want to perform OTJ search, they incur the cost v .²

The purpose of this simple model is to demonstrate how the efficient level of market tightness is influenced by the existence of OTJ search. I derive the generalisation of the Hosios (1990) rule for the worker's bargaining power that ensures an efficient market outcome. When there is OTJ search, the additional surplus of finding a new job for employed workers is lower than the surplus for the unemployed. Therefore, OTJ seekers are relatively more patient to find a better job. In a directed search model like Menzio and Shi (2011) this implies that in equilibrium the market for OTJ seekers features a lower market tightness (i.e. a lower job-finding rate), but a higher wage. As I do not allow for directed search, it is intuitive that the optimal market tightness for matching both unemployed and OTJ seekers with vacancies will lie between the low level of market tightness for OTJ seekers and the high level for the unemployed in a directed search model. To get an efficient market outcome, the worker's bargaining power will then have to be higher than the one suggested by the Hosios (1990) rule. Otherwise, market tightness in the decentralised economy is too high as an OTJ seeker does not fully incorporate the negative impact a change of jobs has on her previous employer.

As I am interested in this efficient level of market tightness, I want to eliminate other decisions made by the worker and the firm, notably when to perform OTJ search and when to destroy a job. Therefore, I assume that the parameters are such that it is both socially and individually optimal to perform OTJ search and not to destroy the job after a λ -shock has decreased the productivity of a match.³ In the full model from section 3.3 onwards, this assumption is relaxed and performing OTJ search becomes an endogenous choice.

First, the efficient solution given by the social planner's problem is derived, then the market outcome is derived and the two are compared.

¹That is $\lim_{\theta \rightarrow 0} a'(\theta) = \infty$ and $\lim_{\theta \rightarrow \infty} a'(\theta) = 0$.

²In this section I impose that OTJ search always takes place at the lower level of productivity. Without loss of generality, OTJ search costs could therefore be set to 0 and be deducted from output, px . I include them to make the equations better comparable to those in the following sections.

³The latter condition is simply $b \leq px - v$ as the low productivity is an absorbing state.

3.2.2 Social planner's problem

A social planner solves the following optimal control problem: the state variables are the fraction of unemployed, employed at productivity p , and employed at productivity px denoted by u , w_1 , and w_x , respectively ($u + w_1 + w_x = 1$). The only control variable in the simple model is labour market tightness θ . Given the initial distribution across states, she wants to maximise the present discounted value of the flow benefits for the unemployed plus output produced by employed less the cost of opening vacancies and of OTJ search:

$$\max \int_0^{\infty} e^{-rt} [ub + p(w_1 + xw_x) - pc\theta(u + w_x) - vw_x] dt, \quad (3.1a)$$

$$\dot{u} = \delta(w_1 + w_x) - a(\theta)u, \quad (3.1b)$$

$$\dot{w}_1 = -(\delta + \lambda)w_1 + a(\theta)(u + w_x), \quad (3.1c)$$

$$\dot{w}_x = -(\delta + a(\theta))w_x + \lambda w_1. \quad (3.1d)$$

The stock of unemployed workers increases, because employed workers lose their job at rate δ , and it decreases, because unemployed find a job at rate $a(\theta)$. The inflow into employment at productivity 1 is given by the fraction of unemployed and OTJ seekers who have found a new job, whereas a fraction $\delta + \lambda$ lose their job due to the shocks. Finally, the inflow into low productivity jobs is only due to λ -shocks whereas these jobs are lost at a rate $\delta + a(\theta)$ due to job destruction and OTJ search, respectively.

Denote by U , W_1 , and W_x the respective costates for the three state variables. The Hamiltonian in the current-value form is given by

$$\begin{aligned} H = & ub + p(w_1 + xw_x) - pc\theta(u + w_x) - vw_x + U[\delta(w_1 + w_x) - a(\theta)u] \\ & + W_1[-(\delta + \lambda)w_1 + a(\theta)(u + w_x)] + W_x[-(\delta + a(\theta))w_x + \lambda w_1]. \end{aligned} \quad (3.2)$$

The first-order condition for the optimal choice of market tightness is given by

$$a'(\theta)[u(W_1 - U) + w_x(W_1 - W_x)] = pc(u + w_x), \quad (3.3a)$$

$$(1 - \eta) \left[\frac{u}{u + w_x} (W_1 - U) + \frac{w_x}{u + w_x} (W_1 - W_x) \right] = \frac{pc}{q(\theta)}, \quad (3.3b)$$

where $\eta \equiv 1 - \frac{a'(\theta)\theta}{a(\theta)}$ is equal to the elasticity of the matching function with respect to the number of job seekers. The left hand side of the first equation describes the marginal benefit of a higher job arrival rate by increasing θ . In optimum it must be equal to the marginal cost on the right hand side. In the rearranged second line, one can see that the actual benefit of higher market tightness is given by the weighted average of the gains for an unemployed and for an employed job seeker, as there is only one market at which job seekers can be matched.

The costates' differential equations are given by

$$\dot{U} = rU - [b - pc\theta + a(\theta)(W_1 - U)], \quad (3.4a)$$

$$\dot{W}_1 = rW_1 - [p + \delta(U - W_1) + \lambda(W_x - W_1)], \quad (3.4b)$$

$$\dot{W}_x = rW_x - [px - pc\theta - v + a(\theta)(W_1 - W_x) + \delta(U - W_x)]. \quad (3.4c)$$

I focus on the description of a steady state, in which the states and costates are constant. Then the equations reduce to:

$$rU = b - pc\theta + a(\theta)(W_1 - U), \quad (3.5a)$$

$$rW_1 = p + \delta(U - W_1) + \lambda(W_x - W_1), \quad (3.5b)$$

$$rW_x = px - pc\theta - \nu + a(\theta)(W_1 - W_x) + \delta(U - W_x). \quad (3.5c)$$

The first equation states that the flow value of an unemployed worker is given by the unemployment benefit less the proportional cost of maintaining vacancies plus the option value of getting matched at rate $a(\theta)$. A match at the high level of productivity benefits from output p but at rate δ the match is destroyed and at rate λ the productivity is downgraded. For a low productivity match, the output px has to be reduced by the direct OTJ search costs ν and the cost for the vacancies. A better job is found at rate $a(\theta)$ and at rate δ the job breaks down. One can rearrange these value equations to obtain the respective surplus:

$$(r + a(\theta) + \delta)(W_x - U) = px - b - \nu, \quad (3.6a)$$

$$(r + a(\theta) + \lambda + \delta)(W_1 - W_x) = p(1 - x) + pc\theta + \nu, \quad (3.6b)$$

$$(r + a(\theta) + \delta)(W_1 - U) = p - b + pc\theta + \lambda(W_x - W_1). \quad (3.6c)$$

From the laws of motion for the states, one can obtain the distribution across states in a steady state:

$$u = \frac{\delta}{\delta + a(\theta)}, \quad (3.7a)$$

$$w_1 = \frac{a(\theta)}{\delta + \lambda + a(\theta)}, \quad (3.7b)$$

$$w_x = \frac{\lambda}{\delta + a(\theta)} w_1. \quad (3.7c)$$

A higher level of market tightness means that job seekers are matched faster, reducing the steady state value of unemployed and of low quality matches relative to high quality matches. The share of high quality matches then increases. Substituting the steady state value of the states and the values obtained from equations (3.6b) and (3.6c) into the first order condition (3.3b) for θ , implicitly yields the efficient steady state value for market tightness.

3.2.3 Decentralised economy

This subsection determines the equilibrium in the decentralised economy. The value of an unemployed and an employed worker at the two levels of productivity is denoted by U , W_1 , and W_x , respectively. Her wage is denoted by \bar{w}_1 and \bar{w}_x . This gives the stationary value equations

$$rU = b + a(\theta)(W_1 - U), \quad (3.8a)$$

$$rW_1 = \bar{w}_1 + \delta(U - W_1) + \lambda(W_x - W_1), \quad (3.8b)$$

$$rW_x = \bar{w}_x - \nu + a(\theta)(W_1 - W_x) + \delta(U - W_x). \quad (3.8c)$$

They differ from the social planner's equations (3.5a), (3.5b), and (3.5c) only insofar as job seekers do not directly bear the cost of vacancies and employed workers receive a wage instead of the whole output. They can be rewritten in terms of the surplus:

$$(r + a(\theta) + \delta)(W_1 - U) = \bar{w}_1 - b + \lambda(W_x - W_1), \quad (3.9a)$$

$$(r + a(\theta) + \lambda + \delta)(W_1 - W_x) = \bar{w}_1 - \bar{w}_x + v. \quad (3.9b)$$

Firms can open and maintain vacancies at cost pc , getting matched at rate $q(\theta)$. The value of a vacancy is denoted by V . The firm's value of a match is denoted by J_1 and J_x , respectively. Its flow profits are given by output less wage. A match with high productivity is destroyed at rate δ and downgraded at rate λ . A low productivity match is destroyed exogenously at rate δ and destroyed endogenously due to the worker leaving for another firm at rate $a(\theta)$. This yields the following value equations:

$$rV = -pc + q(\theta)(J_1 - V), \quad (3.10a)$$

$$rJ_1 = p - \bar{w}_1 - \delta J_1 + \lambda(J_x - J_1), \quad (3.10b)$$

$$rJ_x = px - \bar{w}_x - (\delta + a(\theta))J_x. \quad (3.10c)$$

The last equation can be solved for the value of a low productivity match:

$$J_x = \frac{px - \bar{w}_x}{r + \delta + a(\theta)}. \quad (3.11)$$

In equilibrium, rents from opening vacancies are exhausted such that the zero profit condition holds:

$$J_1 = \frac{pc}{q(\theta)}. \quad (3.12)$$

Using this condition, one can obtain the analogon of equation (3.9b) for the firm:

$$(r + a(\theta) + \lambda + \delta)(J_1 - J_x) = p(1 - x) - (\bar{w}_1 - \bar{w}_x) + pc\theta. \quad (3.13)$$

The equilibrium wage is determined by Nash bargaining. The worker gets a share β , with the remainder going to the firm. It is assumed that the wage is continuously renegotiated. Most importantly, it is adjusted after a negative productivity shock and it is not possible for the firm to prevent OTJ search by committing to a higher wage. From the sharing rule one obtains the respective conditions for the two types of matches:

$$(1 - \beta)(W_1 - U) = \beta J_1, \quad (3.14a)$$

$$(1 - \beta)(W_x - U) = \beta J_x. \quad (3.14b)$$

Taking the difference of the two and substituting from equations (3.9b) and (3.13), one obtains the equation for the wage differential:

$$(1 - \beta)(W_1 - W_x) = \beta(J_1 - J_x), \quad (3.15a)$$

$$(1 - \beta)(\bar{w}_1 - \bar{w}_x + v) = \beta[p(1 - x) - (\bar{w}_1 - \bar{w}_x) + pc\theta], \quad (3.15b)$$

$$\bar{w}_1 - \bar{w}_x = \beta[p(1 - x) + pc\theta] - (1 - \beta)v. \quad (3.15c)$$

Similarly, using equations (3.9a) and (3.10b), the wage at the high productivity is derived:

$$(1 - \beta) [\bar{w}_1 - b + \lambda (W_x - W_1)] = \beta [p - \bar{w}_1 + a(\theta) J_1 + \lambda (J_x - J_1)], \quad (3.16a)$$

$$(1 - \beta) (\bar{w}_1 - b) = \beta (p - \bar{w}_1 + pc\theta), \quad (3.16b)$$

$$\bar{w}_1 = (1 - \beta) b + \beta (p + pc\theta). \quad (3.16c)$$

This is the standard wage equation, which is not affected by the introduction of productivity shocks and OTJ search, as it is renegotiated after shocks. Substituting it into the equation for the wage differential gives the wage at the low level of productivity:

$$\bar{w}_x = (1 - \beta) (b + \nu) + \beta px. \quad (3.17)$$

The worker is compensated for foregoing unemployment benefits and incurring OTJ search cost ν . In addition, she gets a share of the output.

