

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

[JOB MARKET PAPER]

Markus Riegler[†]

London School of Economics

21 November 2014

Abstract

Increases in uncertainty lead to increases in the unemployment rate. Using US data, I show empirically that this is 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 (job creation and job destruction) do not necessarily coincide with worker flows (hires and separations). Costly job creation (in addition to the usual hiring cost) 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.

JEL classification: C63, E24, E32, J63, J64.

Keywords: uncertainty, multi-worker firms, idiosyncratic risk, job-finding rate, separation rate.

*I am very grateful to Wouter den Haan and Philipp Kircher for their supervision during my graduate studies. I thank Monique Ebell, Mike Elsby, Per Krusell, Albert Marcet, Lukas Mayr, Pascal Michaillat, Rachel Ngai, Pedro Pinto, Ricardo Reis, Edouard Schaal, Tony Smith, and seminar participants at LSE, the EDP Jamboree, and the RES conference for helpful comments. I acknowledge financial support by the Economic and Social Research Council.

[†]Department of Economics and Centre for Macroeconomics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK, m.riegler@lse.ac.uk, www.mriegler.com

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 paper, higher uncertainty means a higher expected and higher realized dispersion of idiosyncratic productivity across firms. Bloom et al. (2012) 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

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.

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 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 et al. (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 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 et al. (2007), Elsby and Michaels (2013), and Fujita and Nakajima (2014), whereas Schaal (2012)

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.

and Kaas and Kircher (2013) 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 paper 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 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 paper 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 3 describes the model and derives the optimal choice of firms. Section 4 describes the calibration that is used. In section 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 6. Section 7 describes the algorithm

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.

that I develop to solve the model. Section 8 concludes.

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.

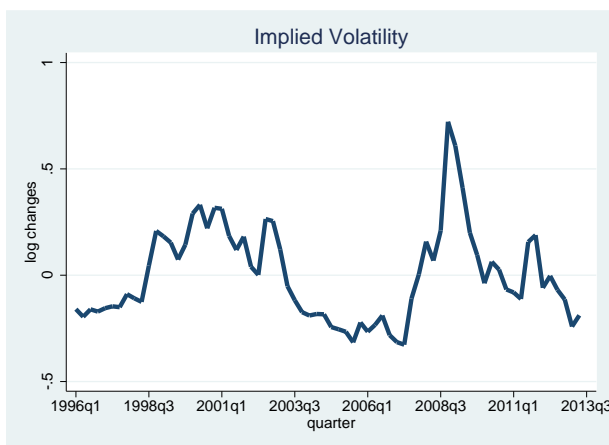


Figure 1: Index of implied volatility constructed from US stock options.

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 one period lag.⁶

Figure 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 sep-

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 et al. (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.

⁶Appendix 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.

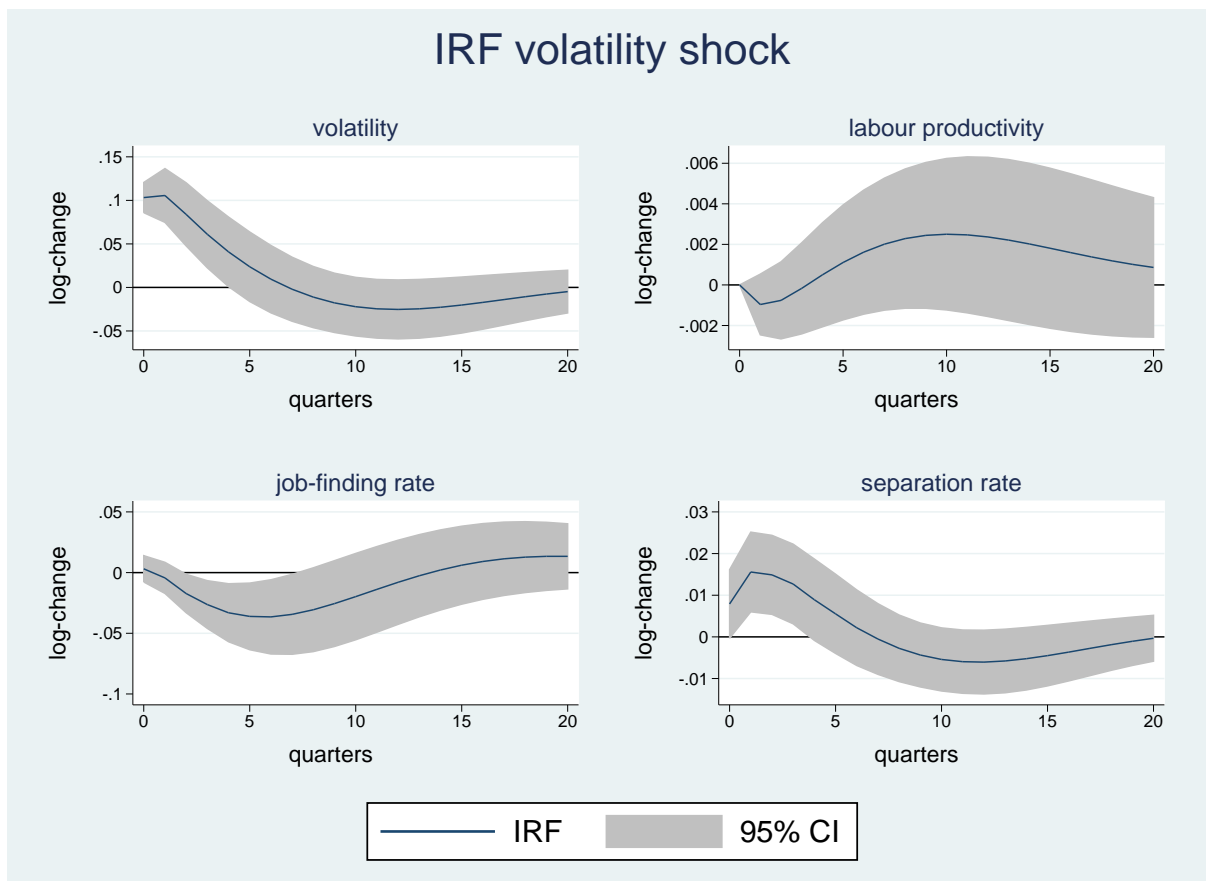


Figure 2: Impulse response functions after a volatility shock in US data.

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 A.2, the following approximation holds:

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

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 2 this means that the peak response of the unemployment rate is slightly less than 40% times the increase in uncertainty.

Table 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.

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 (2013) and Mecikovsky

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

and Meier (2014) use different measures for uncertainty and find similar effects. Guglielminetti (2013) 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 (2014) use the measure of macroeconomic uncertainty in the US proposed by Jurado et al. (2014). They find that job creation falls, whereas job destruction increases in response to higher uncertainty.

3 Model

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

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 the

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 Nakajima (2014) argue that cyclical fluctuations of job flows are mostly accounted for by the expansion or contraction of existing establishments.

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 et al. (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 (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.

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 paper 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 22 in Appendix B.1 shows the results of the comparative statics exercise when a destruction cost is used.

¹⁰In an earlier version of this paper, 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 et al. (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.

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 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

$$V(n_{-1}, x, s) = \max_n zx F(n) - W(n, x, s) - (n - (1 - \lambda) n_{-1})_+ \frac{c}{q(\theta(s))} - (n - n_{-1})_+ H + \dots + \beta \mathbb{E} [V((1 - \delta) n, x', s') | x, s], \quad (2)$$

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

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.

positions n^* and the marginal value of a position in each case.

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

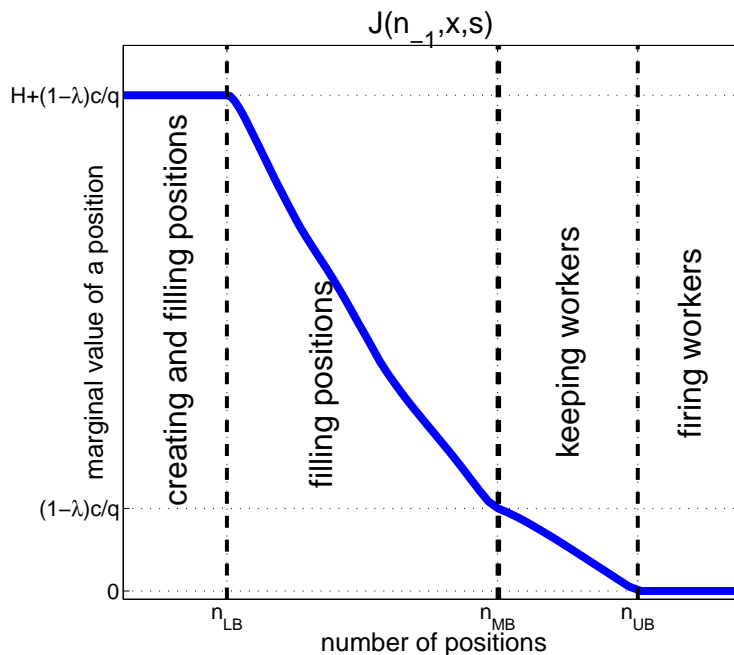


Figure 3: Marginal value of a position (including $1 - \lambda$ workers) for a given level of idiosyncratic productivity x and aggregate state s .