The job creation condition is obtained by equating the value J_1 from equation (3.10b) in combination with equation (3.11) to the value from the zero profit condition (3.12):

$$p - \bar{w}_1 + \lambda \frac{px - \bar{w}_x}{r + \delta + a(\theta)} = (r + \delta + \lambda) \frac{pc}{q(\theta)}. \quad (3.18)$$

Substituting the wage equations derived above into this equation, yields the equilibrium level of market tightness. Given market tightness, the steady state distribution across employment states is given by the same equations (3.7a), (3.7b), and (3.7c) as in the social planner's case. As labour market tightness is the only control variable in this model, the difference between the efficient outcome and the decentralised equilibrium can be analysed by comparing θ in the steady state. This is done in the next subsection.

3.2.4 Efficiency condition

Along the lines of Hosios (1990), I want to find the condition for the bargaining power that must hold in order to make the decentralised equilibrium efficient. This means that $\theta^{SP} = \theta^{DEC}$ shall hold where superscript *SP* denotes variables in the social planner's outcome and superscript *DEC* denotes variables in the decentralised equilibrium. The social planner's first order condition (3.3b) is

$$(1 - \eta) \left[\frac{u}{u + w_x} (W_1^{SP} - U^{SP}) + \frac{w_x}{u + w_x} (W_1^{SP} - W_x^{SP}) \right] = \frac{pc}{q(\theta^{SP})}. \quad (3.19)$$

In the decentralised economy, the zero profit condition (3.12) and the Nash sharing rule imply

$$\frac{pc}{q(\theta^{DEC})} = (1 - \beta) (W_1^{DEC} + J_1^{DEC} - U^{DEC}). \quad (3.20)$$

Note that when market tightness is the same in both cases the surplus of a job is also the same:

$$W_1^{SP}(\theta^{SP}) - U^{SP}(\theta^{SP}) = W_1^{DEC}(\theta^{SP}) + J_1^{DEC}(\theta^{SP}) - U^{DEC}(\theta^{SP}). \quad (3.21)$$

Combining the three equations, yields the condition for efficiency:

$$(1 - \eta) \left[\frac{u}{u + w_x} (W_1^{SP} - U^{SP}) + \frac{w_x}{u + w_x} (W_1^{SP} - W_x^{SP}) \right] = (1 - \beta) (W_1^{SP} - U^{SP}). \quad (3.22)$$

This can be rearranged to express the efficient level of bargaining power:

$$\beta = \eta + (1 - \eta) \frac{w_x}{u + w_x} \frac{W_x^{SP} - U^{SP}}{W_1^{SP} - U^{SP}}. \quad (3.23)$$

The Hosios condition without OTJ search is $\beta = \eta$. This is too low when OTJ search is present (i.e. $w_x > 0$).⁴ The expression $\frac{W_x^{SP} - U^{SP}}{W_1^{SP} - U^{SP}}$ describes how large the surplus of a low-productivity job is relative to a new job. The higher this ratio is and the more relevant it is due to more OTJ seekers, the more relevant is the loss of an existing surplus in low-productivity jobs due to successful OTJ search. Consequently, the larger is the optimal bargaining power of the worker. The intuition behind this result is that a high bargaining power increases wages, thereby reducing the profits of firms and consequently the number of vacancies they open. Otherwise there would be too many vacancies; firms do not account for the fact that when opening a vacancy the rate at which jobs are destroyed is endogenously increased, thereby imposing a negative externality on existing firms. This additional externality is corrected for by adapting the Hosios rule.

The result above proves the following proposition:

Proposition 1 *In an economy with on-the-job search, the efficient level of the worker's bargaining power is larger than the elasticity of the matching function with respect to the number of job seekers:*

$$\beta = \eta + (1 - \eta) \frac{w_x}{u + w_x} \frac{W_x^{SP} - U^{SP}}{W_1^{SP} - U^{SP}}.$$

In this section, I implicitly restricted the parameters such that OTJ search is efficient. In the next section, I discuss the general model, in which both the decision to destroy jobs and the decision to perform OTJ search are determined endogenously.

3.3 Model with endogenous job destruction and on-the-job search

I now relax the assumption that there are only two possible values for the productivity of a match. The model is similar to the one in chapter 4 of Pissarides (2000). Job destruction is endogenised along the lines of Mortensen and Pissarides (1994), which also makes the OTJ search decision endogenous. I assume that there is a finite number of productivity levels $p x_i$ ($x_1 = 1 > x_2 > \dots > x_n$) and jobs are still created at the maximum productivity p . Shocks that arrive at rate λ change the productivity of a match where the new productivity is drawn according to the probability mass function $g(x_i)$ ($\sum_{i=1}^n g(x_i) = 1$).⁵ Jobs are destroyed

⁴Note that the value of a low productivity match is bigger than the value of an unemployed by assumption in this section (i.e. $W_x^{SP} > U^{SP}$). In the model with endogenous job destruction it holds endogenously for all jobs that are not destroyed.

⁵Note that the simple model above is nested in this more general model by setting $x_2 = x$, $x_3 = 0$, $g(x_2) = \frac{\lambda}{\lambda + \delta}$, $g(x_3) = \frac{\delta}{\lambda + \delta}$, and $\bar{\lambda} = \lambda + \delta$.

endogenously if the new productivity level is below a certain reservation productivity such that all jobs with a lower productivity are destroyed. At each level of productivity, a worker can decide whether she wants to perform OTJ search. This will determine a second threshold such that OTJ search takes place for all matches with a lower productivity.

I start again with the discussion of the efficient solution and then I find the market outcome and compare the two.

3.3.1 Social planner's problem

Compared to the simple model in section 3.2.2, the planner has additional controls in the full model. For each level of productivity, let the dummy variable e_i denote the OTJ search decision such that OTJ search takes place when $e_i = 1$. Similarly, let the dummy variable d_i denote the destruction decision for matches with productivity x_i .⁶ The maximisation problem takes the following form:

$$\max \int_0^{\infty} e^{-rt} \left[ub + p \sum_{i=1}^n x_i w_i - pc\theta \left(u + \sum_{i=1}^n e_i w_i \right) - v \sum_{i=1}^n e_i w_i \right] dt, \quad (3.24a)$$

$$\dot{u} = \lambda \sum_{i=1}^n d_i g(x_i) \sum_{j=1}^n w_j - a(\theta) u, \quad (3.24b)$$

$$\dot{w}_1 = \lambda \left(g(x_1) \sum_{j=1}^n w_j - w_1 \right) + a(\theta) \left(u + \sum_{i=1}^n e_i w_i \right), \quad (3.24c)$$

$$\dot{w}_i = \lambda \left((1 - d_i) g(x_i) \sum_{j=1}^n w_j - w_i \right) - e_i a(\theta) w_i. \quad (3.24d)$$

As above, the flow utility is given by unemployment benefits and output less the costs for opening vacancies and OTJ search costs.⁷ Similar to the simple model, the inflow into unemployment is given by workers that lose their jobs after a productivity shock, and the outflow is given by the unemployed that are newly matched. The stock of workers at each level of productivity changes, because after a λ -shock workers' productivity gets changed. Additionally, there is an outflow if OTJ search takes place and at the highest productivity there is an inflow due to newly matched unemployed and OTJ seekers.

The value of an unemployed is still denoted by U , and W_i denotes the value of a match with productivity x_i . Appendix 3.A derives the Bellman equations, which are analogous to equations (3.5a), (3.5b), and (3.5c) in the simple model:

$$rU = b - pc\theta + a(\theta) (W_1 - U), \quad (3.25a)$$

$$rW_1 = p + \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_1 \right), \quad (3.25b)$$

$$rW_i = px_i + e_i [a(\theta) (W_1 - W_i) - pc\theta - v] + \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_i \right). \quad (3.25c)$$

⁶The linearity of the objective function and the constraints in e_i and d_i ensures that a corner solution is optimal.

⁷The laws of motion are written for stationary d_i . If for example an exogenous productivity shock raised the job destruction threshold, there would be a mass of jobs destroyed, which is not incorporated in the differential equations.

The main differences compared to the value equations obtained in the simple model are that the option value of OTJ search is only present when $e_i = 1$ and that there is a new option value after a λ -shock: the latter is the expectation over getting the value W_j of a match at a new productivity level or the value U of an unemployed worker, if the job is optimally destroyed in response to a λ shock.

Appendix 3.A derives the threshold $\bar{S}(\theta)$, such that OTJ search is optimal for all matches with $x \leq \bar{S}(\theta)$:

$$\bar{S}(\theta) = 1 - \frac{(pc\theta + \nu)(r + \lambda)}{a(\theta)p}. \quad (3.26)$$

Without fixed OTJ search costs ν , the threshold would be decreasing in θ because of the decreasing marginal returns of the matching function.⁸ The intuition behind this negative relationship is that the expected cost of getting a new job ($\frac{pc}{q(\theta)}$) is increasing in θ . This makes OTJ search profitable for fewer types of matches. The threshold for OTJ search determines the controls e_i given θ :

$$e_i = \begin{cases} 1 & \text{if } x_i \leq \bar{S}(\theta) \\ 0 & \text{if } x_i > \bar{S}(\theta) \end{cases}. \quad (3.27)$$

Now, the threshold for job destruction can be determined. There exists such a threshold, because matches are identical except for their productivity. Therefore, if matches at some productivity level are destroyed, also matches at lower levels of productivity have to be destroyed in the optimum. I assume that the parameters are such that there is at least some efficient OTJ search at the worst non-destroyed matches.⁹ Appendix 3.A derives that matches are destroyed if

$$\frac{p(1 - x_i) + \nu + pc\theta}{r + a(\theta) + \lambda} \geq W_1 - U. \quad (3.28)$$

The left hand side is the difference between W_1 and W_i which is determined by the actual difference in productivity and the cost of vacancies and OTJ search multiplied by the average (discounted) time until OTJ search is successful or a productivity shock has arrived. If this difference becomes larger than the surplus of a new match, the job is destroyed. Note that $(W_1 - U)$ itself depends on the threshold. Inequality (3.28) thus implicitly determines the threshold $R(\theta)$, which yields the controls:

$$d_i = \begin{cases} 1 & \text{if } x_i \leq R(\theta) \\ 0 & \text{if } x_i > R(\theta) \end{cases}. \quad (3.29)$$

The steady state distribution across employment states can be obtained analogously to the simple model, yielding:

$$u = \frac{\delta}{\delta + a(\theta)}, \quad (3.30a)$$

$$w_1 = g(x_1) \frac{a(\theta)}{\delta + a(\theta)} + \frac{a(\theta)(u + e)}{\lambda}, \quad (3.30b)$$

$$w_i = \begin{cases} \frac{\lambda}{\lambda + e_i a(\theta)} g(x_i) \frac{a(\theta)}{\delta + a(\theta)} & \text{if } d_i = 0 \\ 0 & \text{if } d_i = 1 \end{cases}. \quad (3.30c)$$

⁸In general $\text{sgn}[S'(\theta)] = -\text{sgn}[\eta pc\theta - \nu(1 - \eta)]$ so that it is decreasing for θ large enough. Otherwise, the fixed search cost is too high given a low probability of getting matched.

⁹If OTJ search costs were prohibitively high, the outcome would be identical to the model without OTJ search.

For a given value θ , the optimal values for d_i and e_i , the values for the surplus $W_1 - U$ and $W_1 - W_i$, as well as the distribution across states u and w_i have thus been determined. Using this for the first order condition (3.76b), one can find the possible steady states for labour market tightness θ . Depending on the parameters, there is not necessarily a unique steady state: intuitively, multiple steady states can occur, because higher labour market tightness usually lowers the OTJ search threshold. This could increase the average gain from matching the average job seeker as workers with higher productivity do not perform OTJ search anymore. The higher cost arising from more vacancies per job seeker could thus be compensated by a higher average gain. In the steady state with low market tightness, there would be relatively more unemployed workers but more employed workers at a high level of productivity as more workers perform OTJ search. If there is not a unique steady state, the optimal choice might even depend on the initial distribution across employment states.

3.3.2 Decentralised equilibrium without commitment

In this subsection, the market outcome is determined when there is no commitment in terms of wages or the OTJ search decision. Workers determine privately, if they want to perform OTJ search. As the wage is continuously renegotiated, they base their decisions on the market wage that is prevailing at the different levels of productivity.¹⁰ In this case, a higher wage cannot prevent the worker from OTJ search. This results in too much OTJ search in equilibrium as the worker does not take fully into account the loss that the firm incurs when she leaves. In contrast, the decision for job destruction is still constrained efficient, because wages are determined by Nash bargaining. Then the firm's or the worker's surplus is negative if and only if the combined surplus is negative.

The stationary value equations for the worker are now given by

$$rU = b + a(\theta)(W_1 - U), \quad (3.31a)$$

$$rW_1 = \bar{w}_1 + \lambda \sum_{j=1}^n g(x_j) (\max\{W_j, U\} - W_1), \quad (3.31b)$$

$$rW_i = \bar{w}_i + \lambda \sum_{j=1}^n g(x_j) (\max\{W_j, U\} - W_i) + \max\{a(\theta)(W_1 - W_i) - v, 0\}. \quad (3.31c)$$

They differ from the simple model's equations insofar as the worker can choose whether or not to perform OTJ search, and the value after a λ -shock is given by the expectation over the outcome in all possible states. The equations can be rearranged to obtain expressions for the surplus:

$$(r + \lambda + e_i a(\theta))(W_1 - W_i) = \bar{w}_1 - \bar{w}_i + e_i v, \quad (3.32a)$$

$$(r + \lambda + a(\theta))(W_1 - U) = \bar{w}_1 - b + \lambda \sum_{j=1}^n g(x_j) \max\{W_j - U, 0\}. \quad (3.32b)$$

The firms' equation for the value of a vacancy does not change. The equations for the firms' value of a match have to be similarly adjusted to account for the λ -shock. Additionally, a job

¹⁰Nevertheless, the privately optimal decision of the worker is anticipated in the bargaining problem, and the surplus is calculated based on it.

can be destroyed at rate $a(\theta)$ if the worker performs OTJ search (i.e. $e_i = 1$):

$$rV = -pc + q(\theta)(J_1 - V), \quad (3.33a)$$

$$rJ_1 = p - \bar{w}_1 + \lambda \sum_{j=1}^n g(x_j) (\max\{J_j, 0\} - J_1), \quad (3.33b)$$

$$rJ_i = px_i - \bar{w}_i + \lambda \sum_{j=1}^n g(x_j) (\max\{J_j, 0\} - J_i) - e_i a(\theta) J_i. \quad (3.33c)$$

The first equation implies the same zero profit condition as in the simple problem:

$$J_1 = \frac{pc}{q(\theta)}. \quad (3.34)$$

Making use of it, the firms' surplus of a high productivity match over a lower productivity match can be calculated in equilibrium:

$$(r + \lambda + e_i a(\theta))(J_1 - J_i) = p(1 - x_i) - (\bar{w}_1 - \bar{w}_i) + e_i pc\theta. \quad (3.35)$$

The wage is again determined by Nash-bargaining with the worker's share being β . The possible problem of a non-convex bargaining set as pointed out in Shimer (2006) does not arise here: the continuation value and the OTJ search decision only depend on the equilibrium values. Therefore, the wage bargaining only splits the output less possible OTJ search costs in the current (infinitesimal) period.¹¹ As the combined surplus is not influenced by the outcome of the bargaining process, the bargaining set is convex. In the next section, I relax the assumption of no commitment leading to a possible non-convexity.

The analogon of the wage differential (3.15c) can be derived from equations (3.32a) and (3.35) using the Nash sharing rule:

$$\bar{w}_1 - \bar{w}_i = \beta [p(1 - x_i) + e_i pc\theta] - (1 - \beta) e_i v. \quad (3.36)$$

Now, the worker is only partially compensated for the OTJ search costs and compensates for the cost of vacancies if OTJ search takes place. Similarly, equations (3.32b) and (3.33b) can be used to derive the same wage equations at the highest productivity as in the simple model:

$$\bar{w}_1 = (1 - \beta)b + \beta(p + pc\theta) \quad (3.37)$$

Substituting it into the wage differential, yields the general wage equation:

$$\bar{w}_i = (1 - \beta)(b + e_i v) + \beta(px_i + (1 - e_i)pc\theta). \quad (3.38)$$

Its interpretation is the same as in the simple model: the worker is compensated for potentially incurred OTJ search costs and foregoing unemployment benefits and she gets her share of the output and the cost of vacancies if she does not search.

Before I determine the thresholds for OTJ search and job destruction, I can calculate the

¹¹The assumption of continuous time is not crucial for this result. In discrete time, if the worker can either hide her OTJ search efforts from the firm, or take this decision after the wage is set, the same argument applies. Paying a higher wage would not prevent her from performing OTJ search to find a better job for the next period. The firm would have to be able to commit to the wage paid in the next period to prevent OTJ search. This is the case in Shimer (2006), where the wage remains fixed for the duration of the match after it has been initially bargained.

expressions for the surplus using equations (3.32a) and (3.35) as well as (3.31b) and (3.33b):

$$S_1 - S_i = \frac{p(1 - x_i) + e_i(pc\theta + v)}{r + \lambda + e_i a(\theta)}, \quad (3.39)$$

$$S_1 = \frac{p - b + \lambda \sum_{j=1}^n g(x_j) \max\{S_j, 0\} + pc\theta}{r + \lambda + a(\theta)}. \quad (3.40)$$

I want to find an equilibrium, in which S_i , and hence W_i and J_i , are monotonic in i . Monotonicity of the surplus implies that there is a threshold R such that jobs are destroyed if $x_i \leq R$. Monotonicity of the worker's share implies that there is a threshold \bar{S} such that OTJ search is performed if $x_i \leq \bar{S}$.¹²

The worker will want to perform OTJ search if the expected gains from finding a better job outweigh her OTJ search costs ($a(\theta)(W_1 - W_i) \geq v$). This yields the condition for OTJ search:

$$a(\theta) \beta \frac{p(1 - x_i) + e_i(pc\theta + v)}{r + \lambda + e_i a(\theta)} \geq v. \quad (3.41)$$

For the individual decision to be consistent with the equilibrium outcome e_i , the condition must hold when $e_i > 0$ and the opposite must hold when $e_i < 1$. For the corner solutions, this yields the two conditions

$$a(\theta) \beta \frac{p(1 - x_i) + pc\theta + v}{r + \lambda + a(\theta)} \geq v, \quad (3.42)$$

when there is OTJ search and

$$a(\theta) \beta \frac{p(1 - x_i)}{r + \lambda} \leq v, \quad (3.43)$$

when there is no OTJ search. The first inequality defines a threshold such that OTJ search is an equilibrium strategy when $x_i \leq \bar{S}_1(\theta)$:

$$\bar{S}_1(\theta) = 1 + c\theta - \frac{r + \lambda + (1 - \beta)a(\theta)}{a(\theta)\beta p} v. \quad (3.44)$$

The second inequality defines a threshold such that no OTJ search is an equilibrium strategy when $x_i \geq \bar{S}_0(\theta)$:

$$\bar{S}_0(\theta) = 1 - \frac{r + \lambda}{a(\theta)\beta p} v. \quad (3.45)$$

If $\bar{S}_0(\theta) = \bar{S}_1(\theta)$ held, there would be a unique OTJ search strategy. However, if $\bar{S}_0(\theta) < \bar{S}_1(\theta)$ holds, there are multiple equilibria for $x_i \in [\bar{S}_0(\theta), \bar{S}_1(\theta)]$: if nobody is searching in equilibrium, this implies a higher wage, and it is optimal not to search. But if everyone is searching, the wage is lower, and it is indeed optimal for each worker to search. This is the case that will generally happen as $\bar{S}_0(\theta) < \bar{S}_1(\theta)$ is equivalent to

$$\beta pc\theta > (1 - \beta)v. \quad (3.46)$$

This condition holds in equilibrium, because from the zero profit condition and the Nash sharing rule it follows that the worker's surplus at the best job is $\frac{\beta}{1-\beta} \frac{pc}{q(\theta)}$. The worker's gain from finding a new job must be smaller than this: As her current job is not destroyed, there is a

¹²It could be that a non-monotonic surplus is reconcilable with a non-monotonic OTJ search decision because of the existence of multiple equilibria, as shown below. This would create a further inefficiency of the market outcome. But as such an equilibrium seems arbitrary, I exclude it in the further analysis.

positive surplus, which reduces the gains from finding a new job. Therefore, the worker's gain from OTJ search is less than $\frac{\beta}{1-\beta}pc\theta$. In contrast, the worker's gain must be larger than her cost v . Hence, condition (3.46) holds in any equilibrium, in which OTJ search takes place. Then there are multiple equilibria if $\bar{S}_0(\theta) < 1$ holds, which is equivalent to positive OTJ search costs ($v > 0$). The threshold for OTJ search must thus be in the interval

$$\bar{S} \in \left[1 - \frac{r + \lambda}{a(\theta)\beta p}v, 1 + c\theta - \frac{r + \lambda + (1 - \beta)a(\theta)}{a(\theta)\beta p}v \right]. \quad (3.47)$$

For now, I do not specify which equilibrium threshold shall be chosen. But the next subsection shows that all possible thresholds are larger than the efficient one.