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

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E} [J((1 - \delta)n^*, x', s') | x, s] = 0, \quad (3)$$

$$J(n_{-1}, x, s) = 0. \quad (4)$$

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:

$$n^* = (1 - \lambda)n_{-1}, \quad (5)$$

$$J(n_{-1}, x, s) = (1 - \lambda) \left\{ \frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E} [J((1 - \delta)n^*, x', s') | x, s] \right\} \quad (6)$$

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

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:

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E} [J((1 - \delta)n^*, x', s') | x, s] = \frac{c}{q(\theta(s))}, \quad (7)$$

$$J(n_{-1}, x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (8)$$

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:

$$\begin{aligned} n^* &= n_{-1}, & (9) \\ J(n_{-1}, x, s) &= \frac{\partial \Pi}{\partial n}(n^*, x, s) - \lambda \frac{c}{q(\theta(s))} + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] \\ &\in \left((1 - \lambda) \frac{c}{q(\theta(s))}, H + (1 - \lambda) \frac{c}{q(\theta(s))} \right). \end{aligned} \quad (10)$$

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:

$$\frac{\partial \Pi}{\partial n}(n^*, x, s) + \beta(1 - \delta) \mathbb{E}[J((1 - \delta)n^*, x', s') | x, s] = H + \frac{c}{q(\theta(s))}, \quad (11)$$

$$J(n_{-1}, x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (12)$$

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

$$J(n_{-1}, x, s) = \min \left\{ H + \frac{(1 - \lambda)c}{q(\theta(s))}, \max \left\{ J_F(n_{-1}, x, s), \min \left\{ \frac{(1 - \lambda)c}{q(\theta(s))}, \max \{ J_K(n_{-1}, x, s), 0 \} \right\} \right\} \right\}. \quad (13)$$

where J_K and J_F denote the right hand side values of equations (6) and (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.

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:

$$J_F(n_{LB}(x, s), x, s) = H + (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (14)$$

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 $n_{MB}(x, s)$:

$$J_F(n_{MB}(x, s), x, s) = (1 - \lambda) \frac{c}{q(\theta(s))}. \quad (15)$$

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

$$J_K(n_{UB}(x, s), x, s) = 0. \quad (16)$$

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 2 summarises the optimal behaviour of a firm and Figure 4 illustrates the optimal choice of positions n^* given positions at the beginning of the period n_{-1} .¹⁴

The cutoffs are increasing in x as higher idiosyncratic productivity makes the marginal position more valuable in a given aggregate state. Figure 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 $(1 - \delta) n^*$.

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 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} .

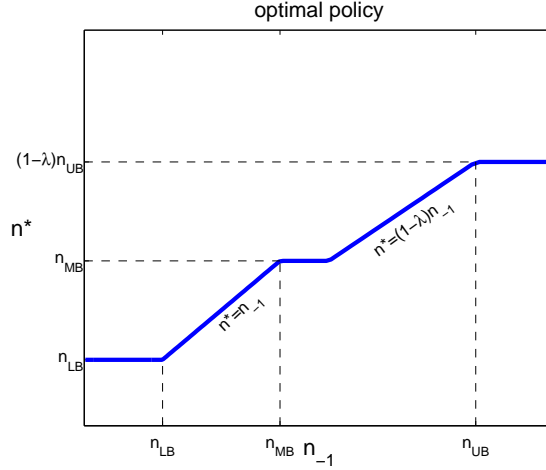


Figure 4: Optimal choice of positions for a given level of idiosyncratic productivity.

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

$$\int_{n_{i,-1}} \int_{x_i} (n^*(n_{i,-1}, x_i, s) - (1 - \lambda) n_{i,-1})_+ dF_s(n_{i,-1}, x_i) = M(u(s), v(s)) \quad (17)$$

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 7, I describe the algorithm that I develop to ensure market clearing exactly.¹⁵

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

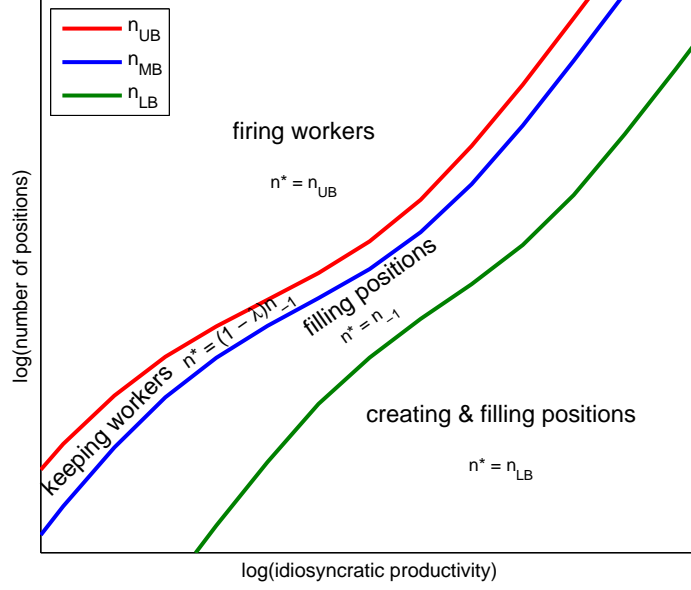


Figure 5: Cutoffs $n_{LB}(x)$, $n_{MB}(x)$, and $n_{UB}(x)$ for a given aggregate state.

4 Parametrization

The model is parametrized to US data. In section 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 3. There are three parameters that I recalibrate in order to meet the following three targets across calibrations: in the absence 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 4.

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), based on estimates in Petrongolo and Pissarides (2001). 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

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

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

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 (Stole-Zwiebel bargaining with bargaining power 0.5)
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)
<i>mass</i>	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 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 4.

be equal to 0.72 in the absence of aggregate 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

$$W(n, x, s) = (1 - \omega_1) z x n^\alpha + (\omega_2 \theta + \omega_3) n. \quad (18)$$

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.

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 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 $std\left(\frac{n_{i,t} - n_{i,t-52}}{0.5(n_{i,t} + n_{i,t-52})}\right)$.

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 parameters ω_1 and ω_2 are 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 resulting values are $\omega_1 = 0.8$ and $\omega_2 = 0.5c$.²¹ 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

²⁰This form can be derived by imposing Stole-Zwiebel 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 et al. (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.

²¹It is not crucial that the wage per worker differs across firms, i.e. that ω_1 is less than 1. A firm's production less wages is equal to $\omega_1 z x n^\alpha - (\omega_2 \theta + \omega_3) n$. 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 can be 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 27 in Appendix B.1 shows that the main results are robust to choosing $\omega_1 = 1$.

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:

$$\log x_t^{new} \sim N(0, \sigma_{x,t}^2). \quad (19)$$

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 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 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:

$$\log \sigma_{x,t} - \log \bar{\sigma}_x = \rho_\sigma (\log \sigma_{x,t-1} - \log \bar{\sigma}_x) + \varepsilon_{\sigma,t}. \quad (20)$$

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 of productivity

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

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

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

$$\int_{n_{i,-1}} \int_{x_i} (\tilde{z}_t x_{i,t})^{\frac{1}{1-\alpha}} dF_s(n_{i,-1}, x_i) = \text{const.} \quad (21)$$

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

$$\log z_t = \rho_z \log z_{t-1} + \varepsilon_{z,t}. \quad (22)$$

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 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 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 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 paper, 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

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 paper with the capital adjustment cost literature is to calculate the aggregate amount of capital adjustment costs spent in the economy in equilibrium. Cooper et al. (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 $H = 5$.

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. 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 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.

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 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

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 23 in Appendix 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 $\frac{\Delta u^{JFR}}{\Delta u^{JFR} + \Delta u^{SR}}$.