A job is destroyed if its surplus becomes negative. I restrict myself to the cases in which there is at least some OTJ search as in the social planner's discussion.¹³ Using equations (3.39) when $e_i = 1$ and (3.40), it follows that the surplus of a match becomes negative if

$$-px_i + v \geq -b + \lambda \sum_{j=1}^n g(x_j) \max\{S_j, 0\}. \quad (3.48)$$

Using equation (3.40) and the no profit condition, the expected surplus after a λ -shock can be determined:

$$\lambda \sum_{j=1}^n g(x_j) \max\{S_j, 0\} = (r + \lambda + a(\theta)) \frac{1}{1 - \beta} \frac{pc}{q(\theta)} - p + b - pc\theta. \quad (3.49)$$

Combining the last two equations, gives the condition for job destruction:

$$p(1 - x_i) + v \geq (r + \lambda) \frac{1}{1 - \beta} \frac{pc}{q(\theta)} + \frac{\beta}{1 - \beta} pc\theta. \quad (3.50)$$

It is monotonous in x and thus defines the job destruction threshold

$$R(\theta) = 1 - \frac{r + \lambda}{1 - \beta} \frac{c}{q(\theta)} - \left(\frac{\beta}{1 - \beta} c\theta - \frac{v}{p} \right), \quad (3.51)$$

and the respective dummy variable for job destruction

$$d_i = \begin{cases} 1 & \text{if } x_i \leq R(\theta) \\ 0 & \text{if } x_i > R(\theta) \end{cases}. \quad (3.52)$$

To find the equilibrium level of market tightness, the analogous job creation condition to equation (3.18) has to be derived. Combining equations (3.33b) and (3.39), making use of the Nash sharing rule, one obtains

$$(r + \delta) J_1 = p - \bar{w}_1 - \lambda \sum_{j=1}^n g(x_j) \frac{(1 - \beta) [p(1 - x_j) + e_j(pc\theta + v)]}{r + \lambda + e_j a(\theta)} (1 - d_j), \quad (3.53)$$

where $\delta \equiv \lambda \sum_{j=1}^n g(x_j) I(J_j \leq 0)$ is the probability of job destruction in equilibrium. Substi-

¹³This assumption is even less restrictive here because there will be more OTJ search than in the efficient case.

tuting the zero profit condition and the wage equation (3.37), yields the condition

$$(r + \delta) \frac{pc}{q(\theta)} = (1 - \beta)(p - b) - \beta pc\theta - \lambda \sum_{j=1}^n g(x_j) \frac{(1 - \beta) [p(1 - x_j) + e_j(pc\theta + \nu)]}{r + \lambda + e_j a(\theta)} (1 - d_j). \quad (3.54)$$

Using the thresholds for OTJ search and job destruction, this equation implicitly defines labour market tightness in the decentralised equilibrium.

3.3.3 Comparison of the decentralised equilibrium with the efficient outcome

In this subsection, I show that the market outcome is in general inefficient regardless of the level of bargaining power. This stems from OTJ search taking place in excess, as the worker does not fully incorporate the loss for a firm if she changes the job. Recall the OTJ search thresholds derived in equations (3.26) and (3.47):

$$\bar{S}^{SP}(\theta) = 1 - \frac{(pc\theta + \nu)(r + \lambda)}{a(\theta)p}, \quad (3.55a)$$

$$1 - \frac{r + \lambda}{a(\theta)\beta p} \nu \leq \bar{S}^{DEC}(\theta) \leq 1 + c\theta - \frac{r + \lambda + (1 - \beta)a(\theta)}{a(\theta)\beta p} \nu. \quad (3.55b)$$

An efficient search threshold would only be possible if the lower bound for \bar{S}^{DEC} was lower than \bar{S}^{SP} , which is equivalent to

$$\beta pc\theta \leq (1 - \beta)\nu \quad (3.56)$$

But in (3.46), exactly the opposite inequality was shown to hold in equilibrium. As a result, if market tightness in the decentralised economy was at the efficient level, the OTJ search threshold would be higher in the decentralised economy and too much OTJ search would be taking place. As this holds regardless of the worker's bargaining power, one can conclude that the outcome in this economy will generally differ from the efficient outcome.¹⁴

Proposition 2 *The decentralised outcome is not efficient: conditional on the level of market tightness, there is too much on-the-job search taking place.*

The reason for this result is that without wage commitment, a worker will perform OTJ search whenever the expected gain from finding a better job exceeds her search cost ν . In particular, in the absence of OTJ search costs there will be OTJ search for all but the highest level of productivity. From a social point of view, the gain of finding a better job is smaller as the old job is destroyed and the firm loses its part of the rent. Even in the absence of OTJ search costs, the social planner's threshold (3.26) shows that not all workers should seek better jobs. If the level of productivity is close to the optimal, it is not worth maintaining the additional vacancies to potentially find a better job for the worker. This result differs from Pissarides

¹⁴As I assume discrete levels of productivity, the efficient outcome can be achieved if there is no possible level of productivity between the two different thresholds. In general, the difference between the two outcomes will be determined by the probability of a match having a productivity level between the efficient and the decentralised threshold.

(2000) and Stevens (2004) who both conclude that there is too little search taking place. First, Pissarides (2000) assumes that the decision for OTJ search is taken jointly by the firm and the worker such that their surplus is maximised. In this section however, there is no commitment and the worker will base the OTJ search decision on her individual optimality conditions. Second, they do not take into account or do not model the cost of maintaining vacancies for OTJ search. In particular, in the absence of (direct) OTJ search costs, the efficient level of OTJ search is that any worker with less than optimal productivity performs OTJ search. This does not take into account that either more vacancies are needed to maintain the level of market tightness or there is a congestion effect decreasing the chances of other job seekers who need a (better) job more strongly.

The second decision of the agents in the economy that could potentially differ from the efficient level, is job destruction. Nash sharing, however, ensures that this decision is jointly optimal and there is no (direct) distortion. In other words, a match will be destroyed endogenously whenever its surplus vanishes. This is the same action a social planner would take, conditional on the OTJ search threshold and market tightness.

In the next section, I allow for (partial) commitment, which can lead to an efficient OTJ search decision and the first best outcome.

3.4 Decentralised equilibrium with (partial) commitment

3.4.1 Setup

I extend the model from above to allow for partial commitment in terms of wages: I assume that, unless productivity changes, the firm and the worker can only renegotiate the wage after a shock with arrival rate ω has arrived. The parameter ω represents the level of commitment, such that if it is 0, there is perfect commitment (until a productivity shock arrives). As ω becomes larger, the two parties are allowed to renegotiate after increasingly shorter intervals. The intuitive consequence of wage commitment is that it becomes feasible for the firm to prevent OTJ search by paying and committing to a higher wage. Without commitment, it was shown that there is too much OTJ search taking place as the worker does not care about the loss the firm incurs. Below, I show that for any finite level of ω , it is possible to achieve the efficient OTJ search threshold in equilibrium. This result hinges on choosing an appropriate subgame perfect equilibrium in the bargaining game over a non-convex bargaining set.

Denote by $W_i(\bar{w}_i)$ and $J_i(\bar{w}_i)$ the worker's and the firm's value of a match when the wage is fixed at \bar{w}_i . The equilibrium values that are expected after renegotiation are denoted by W_i and J_i , respectively. Adapting the stationary value equations then yields

$$\begin{aligned} rW_i(\bar{w}_i) &= \bar{w}_i + \lambda \sum_{j=1}^n g(x_j) (\max\{W_j, U\} - W_i(\bar{w}_i)) \\ &\quad + e_i [a(\theta)(W_1 - W_i(\bar{w}_i)) - \nu] + \omega(W_i - W_i(\bar{w}_i)), \end{aligned} \quad (3.57)$$

$$rJ_i(\bar{w}_i) = px_i - \bar{w}_i + \lambda \sum_{j=1}^n g(x_j) (\max\{J_j, 0\} - J_i(\bar{w}_i)) - e_i a(\theta) J_i(\bar{w}_i) + \omega(J_i - J_i(\bar{w}_i)). \quad (3.58)$$

The worker's optimal search decision is $e_i = 1$ if $a(\theta)(W_1 - W_i(\bar{w}_i)) > \nu$ such that the ex-

pected gain exceeds her search cost.

Before the outcome of the bargaining process is analysed in detail, note that in an equilibrium, in which everyone sets the same wage, the level of commitment ω will not directly influence the wage. It will rather influence the OTJ search decision. It was shown above that there might be multiple equilibria without commitment: there can be an equilibrium featuring a low wage and OTJ search that is Pareto-dominated by an equilibrium with a higher wage and no OTJ search. Since the lack of commitment does not make it feasible for the firm to increase the wage by as much as is needed to prevent the worker from performing OTJ search, these bad equilibria could not be ruled out. With wage commitment, it becomes possible for a match to set a higher wage and prevent the worker from performing OTJ search, even if all other workers with the same level of productivity perform OTJ search. Furthermore, new equilibria are feasible, when the bargaining game is adjusted to the non-convex bargaining set.

The bargaining set becomes potentially non-convex, because raising the wage can prevent the worker from OTJ search and make both the firm and the worker better off. To see this formally, the dependence of the worker's and the firm's surplus on the bargained wage can be analysed. The worker's value equation can be rewritten as

$$W_i(\bar{w}_i) = \frac{\bar{w}_i + e_i [a(\theta) W_1 - \nu] + c_i^W}{r + \lambda + \omega + e_i a(\theta)}, \quad (3.59)$$

where $c_i^W \equiv \lambda \sum_{j=1}^n g(x_j) \max\{W_j, U\} + \omega W_i$ does not depend on the outcome of the bargaining. The worker's surplus is strictly increasing in the bargained wage \bar{w}_i . The threshold \bar{w}_i^{NS} for OTJ search is given by the condition $a(\theta)(W_1 - W_i(\bar{w}_i^{NS})) = \nu$ which yields:

$$\bar{w}_i^{NS} = (r + \lambda + \omega) \left(W_1 - \frac{\nu}{a(\theta)} \right) - c_i^W. \quad (3.60)$$

The worker's optimal choice of this threshold ensures that her value function is continuous at \bar{w}_i^{NS} . But there is a kink as the slope for lower wages (i.e. when there is OTJ search) is $\frac{1}{r + \lambda + \omega + a(\theta)}$ which is less than the slope for higher wages ($\frac{1}{r + \lambda + \omega}$, when there is no OTJ search). Hence, the worker's surplus is continuous, piecewise linear, increasing, and convex in her wage. Using the condition $W_i(\bar{w}_i^R) = U$, her reservation wage \bar{w}_i^R is given by:

$$\begin{aligned} \bar{w}_i^R &= (r + \lambda + \omega) U - e_i [a(\theta)(W_1 - U) - \nu] - c_i^W \\ &= rU - e_i [a(\theta)(W_1 - U) - \nu] - \lambda \sum_{j=1}^n g(x_j) \max\{W_j - U, 0\} - \omega(W_i - U). \end{aligned} \quad (3.61)$$

Likewise, the firm's surplus is given by:

$$(r + \lambda + \omega + e_i a(\theta)) J_i(\bar{w}_i) = px_i - \bar{w}_i + \lambda \sum_{j=1}^n g(x_j) \max\{J_j, 0\} + \omega J_i, \quad (3.62a)$$

$$J_i(\bar{w}_i) = \frac{-\bar{w}_i + c_i^J}{r + \lambda + \omega + e_i a(\theta)}, \quad (3.62b)$$

where the maximum wage that the firm is willing to pay, c_i^J , does not depend on the outcome

of the bargaining:

$$c_i^J = px_i + \lambda \sum_{j=1}^n g(x_j) \max\{J_j, 0\} + \omega J_i. \quad (3.63)$$

The employer's surplus from bargaining is decreasing in the wage, piecewise linear, but not continuous at \bar{w}_i^{NS} : the limit from below (i.e. when there is OTJ search) is $\frac{-\bar{w}_i^{NS} + c_i^J}{r + \lambda + \omega + a(\theta)}$, whereas the limit from above is higher ($\frac{-\bar{w}_i^{NS} + c_i^J}{r + \lambda + \omega}$). This discontinuity reflects that at one side of the threshold the employer risks to lose the worker due to OTJ search but marginally increasing her wage prevents OTJ search and thus increases the expected job duration. Denote by $\bar{w}_i^S < \bar{w}_i^{NS}$ the maximum wage the firm is willing to pay when there is OTJ search without being better off by raising the wage to \bar{w}_i^{NS} . Using the firm's surplus in equation (3.62b), this condition

$$J_i(\bar{w}_i^S | e_i = 1) = J_i(\bar{w}_i^{NS} | e_i = 0) \quad (3.64)$$

yields

$$\bar{w}_i^S = (r + \lambda + \omega + a(\theta)) \left(W_1 - \frac{v}{a(\theta)} \right) - c_i^W - \frac{a(\theta)(c_i^W + c_i^J)}{r + \lambda + \omega}. \quad (3.65)$$

Therefore no wage in the interval $(\bar{w}_i^S, \bar{w}_i^{NS})$ should be the outcome of the bargaining process as both the firm and the worker could do better by raising the wage to \bar{w}_i^{NS} .

Appendix 3.B shows that it is jointly optimal to perform OTJ search if the following condition holds:

$$a(\theta)(W_1 - W_i(\bar{w}_i) - J_i(\bar{w}_i)) \geq v. \quad (3.66)$$

This is a stronger condition than the worker's individual optimality condition:

$$a(\theta)(W_1 - W_i(\bar{w}_i)) \geq v. \quad (3.67)$$

This is, because the worker does not take into account the firm's loss, if she finds a new job. Hence, if it is jointly optimal to perform OTJ search, it will also be individually rational to do so for the worker at any wage that is feasible for the firm.¹⁵ The outcome of the bargaining process will not influence the OTJ search decision and the firm and the worker bargain over the surplus S_i^S . Nash bargaining can be applied to determine the wage in this case as in the previous sections. If on the other hand it is not jointly optimal to perform OTJ search, it is still individually rational for the worker to do so if the wage is below \bar{w}_i^{NS} . Then, the bargaining set becomes non-convex and standard Nash bargaining cannot be applied. Similar to Shimer (2006), I use an alternating offer bargaining game to determine the wage in this case. It is described and solved in the next subsection.

In contrast to the case of no wage commitment, the bargaining surplus depends on the search decision implied by the bargained wage. In the previous section, it did not matter, because the wage was renegotiated after search, making the search decision only dependent on the equilibrium wage, which is taken as given. To see the dependence on ω , the difference between the surplus with and without OTJ search can be calculated from equations (3.86) and

¹⁵Formally, this means that $c_i^J \leq \bar{w}_i^{NS}$ holds. Then, the worker will search even if the firm offers the maximum wage such that its rent becomes 0.

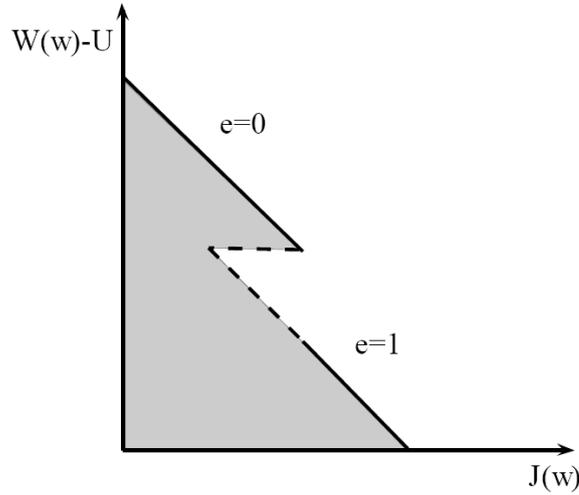


Figure 3.1: Non-convex bargaining set when OTJ search does not maximise the joint surplus: When the worker's share is sufficiently high, she does not perform OTJ search, which maximises $W - U + J$. Below the cutoff, the worker searches OTJ and the firm's value drops discontinuously. The solidly drawn frontier of the bargaining set can be obtained in an alternating bargaining game according to Lemma 3.

(3.87):

$$S_i^{NS} - S_i^S = \frac{a(\theta)}{r + \lambda + \omega + a(\theta)} \left[\frac{-p(1-x_i) + (r + \lambda)S_1 + \omega S_i}{r + \lambda + \omega} - \left(W_1 - U - \frac{v}{a(\theta)} \right) \right]. \quad (3.68)$$

As $\omega \rightarrow \infty$, which corresponds to the case of no commitment, the term in brackets is bounded. Therefore, the difference between S_i^S and S_i^{NS} vanishes in the limit. This proves that, as the wage is renegotiated more frequently, the bargained wage becomes less important and the wages in equilibrium dominate the worker's OTJ search decision.

The next subsection describes the bargaining game used to deal with the potential non-convexity.

3.4.2 Wage bargaining when the bargaining set is non-convex

An example for the non-convex bargaining set when OTJ search is not jointly optimal and thus socially inefficient is depicted in figure 3.1. This set can in general be described as follows. Denote by $x \in [0, S]$ the worker's surplus and by x^{NS} the threshold at which the worker's search behaviour changes. Let T be the joint loss when there is OTJ search. Then the firm's surplus is given by

$$J(x) = \begin{cases} S - x & \text{if } x^{NS} \leq x \leq S \\ S - x - T & \text{if } 0 \leq x < x^{NS} \end{cases}. \quad (3.69)$$

Following Shimer (2006), I consider an alternating offer bargaining game extending the Rubinstein (1982) model: in each stage of the bargaining process the worker is allowed to make an offer with probability $\beta \in [0, 1]$ and the firm makes an offer with probability $1 - \beta$.¹⁶ If the receiver of the offer accepts, bargaining ends with the respective payoffs. If the offer is

¹⁶ β is again interpreted as the worker's bargaining power.

rejected, the bargaining process breaks down with probability $1 - \delta$ before the next stage.

Lemma 3 *When δ is approaching 1, the following values for the worker's surplus can be obtained in a subgame perfect equilibrium:*

- $x = \beta S$ if $\beta > \frac{x^{NS}}{S}$
- $x = x^{NS}$ if $\beta \leq \frac{x^{NS}}{S}$
- $x = \beta(S - T)$ if $\beta < \frac{x^{NS} - T}{S - T}$

Proof. See appendix 3.C ■

Note that there are multiple equilibria if $\beta < \frac{x^{NS} - T}{S - T}$ holds: both an equilibrium without OTJ search (the surplus maximizing outcome) as well as one with OTJ search and Nash sharing could be obtained. Even low β , this is a high bargaining power of the firm, cannot break the surplus maximizing equilibrium, because potential marginal gains for the firm are offset by a discrete loss for the worker. Therefore, as δ approaches 1, the worker would not accept such an offer, even if she potentially had to wait a long time until making an offer again.

The selection of one of the multiple equilibria when $\beta < \frac{x^{NS} - T}{S - T}$ determines the efficiency of the decentralised solution. In the following section, I choose the equilibrium that leads to an efficient OTJ search decision. This means that $x = x^{NS}$ is chosen for all $\beta \leq \frac{x^{NS}}{S}$. The worker's surplus is large enough to prevent her from performing OTJ search. What makes this equilibrium attractive is that it maximises the joint surplus and therefore yields an optimal search decision given market tightness. At the same time, it also maximises the worker's share that does not fall below $\frac{x^{NS}}{S}$.¹⁷

To conclude the discussion of a non-convex bargaining set, note that mixed lotteries could potentially Pareto improve the outcome: by mixing over $x = 0$ and $x = x^{NS}$, both agents can be made better off compared to the equilibria in which OTJ search takes place, when $\beta < \frac{x^{NS}}{S}$. In this case, the bargaining set becomes convex and Nash bargaining can be applied. Following Shimer (2006), I rule out this possibility, as a wage lottery being the outcome of a bargaining process might be difficult to interpret and implement.

3.4.3 Determining the equilibrium

In section 3.4 above, the reservation wage \bar{w}_i^R , the maximum wage c_i^J , the maximum wage with OTJ search \bar{w}_i^S , and the minimum wage without OTJ search \bar{w}_i^{NS} were determined in equations (3.61), (3.63), (3.65), and (3.60), respectively. These values only depend on the equilibrium outcome but not on the individually bargained wage. It always holds that \bar{w}_i^S is smaller than \bar{w}_i^{NS} but the relative ranking of reservation wage and maximum wage determine the outcome of the bargaining process. I assume that $\bar{w}_i^R \leq \bar{w}_i^{NS}$ holds such that OTJ search takes place at some levels of productivity in equilibrium.¹⁸ Three cases have to be distinguished:

1. Job destruction zone ($\bar{w}_i^R > c_i^J$): if the reservation wage is larger than the maximum wage the firm is willing to pay, the result is immediate job destruction. When OTJ search takes

¹⁷ Alternative selections of an equilibrium in the bargaining game could be the equilibrium that maximises the firm's share, i.e. $x = \beta(S - T)$ if $\beta < \frac{x^{NS} - T}{S - T}$, or the one that maximises the weighted product of returns $x^\beta J(x)^{1-\beta}$.

¹⁸ The condition $\bar{w}_i^R \leq \bar{w}_i^{NS}$ is equivalent to $a(\theta)(W_1 - U) \geq v$ as the reservation wage reduces the worker's value to U and she will only want to search if the option value of finding a job is larger than the cost. If it did not hold, the OTJ search costs are prohibitively large and there could never be OTJ search as $W_i \geq U$ holds.

place for low levels of productivity, jobs are destroyed if their output plus the option value of getting hit by a λ -shock is less than unemployment benefit plus OTJ search costs. This is the same condition as in the case without commitment.

2. On-the-job search zone ($\bar{w}_i^R \leq c_i^J < \bar{w}_i^{NS}$): this is the case, in which it is jointly optimal for the worker and the firm and also individually rational for the worker to perform OTJ search. Therefore, the worker will perform OTJ search after any outcome of the bargaining process and Nash bargaining can be used again. Using the above definitions, the condition $c_i^J < \bar{w}_i^{NS}$ can be rearranged, yielding

$$a(\theta)(W_1 - W_i - J_i) > v. \quad (3.70)$$

This means exactly that the joint benefit from OTJ search is higher than the cost.