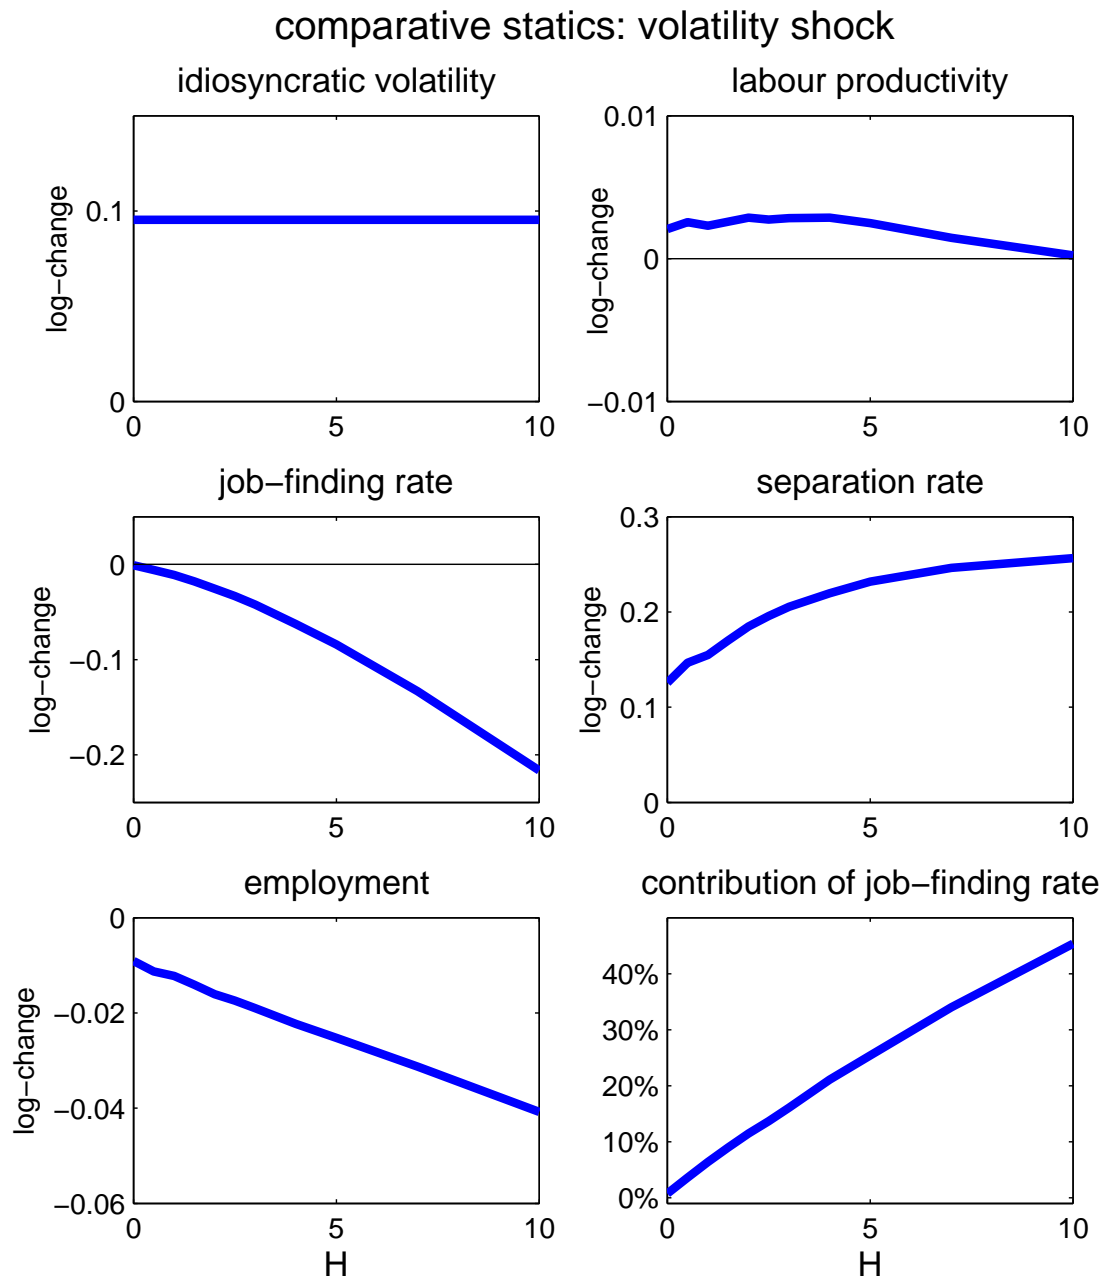


Figure 6: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The job-finding rate only falls and thus contributes to lower employment in the presence of creation cost $H > 0$.

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 11.

If we want to explain the effect of uncertainty on the creation of positions, we need to 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 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 $H + (1 - \lambda) \frac{\varepsilon}{q}$, 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 of the presence of

³¹This effect on measured TFP has also been discussed by Bloom et al. (2012).

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.

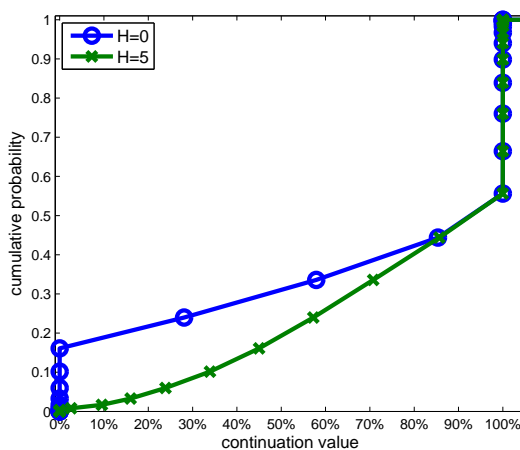


Figure 7: Cumulative distribution function of continuation value of a firm with median productivity at the creation cutoff when a shock arrives.

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.

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.

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 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 .

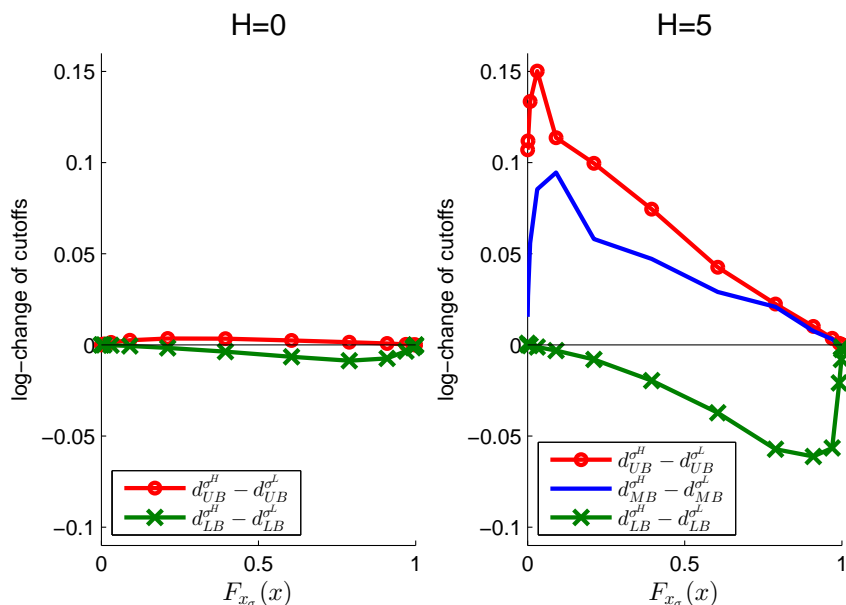


Figure 8: Wait-and-see effect: change in cutoffs after a 10% increase in uncertainty.

Figure 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 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 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:

$$\begin{aligned} \log n_j^{W\&S}(F_{\sigma_x}^{-1}(p)) &= \log n_j(F_{\sigma_x}^{-1}(p)) + \left[d_j^{\sigma^H}(p) - d_j^{\sigma^L}(p) \right], \\ j &\in \{LB, MB, UB\}. \end{aligned}$$

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 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 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 aggre-

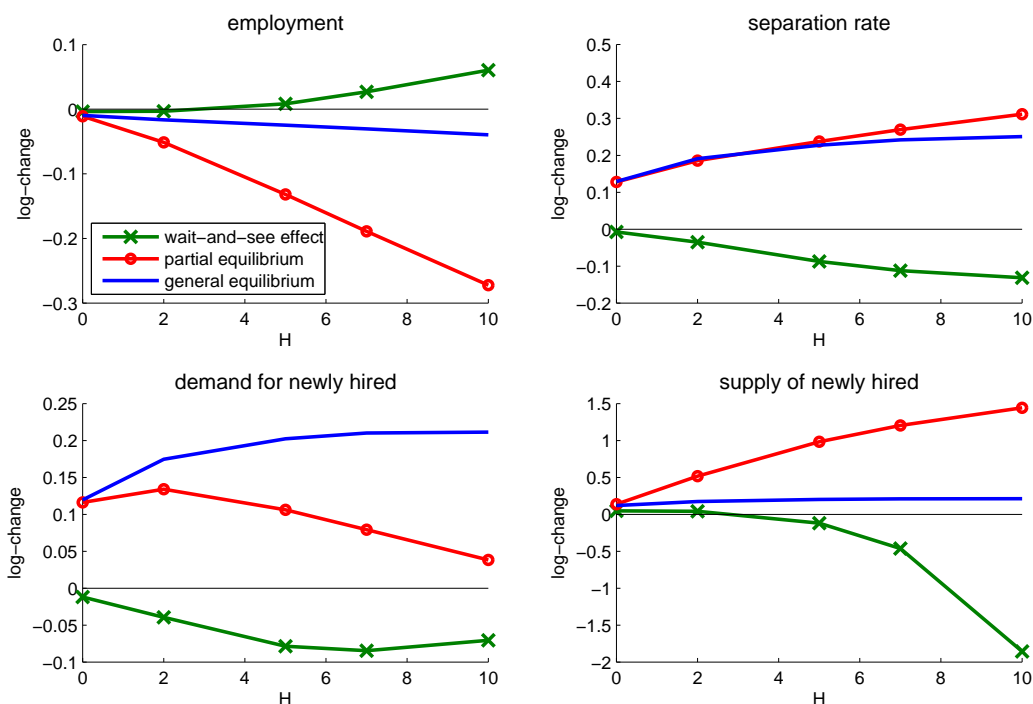


Figure 9: Decomposition of a 10% increase in uncertainty into wait-and-see effect (green curves with crosses), realized volatility effect (difference between red curves with dots and green curves with crosses), and general equilibrium effect (difference between solid blue curves and red curves with dots).

gate 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 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.

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 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 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.

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 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

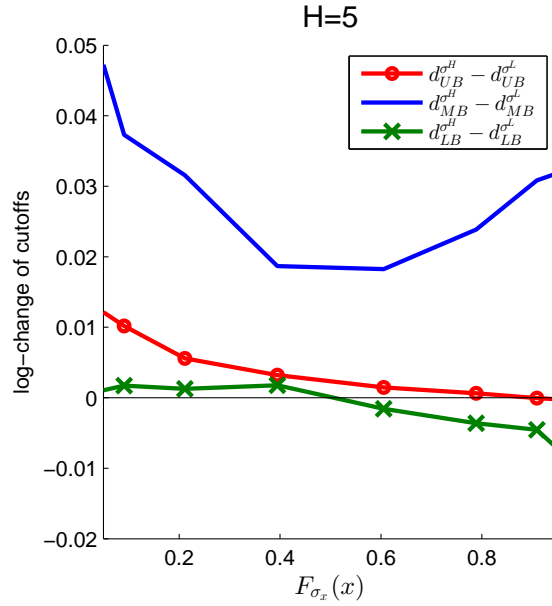


Figure 10: General equilibrium effect: change in cutoffs due to a change in market tightness that is necessary to clear markets after a 10% increase in uncertainty.

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 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 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

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:

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

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 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 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 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 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.

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 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.

IRFs volatility shock

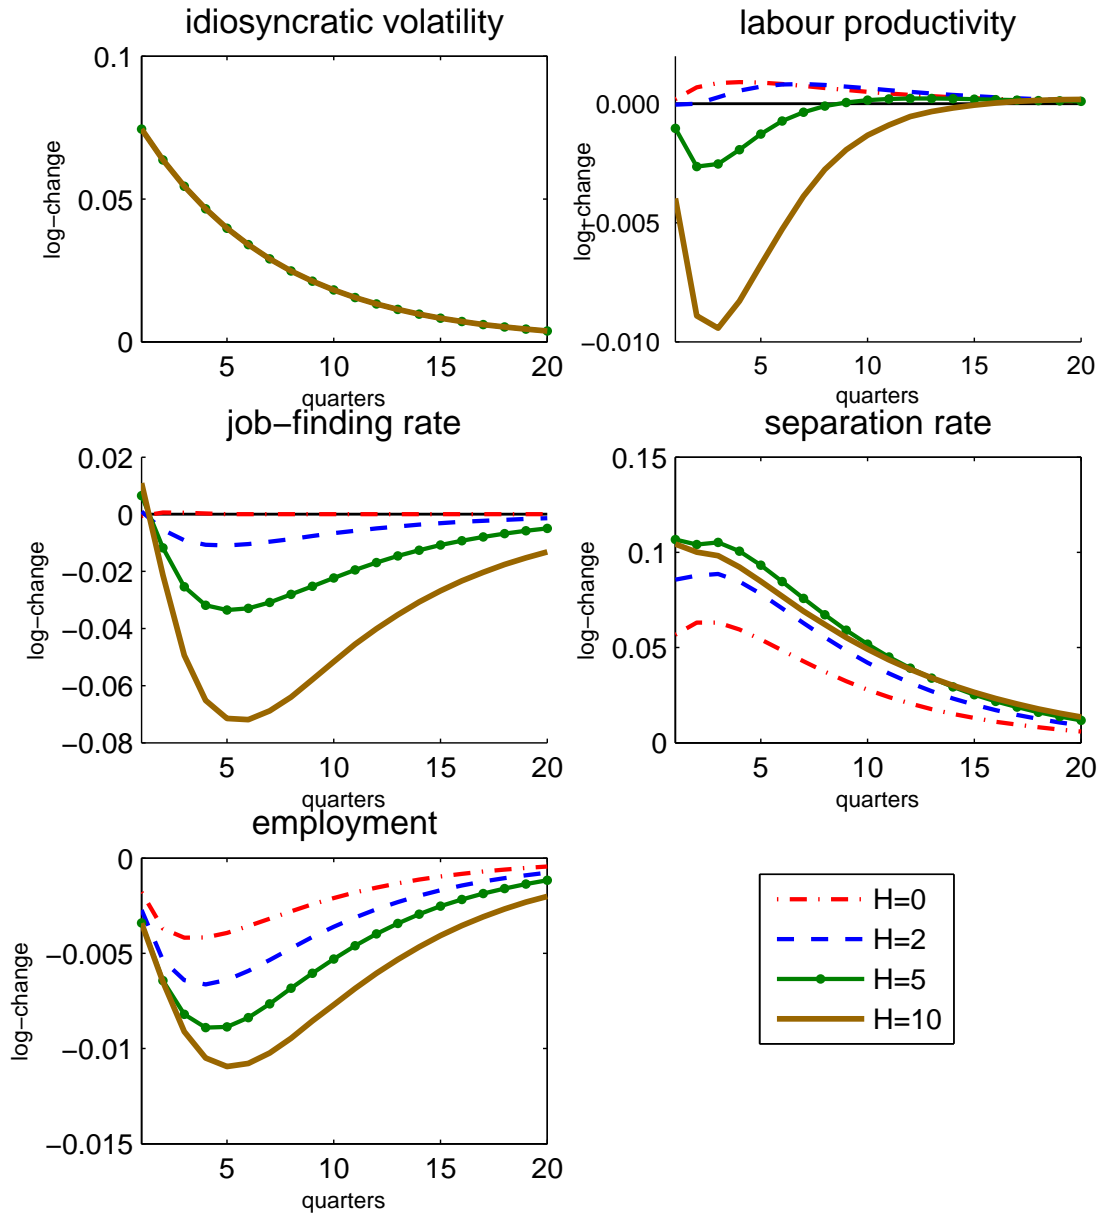


Figure 11: Impulse response functions for idiosyncratic volatility shock in the model.

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 12.

Whereas Figure 11 plots the median impulse response functions in the model, the observed 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

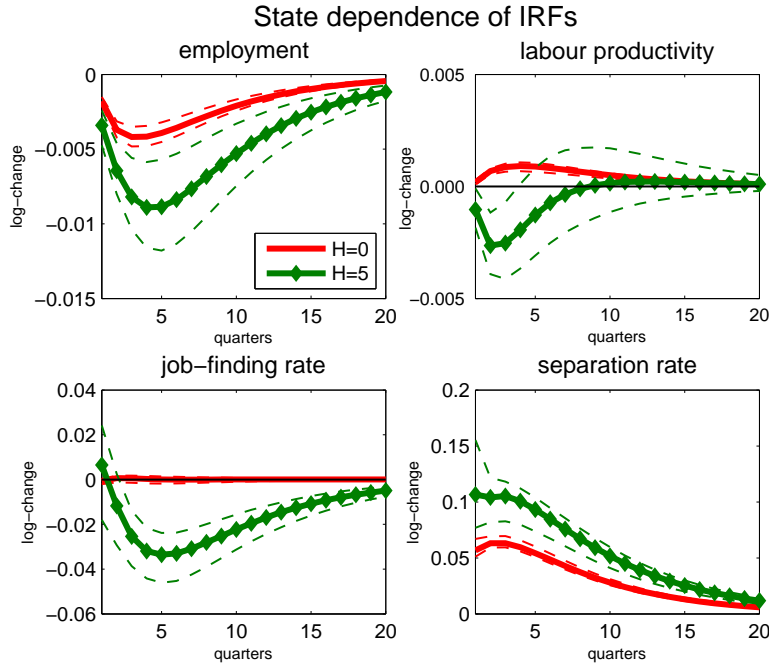


Figure 12: Impulse response functions for idiosyncratic volatility shock in the model. Solid lines depict the median change of the respective variable. Dashed lines draw the 5th and the 95th percentiles of the responses.

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 12 repeats the median IRFs from Figure 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	$d \log n$		
	n^L	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 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 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 unemploy-

ment 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 (44) in Appendix 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.

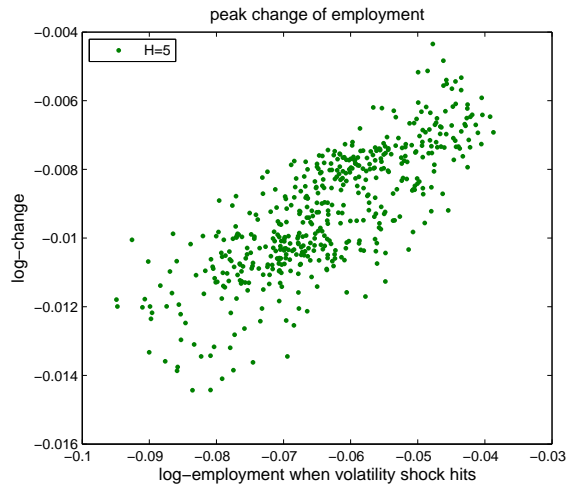


Figure 13: Each dot corresponds to the peak response of log-employment in one impulse response function after a positive uncertainty shock for $H = 5$.

But even if the initial unemployment rate is taken into account, some state dependency remains. Figure 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 13 also gives an indication that the first moment of the distribution is not enough to forecast aggregate outcomes. Section 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.

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 14. At a first glance, the unemployment 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 32 in Appendix 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.

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 volatility, shown in Figure 1.³⁷ Whereas the idiosyncratic volatility shock is feeded 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 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 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

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

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.

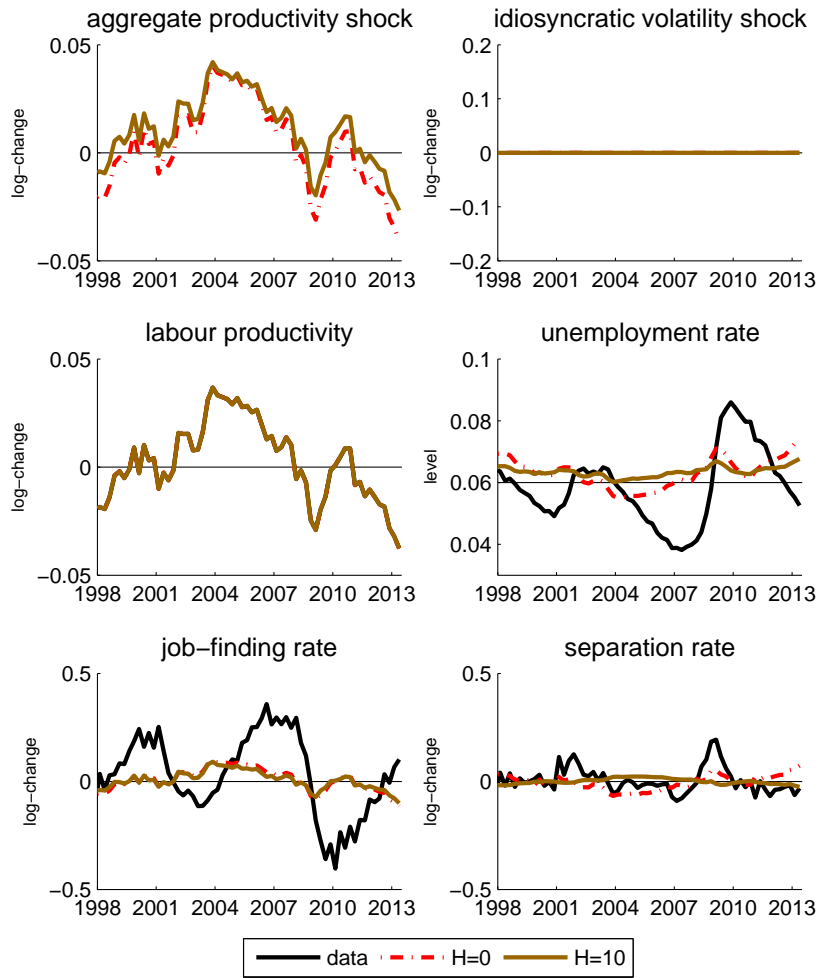


Figure 14: Model without uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data (linear trend removed).

is possible to back out the idiosyncratic volatility shocks that would be needed to match the unemployment rate. The resulting shocks are in Figure 33 in Appendix 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 $H = 10$.