3. Non-convex bargaining set ($\bar{w}_i^{NS} \leq c_i^J$): it would be jointly optimal not to perform OTJ search ($a(\theta)(W_1 - W_i - J_i) \leq v$) but for low values of the wage it is individually rational for the worker to do so ($\bar{w}_i^R \leq \bar{w}_i^{NS}$). Therefore, the bargaining set becomes non-convex. In this case, I apply the bargaining process from above to determine the wage. In the notation of section 3.4.2, $S \equiv S_i^{NS}$, $T \equiv S_i^{NS} - S_i^S$, and $x^{NS} \equiv W_i(\bar{w}_i^{NS}) - U = W_1 - U - \frac{v}{a(\theta)}$. Applying Lemma 3 and choosing the equilibrium as discussed above, the worker's surplus is

$$\begin{aligned} W_i - U &= \begin{cases} \beta S_i^{NS} & \text{if } \beta S_i^{NS} \geq W_1 - U - \frac{v}{a(\theta)} \\ W_1 - U - \frac{v}{a(\theta)} & \text{else} \end{cases} \\ &= \max\left(\beta S_i^{NS}, W_1 - U - \frac{v}{a(\theta)}\right). \end{aligned} \quad (3.71)$$

Having determined the individual behaviour in the bargaining process, I can find the outcome in equilibrium, where $W_i(\bar{w}_i) = W_i$ and $J_i(\bar{w}_i) = J_i$ hold. Then the value equations for given \bar{w}_i and e_i are the same as in section 3.3.2 regardless of the level of commitment. In particular, also the surplus is given by equations (3.39) and (3.40).

Appendix 3.D shows that the same threshold for OTJ search as a function of market tightness as in the social planner's problem is obtained. The firing threshold is also efficient conditional on market tightness. Depending on the level of productivity, the worker's surplus and optimal decision are summarised in the following table.

Decision	$W_i - U$	Range of productivity
No OTJ search	βS_i^{NS}	$x \in \left[1 - \frac{r+\lambda}{a(\theta)\beta p} \nu, 1\right]$
No OTJ search	$\beta S_1^{NS} - \frac{v}{a(\theta)}$	$x \in \left[1 - \frac{r+\lambda}{a(\theta)p} (\nu + pc\theta), 1 - \frac{r+\lambda}{a(\theta)\beta p} \nu\right]$
OTJ search	βS_i^S	$x \in \left[1 - \frac{r+\lambda}{1-\beta} \frac{c}{q(\theta)} - \left(\frac{\beta}{1-\beta} c\theta - \frac{v}{p}\right), 1 - \frac{r+\lambda}{a(\theta)p} (\nu + pc\theta)\right]$
Job destruction	0	$x \in \left[0, 1 - \frac{r+\lambda}{1-\beta} \frac{c}{q(\theta)} - \left(\frac{\beta}{1-\beta} c\theta - \frac{v}{p}\right)\right]$

Table 3.1: Optimal decision and worker surplus with (partial) commitment

Given the worker's surplus at all levels of productivity, her wage can be calculated from equation (3.57):

$$\bar{w}_1 = b + (r + \lambda + a(\theta)) \frac{\beta}{1 - \beta} \frac{pc}{q(\theta)} + \lambda \sum_{j=1}^n g(x_j) \max\{W_j - U, 0\}. \quad (3.72)$$

As the surplus is not split in the same proportion at all levels any more, the expression for \bar{w}_1 becomes more complicated. The wage differential, however, can still be easily obtained:

$$\bar{w}_1 - \bar{w}_i = (r + \lambda)(W_1 - U) - (r + \lambda + e_i a(\theta))(W_i - U) + \frac{\beta}{1 - \beta} e_i pc\theta - e_i v \quad (3.73)$$

$$= \begin{cases} \beta p(1 - x_i) & \text{No OTJ search} \\ (r + \lambda) \frac{v}{a(\theta)} & \text{Corner Solution} \\ \beta p(1 - x_i) + \beta pc\theta - (1 - \beta)v & \text{OTJ search} \end{cases} \quad (3.74)$$

The wage differential is the same as before in the first and the third case, as the surplus is still shared according to the Nash rule. In the intermediate case, when the worker is just made indifferent between searching and not, the differential reflects exactly this indifference: her expected cost of finding a job that pays \bar{w}_1 is $\frac{v}{a(\theta)}$ whereas the expected loss from the lower wage is $\frac{\bar{w}_1 - \bar{w}_i}{r + \lambda}$.

The main result of this section is summarised in the following proposition.

Proposition 4 *The efficient threshold for on-the-job search conditional on market tightness can be obtained at any positive level of wage commitment.*

This result is a consequence of the non-convex bargaining set that is obtained whenever OTJ search is not efficient. Of course, it crucially depends on the choice of the equilibrium in the bargaining game: even if the worker's bargaining power and the possible degree of inefficiency is small, the equilibrium at the kink of the bargaining set was chosen. On the one hand, this is a strong assumption when the worker and the firm enter such a bargaining game after a shock. On the other hand, it maximises the expected surplus when the job is created. The firm is thus compensated by having to pay a lower wage before such a shock, in anticipation of possible future shocks. The chosen equilibrium implied that the worker's share is higher than β in some cases to prevent her from OTJ search. If this implicit non-constant bargaining power was also used in the Nash bargaining in the no-commitment case, the efficient OTJ search decision could also be obtained.

3.5 Conclusion

In this chapter, I have discussed two possible inefficiencies that are present in a search model with OTJ search. First, opening vacancies imposes externalities on other firms and job seekers. It was shown that the worker's bargaining power as suggested by the Hosios rule is too low in the presence of OTJ search. Second, a worker that quits her firm imposes an externality on it. In general, this leads to too much job turnover. It was shown that allowing for wage commitment can help to reduce this externality. Using a suitable subgame perfect equilibrium the efficient OTJ search decision could be obtained.

As pointed out above, the choice of the respective equilibrium is crucial for this result. In future research, the consequences of different bargaining mechanisms may be analysed. A second limitation of this chapter is that it focuses only on the outcome in a steady state. Compared to Menzio and Shi (2011), it is more difficult to study aggregate shocks and transitional dynamics; for example, the distribution of matches over productivity influences the efficient level of labour market tightness. In the presence of aggregate shocks, this distribution will depend on the past history of shocks, because the thresholds for destruction and OTJ search become time-varying.

Appendices to Chapter 3

3.A Solution of social planner's problem in the full model

The costates are denoted by U and W_i , respectively, which yields the Hamiltonian in the current-value form:

$$\begin{aligned}
 H = & ub + p \sum_{i=1}^n x_i w_i - pc\theta \left(u + \sum_{i=1}^n e_i w_i \right) - v \sum_{i=1}^n e_i w_i \\
 & + U \left[\lambda \sum_{i=1}^n d_i g(x_i) \sum_{j=1}^n w_j - a(\theta) u \right] + \\
 & + W_1 \left[\lambda \left(g(x_1) \sum_{j=1}^n w_j - w_1 \right) + a(\theta) \left(u + \sum_{i=1}^n e_i w_i \right) \right] + \\
 & + \sum_{i=2}^n W_i \left[\lambda \left((1-d_i) g(x_i) \sum_{j=1}^n w_j - w_i \right) - e_i a(\theta) w_i \right]. \tag{3.75}
 \end{aligned}$$

By differentiating with respect to θ , the analogous first-order condition to equation (3.3b) is obtained:

$$a'(\theta) \left[u(W_1 - U) + \sum_{i=1}^n e_i w_i (W_1 - W_i) \right] = pc(u + e), \tag{3.76a}$$

$$(1 - \eta) \left[\frac{u}{u + e} (W_1 - U) + \sum_{i=1}^n \frac{e_i w_i}{u + e} (W_1 - W_i) \right] = \frac{pc}{q(\theta)}, \tag{3.76b}$$

where $e \equiv \sum_{i=1}^n e_i w_i$ denotes the mass of OTJ seekers. Again, the marginal cost of increasing θ must equal the marginal benefit from matching more unemployed and employed workers, where the weighted average is now over all job seekers.

Maximisation with respect to e_i gives the condition for OTJ search

$$a(\theta) (W_1 - W_i) \geq pc\theta + v. \tag{3.77}$$

OTJ search takes place at productivity x_i , if the expected gain from improving the productivity outweighs the cost of opening additional vacancies and of the OTJ search cost.

A match is destroyed, if the value of an unemployed agent becomes larger than the value of keeping the job:

$$U \geq W_i. \tag{3.78}$$

The differential equations for the costates become:

$$\dot{U} = rU - [b - pc\theta + a(\theta)(W_1 - U)], \quad (3.79a)$$

$$\dot{W}_1 = rW_1 - \left[p + \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_1 \right) \right], \quad (3.79b)$$

$$\begin{aligned} \dot{W}_i &= rW_i - [px_i + e_i [a(\theta)(W_1 - W_i) - pc\theta - v]] \\ &\quad - \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_i \right), \end{aligned} \quad (3.79c)$$

and in a steady state the analogous equations to (3.5a), (3.5b), and (3.5c) are obtained:

$$rU = b - pc\theta + a(\theta)(W_1 - U), \quad (3.80a)$$

$$rW_1 = p + \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_1 \right), \quad (3.80b)$$

$$\begin{aligned} rW_i &= px_i + e_i [a(\theta)(W_1 - W_i) - pc\theta - v] \\ &\quad + \lambda \left(\sum_{j=1}^n g(x_j) [d_j U + (1 - d_j) W_j] - W_i \right). \end{aligned} \quad (3.80c)$$

Subtracting the equations from each other, one can obtain an expression for the respective surplus:

$$\begin{aligned} (r + e_i a(\theta) + \delta)(W_i - U) &= px_i - e_i v - b - (1 - e_i)(a(\theta)(W_1 - U) - pc\theta) \\ &\quad + \lambda \sum_{j=1}^n (1 - d_j) g(x_j) (W_j - W_i), \end{aligned} \quad (3.81a)$$

$$(r + e_i a(\theta) + \lambda)(W_1 - W_i) = p(1 - x_i) + e_i (pc\theta + v), \quad (3.81b)$$

$$(r + a(\theta) + \delta)(W_1 - U) = p - b + pc\theta + \lambda \sum_{j=1}^n (1 - d_j) g(x_j) (W_j - W_1), \quad (3.81c)$$

where $\delta \equiv \lambda \sum_{i=1}^n d_i g(x_i)$ denotes the (endogenous) rate of job destruction. The second set of equations could be substituted into the last to obtain $W_1 - U$ depending only on the control variables.

Substituting equation (3.81b) into the OTJ search condition (3.77) yields

$$a(\theta) \frac{p(1 - x_i)}{r + \lambda} \geq pc\theta + v, \quad (3.82)$$

which is monotonous in x_i . Hence, for each θ there exists a threshold $\bar{S}(\theta)$, such that there is OTJ search if $x \leq \bar{S}(\theta)$:

$$\bar{S}(\theta) = 1 - \frac{(pc\theta + v)(r + \lambda)}{a(\theta)p}. \quad (3.83)$$

3.B Derivation of jointly optimal OTJ search decision with (partial) commitment

Under partial commitment, the joint value of a match is given by

$$(r + \lambda + \omega) (W_i(\bar{w}_i) + J_i(\bar{w}_i)) = e_i [a(\theta) (W_1 - W_i(\bar{w}_i) - J_i(\bar{w}_i)) - \nu] + c_i^W + c_i^J, \quad (3.84)$$

where

$$c_i^W + c_i^J = -p(1 - x_i) + (r + \lambda) (W_1 + J_1) + \omega (W_i + J_i). \quad (3.85)$$

It is piecewise constant and has a jump discontinuity at \bar{w}_i^{NS} : when the worker does not want to search, the joint surplus is

$$S_i^{NS} = \frac{-p(1 - x_i) + (r + \lambda) S_1 + \omega S_i}{r + \lambda + \omega}, \quad (3.86)$$

and when the worker wants to search it is

$$S_i^S = \frac{-p(1 - x_i) + (r + \lambda) S_1 + \omega S_i + a(\theta) (W_1 - U) - \nu}{r + \lambda + \omega + a(\theta)}. \quad (3.87)$$

The ranking of the two determines whether it is jointly optimal for the firm and the worker to perform OTJ search. They should jointly agree on OTJ search if

$$a(\theta) (W_1 - W_i(\bar{w}_i) - J_i(\bar{w}_i)) \geq \nu \quad (3.88)$$

holds.

3.C Proof of Lemma 3

Proof. Consider the following strategies: when they get the chance, the worker proposes x^W and the firm proposes x^F . The worker accepts any proposal $x \geq x^F$ and the firm accepts if $J(x) \geq J(x^W)$. I want to find x^W and x^F , such that these strategies constitute a subgame perfect equilibrium. The firm must be indifferent between accepting or not. This condition arises, because if the firm was strictly better off from accepting, the worker could marginally raise her share. This would make the worker better off and the firm at most only marginally worse off. In contrast, it can be the case that the worker is strictly better off from accepting the firm's offer. This difference is caused by the non-convexity: if x^W is larger than $x^{NS} - T$, it is best for the firm to offer at least x^{NS} and thereby prevent the worker from OTJ search. Therefore, the firm will not propose $x^F \in (x^{NS} - T, x^{NS})$. This yields the conditions

$$J(x^W) = \delta [\beta J(x^W) + (1 - \beta) J(x^F)], \quad (3.89a)$$

$$x^F \geq \delta [\beta x^W + (1 - \beta) x^F]. \quad (3.89b)$$

being equivalent to

$$(1 - \delta\beta)J(x^W) = \delta(1 - \beta)J(x^F), \quad (3.90a)$$

$$x^F \geq \frac{\delta\beta x^W}{1 - \delta(1 - \beta)}. \quad (3.90b)$$

The inequality can only be strict if $x^F = x^{NS}$.

First, I discuss the case of equality; it immediately follows that $x^W > x^F$ holds as there is an advantage of being allowed to make an offer. Furthermore, both offers must lie on the same side of the discontinuity at x^{NS} , if δ is sufficiently large: if the worker's offer involved no OTJ search but the firm's offer did, they cannot both be indifferent between accepting or not as the (discounted) surplus is bigger when the worker makes the offer. Substituting x^F into the firm's payoff function yields

$$(1 - \delta) \left(S - I(x^W < x^{NS})T \right) - (1 - \delta\beta)x^W = -\frac{\delta(1 - \beta)\delta\beta}{1 - \delta(1 - \beta)}x^W. \quad (3.91)$$

Rearranging it, gives the worker's surplus proposed by a worker and consequently by a firm:

$$x^W = [1 - \delta(1 - \beta)] \left(S - I(x^W < x^{NS})T \right), \quad (3.92a)$$

$$x^F = \delta\beta \left(S - I(x^W < x^{NS})T \right). \quad (3.92b)$$

As $\delta \rightarrow 1$, the worker's surplus converges to

$$x = \begin{cases} \beta S & \text{if } \beta \geq \frac{x^{NS}}{S} \\ \beta(S - T) & \text{if } \beta < \frac{x^{NS} - T}{S - T} \end{cases} \quad (3.93)$$

The bargained share is such that x^F indeed lies outside the interval $(x^{NS} - T, x^{NS})$ for δ sufficiently large. This makes both strategies a best response and a subgame perfect equilibrium is found.

Second, I discuss the alternative where $x^F = x^{NS}$ holds, so that the firm just induces the worker not to perform OTJ search. As the firm must be indifferent between accepting or not, the worker's offer cannot be smaller than x^{NS} as well. Using condition (3.90a) for the firm's indifference, yields the worker's offer

$$x^W = \frac{(1 - \delta)S + \delta(1 - \beta)x^{NS}}{1 - \delta\beta}. \quad (3.94)$$

As $\delta \rightarrow 1$, the worker's surplus converges to $x = x^{NS}$, which is the claimed surplus in equilibrium for $\beta \leq \frac{x^{NS}}{S}$. For it to be a best response by the worker $x^{NS} = x^F \geq \frac{\delta\beta x^W}{1 - \delta(1 - \beta)}$ has to hold. This gives the condition

$$x^{NS} \geq \frac{\delta\beta}{1 - \delta(1 - \beta)} \frac{(1 - \delta)S + \delta(1 - \beta)x^{NS}}{1 - \delta\beta}, \quad (3.95)$$

which is equivalent to

$$\beta \leq \frac{x^{NS}}{\delta S}. \quad (3.96)$$

Therefore, $x = x^{NS}$ can indeed be obtained as δ approaches 1 for all $\beta \leq \frac{x^{NS}}{S}$. ■

3.D Derivation of equilibrium with (partial) commitment

Workers in matches with high productivity do not perform OTJ search and the surplus is divided by Nash bargaining. The range of such productivities can be determined using equations (3.86) and (3.39) when $e_i = 0$ as well as that the Nash sharing rule holds for the maximum productivity:

$$\beta S_i^{NS} \geq W_1 - U - \frac{v}{a(\theta)}, \quad (3.97a)$$

$$\beta \frac{p(1-x_i)}{r+\lambda} \leq \frac{v}{a(\theta)}. \quad (3.97b)$$

This is the same condition as condition (3.43) in the case of no commitment that defined the lower bound $S_0(\theta) = 1 - \frac{r+\lambda}{a(\theta)\beta p}v$ for the range of productivities not featuring OTJ search in equilibrium.

If productivity is below this threshold, there is still no OTJ search but the wage is not determined by Nash bargaining but using the bargaining process for non-convex sets. The worker's value is then at the corner solution $W_i = W_1 - \frac{v}{a(\theta)}$ as long as $a(\theta)(W_1 - W_i - J_i) \leq v$ holds, which yields the condition

$$a(\theta)(S_1 - S_i - J_1) \leq v, \quad (3.98a)$$

$$a(\theta) \frac{p(1-x_i)}{r+\lambda} \leq v + pc\theta. \quad (3.98b)$$

This determines the OTJ search threshold $S(\theta) = 1 - \frac{r+\lambda}{a(\theta)p}(v + pc\theta)$, which is indeed the threshold obtained in the case of a social planner in equation (3.26).

Below this threshold, OTJ search will take place until the productivity becomes too small and the job is destroyed. The job is destroyed if $\bar{w}_i^R > c_i^J$, which using equations (3.85) and (3.61) yields

$$(r + \lambda + \omega)U - e_i[a(\theta)(W_1 - U) - v] > c_i^W + c_i^J, \quad (3.99a)$$

$$-e_i a(\theta) \left[S_1 - \frac{pc\theta + v}{a(\theta)} \right] > -p(1-x_i) + (r+\lambda)S_1 + \omega S_i. \quad (3.99b)$$

In an equilibrium with at least some OTJ search, $e_i = 1$ and $S_R = 0$ hold, which yields the condition

$$p(1-x_i) + pc\theta + v > (r + \lambda + a(\theta)) \frac{pc}{(1-\beta)q(\theta)}, \quad (3.100)$$

and hence the reservation productivity

$$R(\theta) = 1 - \frac{r+\lambda}{1-\beta} \frac{c}{q(\theta)} - \left(\frac{\beta}{1-\beta} c\theta - \frac{v}{p} \right). \quad (3.101)$$

It is exactly the threshold obtained in equation (3.51) in the case without commitment. It is not surprising that it does not depend on the level of commitment in equilibrium as the firm and the worker jointly decide to destroy the job, which only happens if the surplus is negative.

Using equations (3.86) and (3.87) as well as the zero-profit condition one obtains the surplus with and without OTJ search in equilibrium:

$$S_i^S = \frac{pc}{(1-\beta)q(\theta)} - \frac{p(1-x_i) + pc\theta + v}{r + \lambda + a(\theta)}, \quad (3.102)$$

$$S_i^{NS} = \frac{pc}{(1-\beta)q(\theta)} - \frac{p(1-x_i)}{r + \lambda}. \quad (3.103)$$

Bibliography

- ALAN, S., T. CROSSLEY, AND H. LOW (2012): "Saving on a Rainy Day, Borrowing for a Rainy Day," Unpublished manuscript, University of Cambridge.
- AMARAL, P., AND J. ICE (2014): "Reassessing the Effects of Extending Unemployment Insurance Benefits," *Federal Reserve Bank of Cleveland Economic Commentary*, 2014-23.
- BARATTIERI, A., S. BASU, AND P. GOTTSCHALK (2010): "Some Evidence on the Importance of Sticky Wages," NBER working paper 16130.
- BAYER, C., R. LÜTTICKE, L. PHAM-DAO, AND V. TJADEN (2014): "Precautionary Savings, Illiquid Assets and the Aggregate Consequences of Shocks to Household Income Risk," Unpublished Manuscript, University of Bonn.
- BENITO, A. (2004): "Does Job Insecurity affect Household Consumption," Bank of England Working Paper no. 220.
- BILS, M., Y. CHANG, AND S.-B. KIM (2011): "Worker Heterogeneity and Endogenous Separations in a Matching Model of Unemployment Fluctuations," *American Economic Journal: Macroeconomics*, 3(1), 128–154.
- BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.
- BLOOM, N., S. BOND, AND J. VAN REENEN (2007): "Uncertainty and Investment Dynamics," *The Review of Economic Studies*, 74(2), 391–415.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2014): "Really Uncertain Business Cycles," CES Working Paper 14-18.
- CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage Bargaining with On-the-Job Search: Theory and Evidence," *Econometrica*, 74(2), 323–364.
- CARCELES-POVEDA, E., AND D. COEN-PIRANI (2009): "Shareholders' Unanimity with Incomplete Markets," *International Economic Review*, 50(1), 577–606.
- CARD, D., R. CHETTY, AND A. WEBER (2007): "Cash-On-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market," *Quarterly Journal of Economics*, 122(4), 1511–1560.
- CARROLL, C. (1992): "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 2, 61–156.