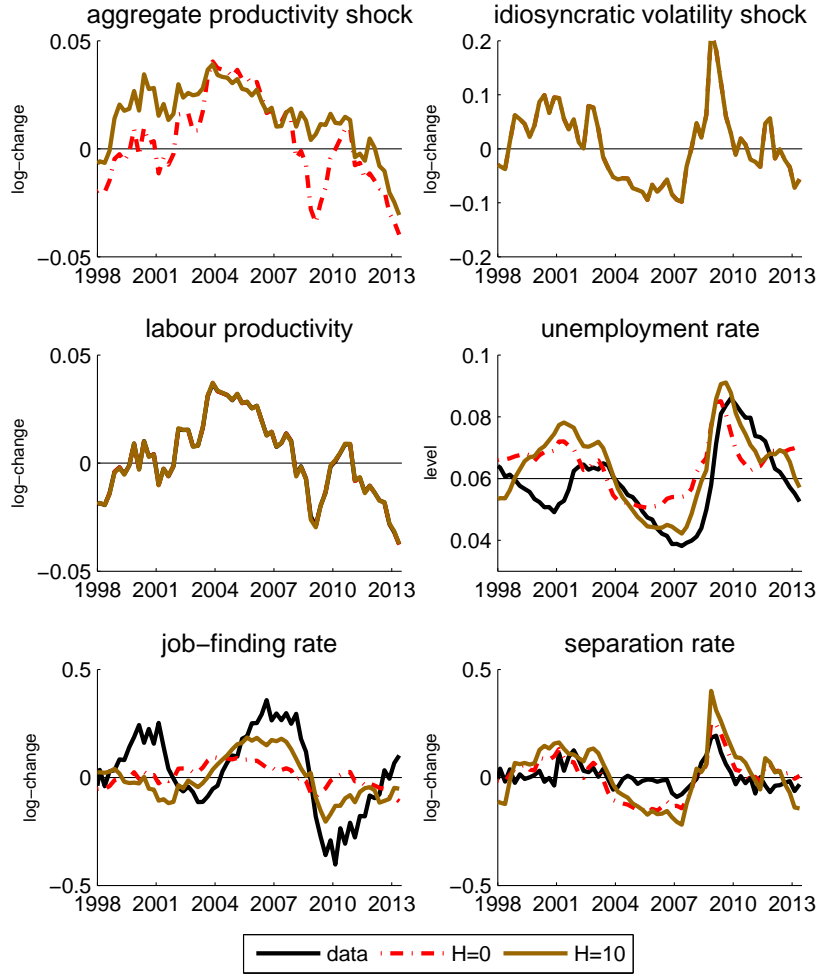


Figure 15: Model with aggregate productivity and uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data. Log-changes in idiosyncratic volatility are 20% of changes of the implied volatility index estimated in Section 2. Linear trends are removed from US data series.

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.³⁸

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:

$$m' = \Psi_m(m, z). \quad (23)$$

In addition, it might be necessary to approximate a set of prices p as a function of the state variables:^{39,40}

$$p = \Psi_p(m, z). \quad (24)$$

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

$$\tilde{m}' = \Psi_m(\tilde{m}, \tilde{z}) + x_m, \quad (25)$$

$$\tilde{p} = \Psi_p(\tilde{m}, \tilde{z}) + x_p. \quad (26)$$

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

$$m' = \Psi_m(m, z) + x_m, \quad (27)$$

$$p = \Psi_p(m, z) + x_p, \quad (28)$$

$$x' = \Psi_x(m, z, x). \quad (29)$$

³⁸Riegler (2014) describes its application to simpler heterogeneous agents models and compares it with other methods of solving these models in the presence of aggregate uncertainty.

³⁹For 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.

⁴⁰Note 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 depend on the current period's state variables.

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 (28). If it is added to intertemporal law of motions like in equation (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.

7.2 Application of the residual method in this paper

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 (17) must hold:

$$\int_{n_{i,-1}} \int_{x_i} (n^*(n_{i,-1}, x_i, s) - (1 - \lambda) n_{i,-1})_+ dF_s(n_{i,-1}, x_i) = a(\theta(s)) u(s). \quad (30)$$

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:

$$\log \theta = \Psi_{\theta,0} + \Psi_{\theta,1} \log z + \Psi_{\theta,2} \log N_{-1} + \Psi_{\theta,3} x_\sigma + \Psi_{\theta,4} m^x + x_\theta, \quad (31)$$

$$x_\theta = \Psi_{x_\theta,0} + \Psi_{x_\theta,1} \log z_{-1} + \Psi_{x_\theta,2} \log N_{-2} + \Psi_{x_\theta,3} x_{\sigma-1} + \Psi_{x_\theta,4} m_{-1}^x + \Psi_{x_\theta,5} x_{\theta-1}, \quad (32)$$

$$\log N = \Psi_{N,0} + \Psi_{N,1} \log z_{-1} + \Psi_{N,2} \log N_{-1} + \Psi_{N,3} x_{\sigma-1} + \Psi_{N,4} m_{-1}^x + \Psi_{N,5} x_{\theta-1}. \quad (33)$$

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:

$$(n^*(n_{i,-1}, x_i, z, x_\sigma, m^x, N_{-1}, x_\theta) - (1 - \lambda) n_{i,-1})_+. \quad (34)$$

Then, I simulate the economy using the histogram method developed by Young (2010).⁴¹ In the simulation, I impose the market clearing condition (30) using the residual method. All arguments except x_θ are predetermined in (34). I can solve for the unknown x_θ such that the market clearing condition (30) holds. Using the simulated values, the coefficients in the laws of motion (31), (32), and (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 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 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.⁴²

⁴¹This 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.

⁴²The 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 (31). The predicted residual is given by equation (32) using lagged state variables. 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 (32), using the realized residual x_θ .

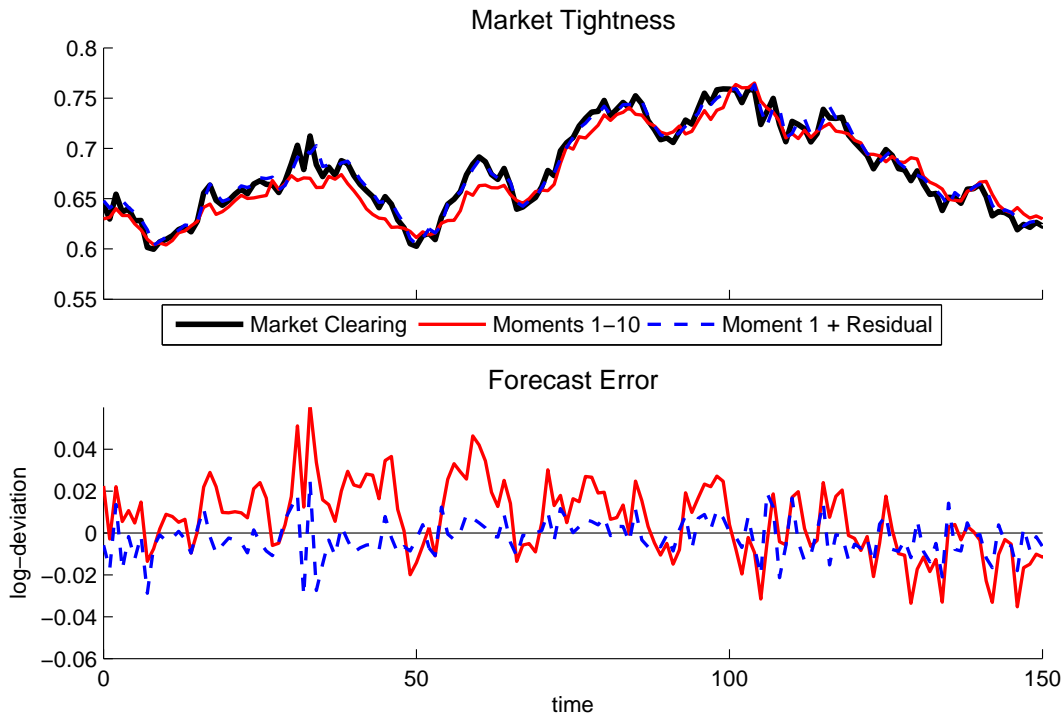


Figure 16: The upper part plots an example of market clearing values for market tightness in a simulation when $H = 5$. The dashed blue line draws the estimated values for market tightness using the residual method and the first moment of the distribution of positions. The red solid curve draws the estimated 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.

8 Conclusion

In this paper, 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 paper 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 to a low job-finding rate. Hence, this paper suggests that higher uncertainty is particularly bad for workers, because it decreases their chances of finding a job when they become unemployed. In den Haan et al. (2014), 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 paper suggests that macroeconomic variables will react stronger in response to policy interventions in times of higher unemployment.