- CHALLE, E., J. MATHERON, X. RAGOT, AND J. F. RUBIO-RAMIREZ (2014): "Precautionary Saving and Aggregate Demand," unpublished manuscript.
- CHALLE, E., AND X. RAGOT (2014): "Precautionary Saving over the Business Cycle," unpublished manuscript.
- COLES, M. G., AND A. M. KELISHOMI (2011): "New Business Start-ups and the Business Cycle," CEPR Discussion Paper DP8588.
- COOPER, R., J. HALTIWANGER, AND J. L. WILLIS (2007): "Search frictions: Matching aggregate and establishment observations," *Journal of Monetary Economics*, 54, 56–78.
- COOPER, R. W., J. C. HALTIWANGER, AND J. WILLIS (2004): "Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications," NBER working paper 10297.
- COSTAIN, J. S., AND M. REITER (2008): "Business cycles, unemployment insurance, and the calibration of matching models," *Journal of Economic Dynamics and Control*, 32(4), 1120–1155.
- DALY, M. C., AND HOBIJN (2013): "Downward Nominal Wage Rigidities Bend the Phillips Curve," Federal Reserve Bank of San Francisco Working Paper 2013-08.
- DALY, M. C., B. HOBIJN, AND B. LUCKING (2012): "Why Has Wage Growth Stayed Strong?," *Federal Reserve Bank of San Francisco Economic Letter*, 10.
- DAVIS, S. J., R. J. FABERMAN, AND J. HALTIWANGER (2012): "Labor market flows in the cross section and over time," *Journal of Monetary Economics*, 59(1), 1–18.
- DAVIS, S. J., AND J. HALTIWANGER (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *The Quarterly Journal of Economics*, 107(3), 819–863.
- DEN HAAN, W. J., G. RAMEY, AND J. WATSON (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(3), 482–498.
- DICKENS, W. T., L. GOETTE, E. L. GROSHEN, S. HOLDEN, J. MESSINA, M. E. SCHWEITZER, J. TURUNEN, AND M. E. WARD (2007): "How Wages Change: Micro Evidence from the International Wage Flexibility Project," *Journal of Economic Perspectives*, 21(2), 195–214.
- DRUANT, M., S. FABIANI, G. KEZDI, A. LAMO, F. MARTINS, AND R. SABBATINI (2009): "How are Firms' Wages and Prices Linked," ECB Working Paper 1084.
- EKERN, S., AND R. WILSON (1974): "On the Theory of the Firm in an Economy with Incomplete Markets," *Bell Journal of Economics and Management Science*, 5(1), 171–180.
- ELMENDORE, D. W., AND M. S. KIMBALL (2000): "Taxation of Labor Income and the Demand for Risky Assets," *International Economic Review*, 41(3), 801–832.
- ELSBY, M. W. L., AND R. MICHAELS (2013): "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows," *American Economic Journal: Macroeconomics*, 5(1), 1–48.
- ESSER, I., T. FERRARINI, K. NELSON, J. PALME, AND O. SJÖBERG (2013): "Unemployment Benefits in the EU Member States," European Commission report.

- FALLICK, B., AND C. A. FLEISCHMAN (2004): "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Finance and Economics Discussion Series 2004-34. Board of Governors of the Federal Reserve System (U.S.).
- FAVILUKIS, J., S. C. LUDVIGSON, AND S. VAN NIEUWERBURGH (2013): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," unpublished manuscript.
- FUJITA, S., AND M. NAJAJIMA (2014): "Worker Flows and Job Flows: A Quantitative Investigation," Federal Reserve Bank of Philadelphia Working Paper 13-9/R.
- FUJITA, S., AND G. RAMEY (2007): "Job matching and propagation," *Journal of Economic Dynamics and Control*, 31(11), 3671–3698.
- GAUTIER, P. A., C. N. TEULINGS, AND A. VAN VUUREN (2010): "On-the-Job Search, Mismatch and Efficiency," *The Review of Economic Studies*, 77(1), 245–272.
- GERTLER, M., C. HUCKFELDT, AND A. TRIGARI (2014): "Unemployment Fluctuations, Match Quality, and the Wage Cyclicity of New Hires," Slides available at "<https://files.nyu.edu/ckh249/public/papers.html>".
- GOLLIER, C., AND J. W. PRATT (1996): "Risk Vulnerability and the Tempering Effect of Background Risk," *Econometrica*, 64(5), 1109–1123.
- GORNEMAN, N., K. KUESTER, AND M. NAKAJIMA (2012): "Monetary Policy with Heterogeneous Agents," Unpublished manuscript, Federal Reserve Bank of Philadelphia.
- GRUBER, J. (2001): "The Wealth of the Unemployed," *Industrial and Labor Relations Review*, 55(1), 79–94.
- GUGLIELMINETTI, E. (2015): "The Effects of Uncertainty Shocks on the Labor Market: A Search Approach," Unpublished manuscript, Sciences Po.
- HAEFKE, C., M. SONNTAG, AND T. VAN RENS (2013): "Wage Rigidity and Job Creation," *Journal of Monetary Economics*, 60(8), 887–899.
- HAGEDORN, M., F. KARAHAN, I. MANOVSKII, AND K. MITMAN (2015): "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects," Unpublished manuscript.
- HALL, R. E. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), 50–65.
- HALL, R. E., AND P. R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98(4), 1653–1674.
- HOSIOS, A. J. (1990): "On the Efficiency of Matching and Related Models of Search and Unemployment," *The Review of Economic Studies*, 57(2), 279–298.
- IRELAND, P. (2004): "Money's Role in the Monetary Business Cycle," *Journal of Money Credit and Banking*, 36(6), 969–983.

- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105(3), 1177–1216.
- KAAS, L., AND P. KIRCHER (2014): "Efficient Firm Dynamics in a Frictional Labor Market," Unpublished manuscript, University of Konstanz, University of Edinburgh.
- KHAN, A., AND J. K. THOMAS (2013): "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," *Journal of Political Economy*, 121(6), 1055–1107.
- KIMBALL, M. S. (1990): "Precautionary Saving in the Small and in the Large," *Econometrica*, 58(1), 53–73.
- (1992): "Precautionary Motives for Holding Assets," in *The New Palgrave Dictionary of Money and Finance*, ed. by P. Newman, M. Milgate, and J. Eatwell, pp. 158–161. Stockton Press, New York.
- KOLSRUD, J., C. LANDAIS, P. NILSSON, AND J. SPINNEWIJN (2015): "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden," Unpublished manuscript, London School of Economics.
- KRUEGER, A. B., J. CRAMER, AND D. CHO (2014): "Are the Long-Term Unemployed on the Margins of the Labor Market?," *Brookings Papers on Economic Activity*, Spring, 229–280.
- KRUSELL, P., T. MUKOYAMA, AND A. SAHIN (2010): "Labour-Market Matching with Precautionary Savings and Aggregate Fluctuations," *Review of Economic Studies*, 77, 1477–1507.
- KRUSELL, P., AND A. A. SMITH, JR. (1997): "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics*, 1, 387–422.
- (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867–896.
- LALIVE, R. (2007): "Unemployment Benefits, Unemployment Duration, and Postunemployment Jobs: A Regression Discontinuity Approach," *American Economic Review*, 97(2), 108–112.
- LE BARBANCHON, T. (2012): "The Effect of the Potential Duration of Unemployment Benefits on Unemployment Exits to Work and Match Quality in France," Crest Document de Travail 2012-21.
- LEAHY, J. V., AND T. M. WHITED (1996): "The Effect of Uncertainty on Investment: Some Stylized Facts," *Journal of Money, Credit and Banking*, 28(1), 64–83.
- LUCAS, JR., R. E. (2000): "Inflation and Welfare," *Econometrica*, 68(2), 247–274.
- MCKAY, A., AND R. REIS (2013): "The Role of Automatic Stabilizers in the U.S. Business Cycle," Unpublished Manuscript, Boston University and Columbia University.
- MECIKOVSKY, A. M., AND M. MEIER (2015): "Do plants freeze upon uncertainty shocks?," Unpublished manuscript, University of Bonn.
- MENZIO, G., AND S. SHI (2010): "Block recursive equilibria for stochastic models of search on the job," *Journal of Economic Theory*, 145(4), 1453–1494.

- (2011): “Efficient Search on the Job and the Business Cycle,” *Journal of Political Economy*, 119(3), 468–510.
- MICHAILLAT, P. (2014): “A Theory of Countercyclical Government Multiplier,” *American Economic Journal: Macroeconomics*, 6(1), 190–217.
- MOEN, E. R., AND A. ROSEN (2011): “Incentives in Competitive Search Equilibrium,” *The Review of Economic Studies*, 78(2), 733–761.
- MORTENSEN, D. T., AND C. A. PISSARIDES (1994): “Job Creation and Job Destruction in the Theory of Unemployment,” *The Review of Economic Studies*, 61(3), 397–415.
- NAKAJIMA, M. (2010): “Business Cycles in the Equilibrium Model of Labor Market Search and Self-Insurance,” Federal Reserve Bank of Philadelphia Working Paper 2010-24.
- NEKOEI, A., AND A. WEBER (2015): “Does Extending Unemployment Benefits Improve Job Quality,” CEPR Discussion Paper no. 10568.
- PETRONGOLO, B., AND C. A. PISSARIDES (2001): “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39(2), 390–431.
- PISSARIDES, C. A. (1994): “Search Unemployment with On-the-job Search,” *The Review of Economic Studies*, 61(3), 457–475.
- (2000): *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, 2nd edn.
- (2009): “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?,” *Econometrica*, 77(5), 1339–1369.
- POSTEL-VINAY, F., AND J.-M. ROBIN (2002): “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity,” *Econometrica*, 70(6), 2295–2350.
- RAVN, M. O., AND V. STERK (2015): “Job Uncertainty and Deep Recessions,” unpublished manuscript.
- REITER, M. (2009): “Solving Heterogenous-Agent Models by Projection and Perturbation,” *Journal of Economic Dynamics and Control*, 33(3), 649–665.
- RENDAHL, P. (2012): “Fiscal Policy in an Unemployment Crisis,” Unpublished manuscript, University of Cambridge.
- RUBINSTEIN, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50(1), 97–109.
- RUDANKO, L. (2009): “Labor Market Dynamics under Long-Term Wage Contracting,” *Journal of Monetary Economics*, 56, 170–183.
- SAPORTA-EKSTEN, I. (2014): “Job Loss, Consumption and Unemployment Insurance,” Unpublished manuscript, Stanford University.
- SCHAAL, E. (2012): “Uncertainty, Productivity and Unemployment in the Great Recession,” Unpublished manuscript, Federal Reserve Bank of Minneapolis.

- SCHMIEDER, J., T. VON WACHTER, AND S. BENDER (2014): "The Causal Effect of Unemployment Duration on Wages: Evidence from Unemployment Insurance Extensions," unpublished manuscript.
- SHAO, E., AND P. SILOS (2007): "Uninsurable Individual Risk and the Cyclical Behavior of Unemployment and Vacancies," Federal Reserve Bank of Atlanta working paper 2007-5.
- SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25–49.
- (2006): "On-the-job search and strategic bargaining," *European Economic Review*, 50(4), 811–830.
- (2012): "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 15(2), 127–148.
- STEPHENS, JR., M. (2004): "Job Loss Expectations, Realizations, and Household Consumption Behavior," *The Review of Economics and Statistics*, 86(1), 253–269.
- STEVENS, M. (2004): "Wage-Tenure Contracts in a Frictional Labour Market: Firms' Strategies for Recruitment and Retention," *The Review of Economic Studies*, 71(2), 535–551.
- STOLE, L. A., AND J. ZWIEBEL (1996): "Intra-firm Bargaining under Non-binding Contracts," *The Review of Economic Studies*, 63(3), 375–410.
- TELYUKOVA, I. A. (2013): "Household Need for Liquidity and the Credit Card Debt Puzzle," *Review of Economic Studies*, 80, 1148–1177.
- VAN OURS, J., AND G. RIDDER (1992): "Vacancies and the Recruitment of New Employees," *Journal of Labor Economics*, 10(2), 138–155.
- VAN OURS, J. C., AND M. VODOPIVEC (2008): "Does Reducing Unemployment Insurance Generosity Reduce Job Match Quality," *Journal of Public Economics*, 92, 684–695.
- YOUNG, E. R. (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-Stochastic Simulations," *Journal of Economic Dynamics and Control*, 34(1), 36–41.