References

- Bloom, Nick, Stephen Bond, and John Van Reenen (2007). Uncertainty and Investment Dynamics, *Review of Economic Studies*, 74: 391-415.
- Bloom, Nicholas (2009). The Impact of Uncertainty Shocks, *Econometrica*, 77(3): 623-685.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2012). Really Uncertain Business Cycles, *NBER Working Paper*, No. 18245.
- Coles, Melvyn G., and Ali Moghaddasi Kelishomi (2011). New Business Start-ups and the Business Cycle, *CEPR Discussion Paper*, No. 8588.
- Cooper, Russell W., John C. Haltiwanger, and Jonathan Willis (2004). Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications, *NBER Working Paper*, No. 10297.
- Cooper, Russell, John Haltiwanger, and Jonathan L. Willis (2007). Search frictions: Matching aggregate and establishment observations, *Journal of Monetary Economics*, 54: 56-78.
- Costain, James S., and Michael Reiter (2008). Business cycles, unemployment insurance, and the calibration of matching models, *Journal of Economic Dynamics and Control*, 2008(4): 1120-1155.
- Davis, Stephen J., R. Jason Faberman, and John Haltiwanger (2012). Labor market flows in the cross section and over time, *Journal of Monetary Economics*, 59: 1:18.
- Davis, Stephen J., and John Haltiwanger (1992). Gross Job Creation, Gross Job Destruction, and Employment Reallocation, *Quarterly Journal of Economics*, 107(3): 819-863.
- den Haan, Wouter J., Pontus Rendahl, and Markus Riegler (2014). Unemployment (Fears), Precautionary Savings, and Aggregate Demand.
- den Haan, Wouter J., Garey Ramey, and Joel Watson (2000), Job Destruction and Propagation of Shocks, *American Economic Review*, 90(3): 482-498.
- Elsby, Michael W. L., and Ryan Michaels (2013). Marginal Jobs, Heterogeneous Firms, and Unemployment Flows, *American Economic Journal: Macroeconomics*, 5(1): 1-48.
- Fujita, Shigeru, and Makoto Nakajima (2014). Worker Flows and Job Flows: A Quantitative Investigation, *Federal Reserve Bank of Philadelphia Working Paper*, No. 13-9/R.
- Fujita, Shigeru, and Garey Ramey (2007). Job Matching and Propagation, *Journal of Economic Dynamics and Control*, 31(11): 3671-3698.
- Guglielminetti, Elisa (2013). The Effects of Uncertainty Shocks on the Labor Market: A Search Approach, <https://sites.google.com/site/elisaguglielminetti/>.

- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng (2014). Measuring Uncertainty, *American Economic Review*, forthcoming.
- Hall, Robert E. (2005). Employment Fluctuations with Equilibrium Wage Stickiness, *American Economic Review*, 95(1): 50-65.
- Kaas, Leo, and Philipp Kircher (2013). Efficient Firm Dynamics in a Frictional Labor Market, <http://homepages.econ.ed.ac.uk/~pkircher/>.
- Khan, Aubhik, and Julia K. Thomas (2013). Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity, *Journal of Political Economy*, 121(6): 1055-1107.
- Krusell, Per, and Anthony Smith (1998). Income and Wealth Heterogeneity in the Macroeconomy, *Journal of Political Economy*, 106(5): 867-896.
- Leahy, John V., and Toni M. Whited (1997). The Effect of Uncertainty on Investment: Some Stylized Facts, *Journal of Money, Credit and Banking*, 28(1), 64-83.
- Mecikovsky, Ariel M., and Matthias Meier (2014). Do plants freeze upon uncertainty shocks?, <https://sites.google.com/site/amecikovsky/publications>.
- Michaillat, Pascal (2014). A Theory of Countercyclical Government Multiplier, *American Economic Journal: Macroeconomics*, 6(1): 190-217.
- Mortensen, Dale T., and Christopher A. Pissarides (1994). Job Creation and Job Destruction in the Theory of Unemployment, *Review of Economic Studies*, 61(3): 397-415.
- Petrongolo, Barbara, and Christopher A. Pissarides (2001). Looking into the Black Box: A Survey of the Matching Function, *Journal of Economic Literature*, 39(2): 390-431.
- Pissarides, Christopher A. (2009). The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?, *Econometrica*, 77(5): 1339-1369.
- Postel-Vinay, Fabien, and Jean-Marc Robin (2002). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity, *Econometrica*, 70(6): 2295-2350.
- Riegler, Markus (2014). Solving Heterogeneous Agents Models using Past Error Terms of the Law of Motion.
- Schaal, Edouard (2012). Uncertainty, Productivity and Unemployment in the Great Recession, <https://sites.google.com/site/edouardschaal/research>.
- Shimer, Robert (2005). The Cyclical Behavior of Equilibrium Unemployment and Vacancies, *American Economic Review*, 95(1): 25-49.
- Shimer, Robert (2012). Reassessing the ins and outs of unemployment, *Review of Economic Dynamics*, 15: 127-148.

- Stole, Lars A., and Jeffrey Zwiebel (1996). Intra-firm Bargaining under Non-Binding Contracts, *Review of Economic Studies*, 63(3): 375-410.
- van Ours, Jan, and Geert Ridder (1992). Vacancies and the Recruitment of New Employees, *Journal of Labor Economics*, 10(2): 138-155.
- Young, Eric (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.

Appendices

A Empirical part

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 paper.

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 by⁴³

$$F_t = 1 - \frac{u_{t+1} - u_{t+1}^S}{u_{t+1}}, \quad (35)$$

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

$$f_t = -\log(1 - F_t). \quad (36)$$

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

$$u_{t+1} = \frac{1 - \exp(-f_t - s_t) s_t}{f_t + s_t} l_t + \exp(-f_t - s_t) u_t. \quad (37)$$

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

⁴³In what follows, u_t denotes unemployment, u_t^S denotes unemployment less than 5 weeks, and l_t denotes employment.

taken.

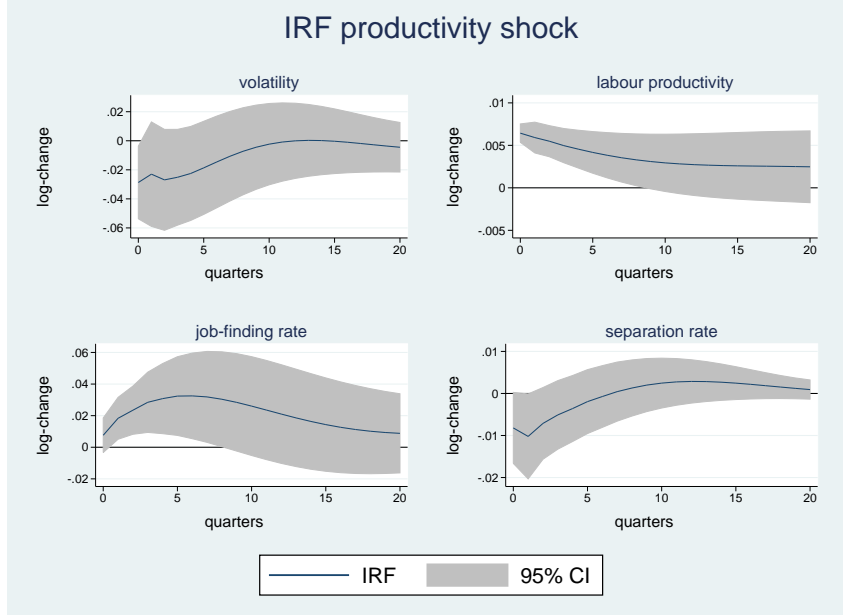


Figure 17: Impulse response functions after a shock to labour productivity in US data

A.2 Approximation of changes in unemployment

The law of motion of the unemployment rate is given by

$$u_{t+1} = s_t (1 - u_t) + (1 - f_t) u_t. \quad (38)$$

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

$$\bar{u} = \frac{\bar{s}}{\bar{s} + \bar{f}}. \quad (39)$$

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

$$u_{t+1} - \bar{u} = [1 - (\bar{s} + \bar{f})] (u_t - \bar{u}). \quad (40)$$

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):

$$u_{t+1} \approx \frac{s_t}{s_t + f_t}. \quad (41)$$

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

$$\begin{aligned}
\log \frac{u_{t+1}}{\bar{u}} &\approx \log \frac{s_t}{\bar{s}} - \frac{1}{s_t + f_t} \left(s_t \log \frac{s_t}{\bar{s}} + f_t \log \frac{f_t}{\bar{f}} \right) \\
&= \frac{f_t}{s_t + f_t} \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right) \\
&= (1 - \bar{u}) \left(\log \frac{s_t}{\bar{s}} - \log \frac{f_t}{\bar{f}} \right),
\end{aligned} \tag{42}$$

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

An analogous relationship can be derived for employment:

$$n_{t+1} \approx \frac{f_t}{s_t + f_t}, \tag{43}$$

$$\begin{aligned}
\log \frac{n_{t+1}}{\bar{n}} &\approx (1 - \bar{n}) \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right) \\
&= \bar{u} \left(-\log \frac{s_t}{\bar{s}} + \log \frac{f_t}{\bar{f}} \right).
\end{aligned} \tag{44}$$

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).

A.3 Robustness checks

This section shows that the impulse response function in Figure 2 is robust to alternative identification, to an alternative volatility measure, and to alternative detrending of the variables. In Figure 18, the volatility index is ordered first in the Cholesky decomposition, which implies that labour productivity can react instantaneously to volatility shocks. Figure 19 uses the volatility index VIX instead of the constructed index based on single stocks. Figures 20 and 21 use the Hodrick-Prescott filter to remove trends instead of removing a linear trend in the baseline. Figure 20 uses the smoothing parameter $\lambda = 100,000$ as preferred by Shimer (2005), whereas Figure 21 uses the usual value of $\lambda = 1,600$ for quarterly data.

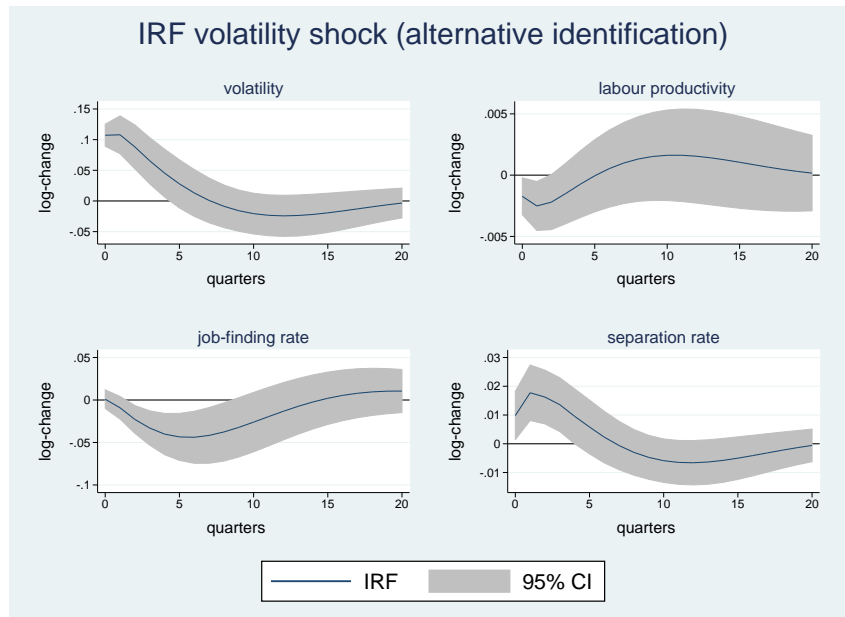


Figure 18: Impulse response functions after a volatility shock in US data using an alternative identification, in which volatility is ordered first in the Cholesky decomposition

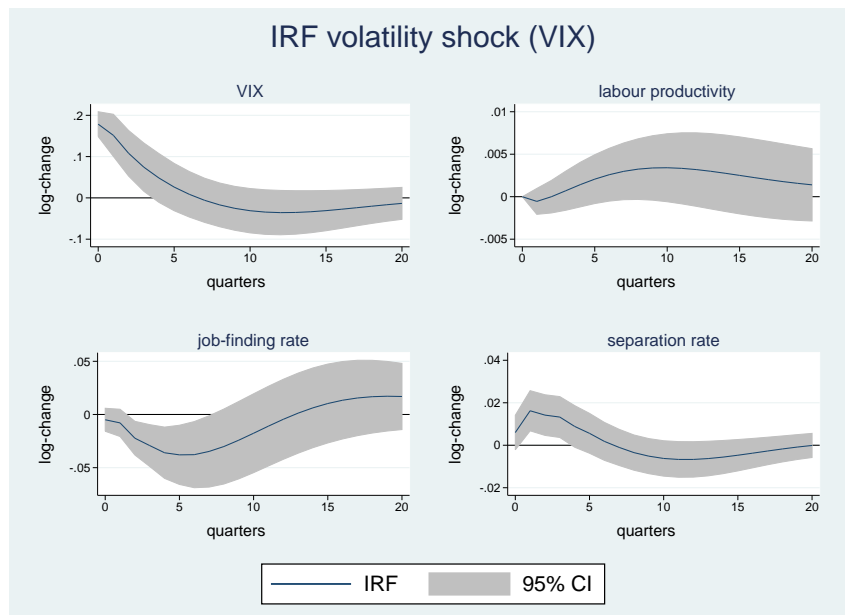


Figure 19: Impulse response functions after a volatility shock in US data. Instead of the volatility measure constructed from single stocks, the VIX is used as the volatility measure.

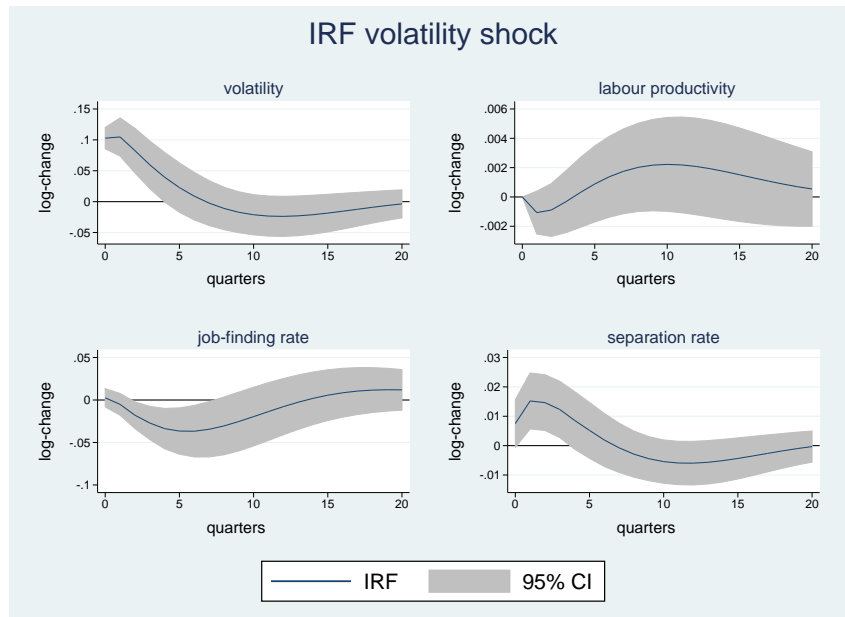


Figure 20: Impulse response functions after a volatility shock in US data. Instead of detrending using a linear trend, the series are detrended using the HP-filter with parameter $\lambda = 100,000$.

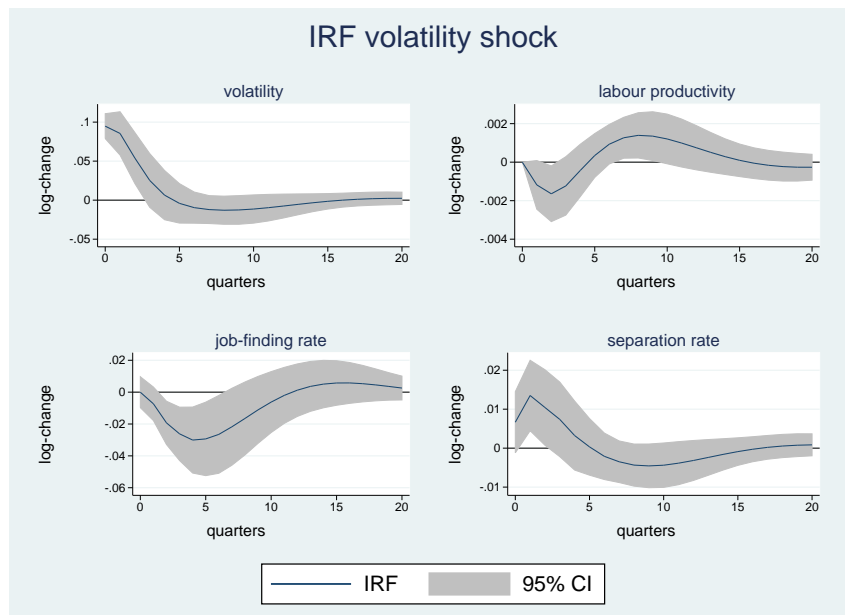


Figure 21: Impulse response functions after a volatility shock in US data. Instead of detrending using a linear trend, the series are detrended using the HP-filter with parameter $\lambda = 1,600$.

B Comparative statics

B.1 Sensitivity analysis

The purpose of this section is to show that the main result of this paper 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 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 (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 paper, 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 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 24 shows the results when the cost of creating positions is proportional to the filling cost $\frac{c}{q(\theta)}$.⁴⁴ The result is that for high values 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 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 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 27 shows the results for alternative parameters in the wage equation (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

⁴⁴Remember that in the initial steady state without aggregate uncertainty, $\frac{c}{q(\theta)}$ is normalized to be 1. Now, H changes in proportion to changes in $\frac{c}{q(\theta)}$ in the new steady state.

revenue function has been used. In the baseline scenario, I use an intermediate value of $\alpha = 0.75$. Figure 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 29 shows the results for different values of 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 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.

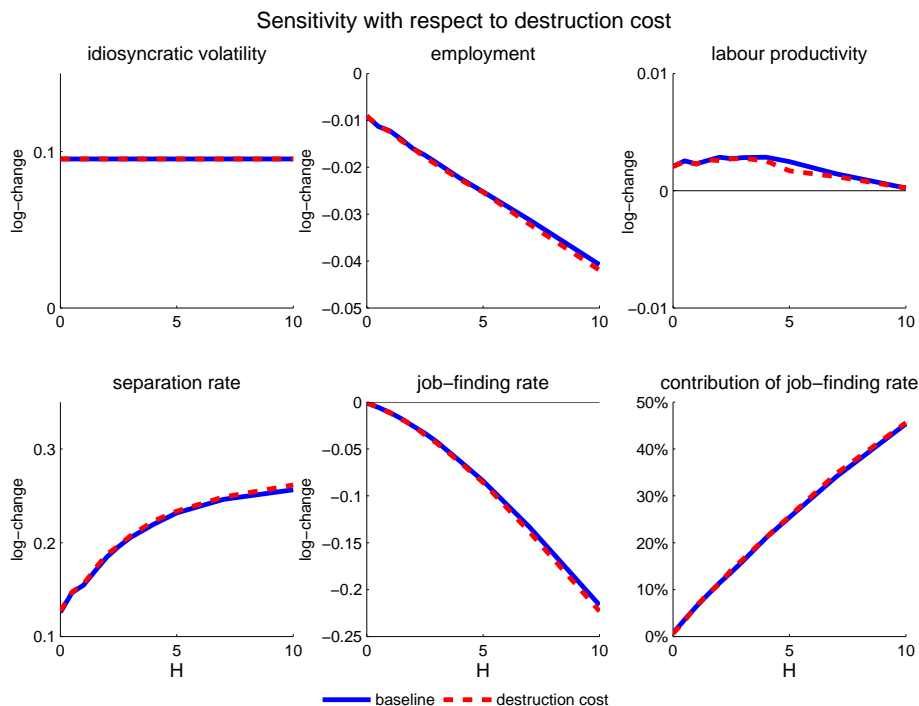


Figure 22: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the cost is not paid upon creation but upon voluntary destruction of a position.

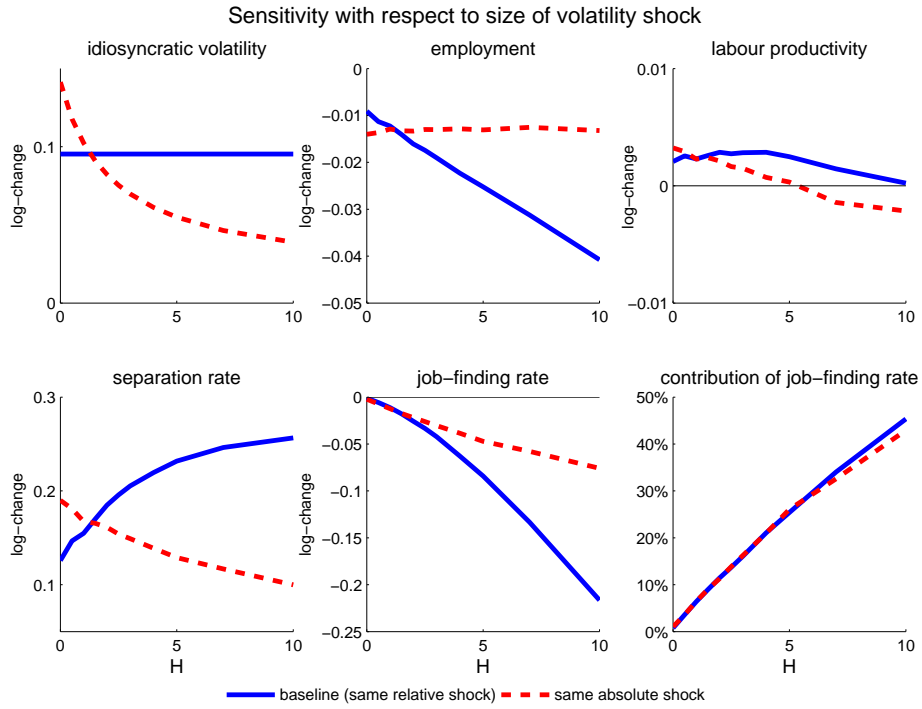


Figure 23: Responses after a permanent 10% increase compared with a 0.02 increase of idiosyncratic volatility for varying H .

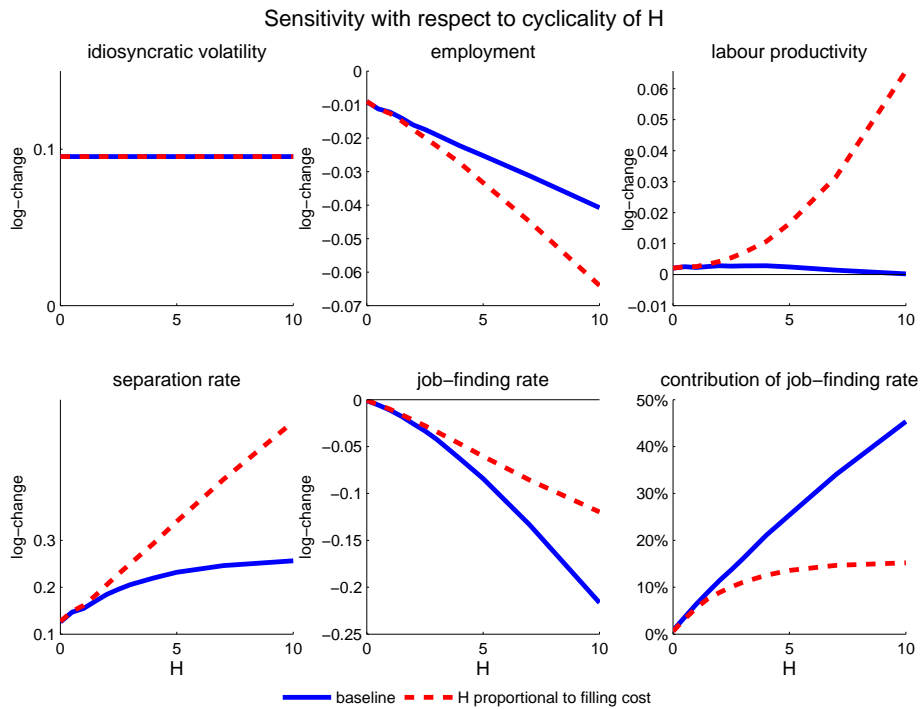


Figure 24: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the creation cost H is proportional to the filling cost $\frac{c}{q}$.

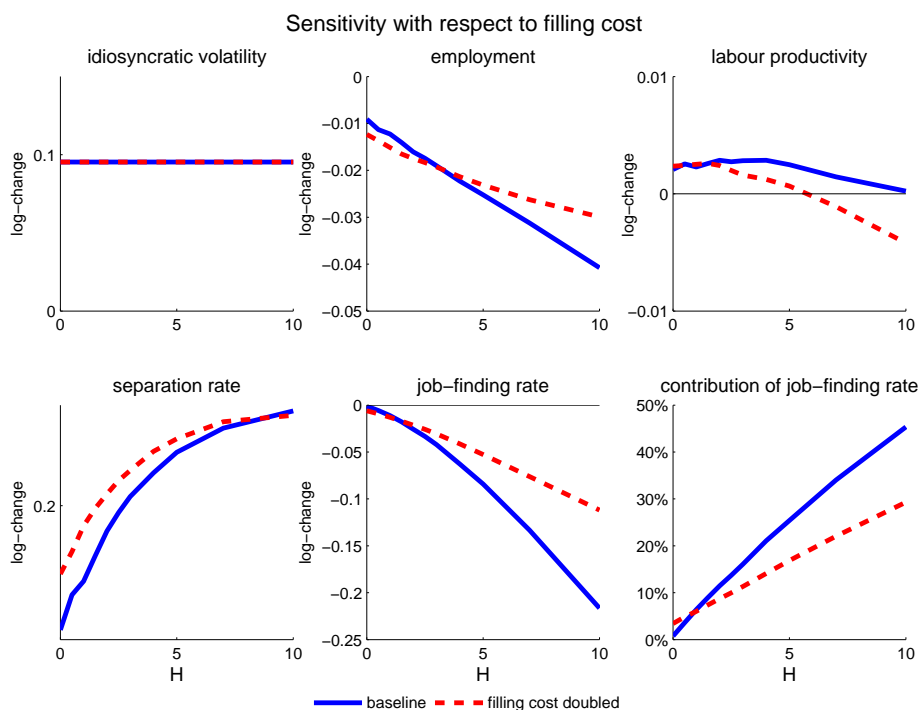


Figure 25: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solution, the target for the filling cost is doubled relative to the baseline scenario. The creation cost remains unchanged, such that $H = 10$ now corresponds to the creation cost being 5 times the filling cost in steady state.

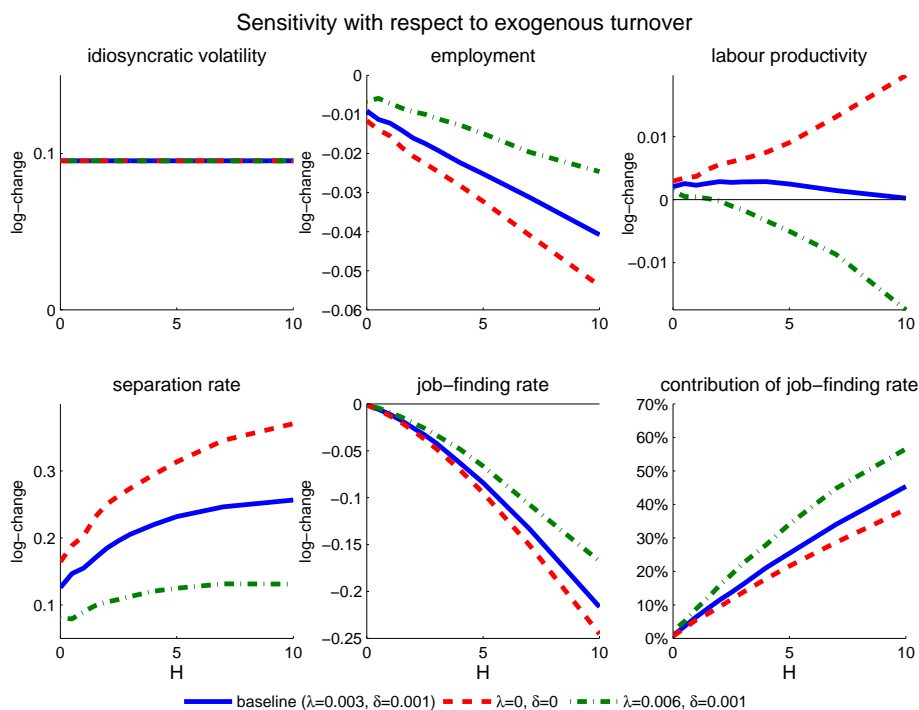


Figure 26: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution solves the model without depreciation of positions and without quits. The green dash-dotted solution solves the model for a doubled rate of quits.

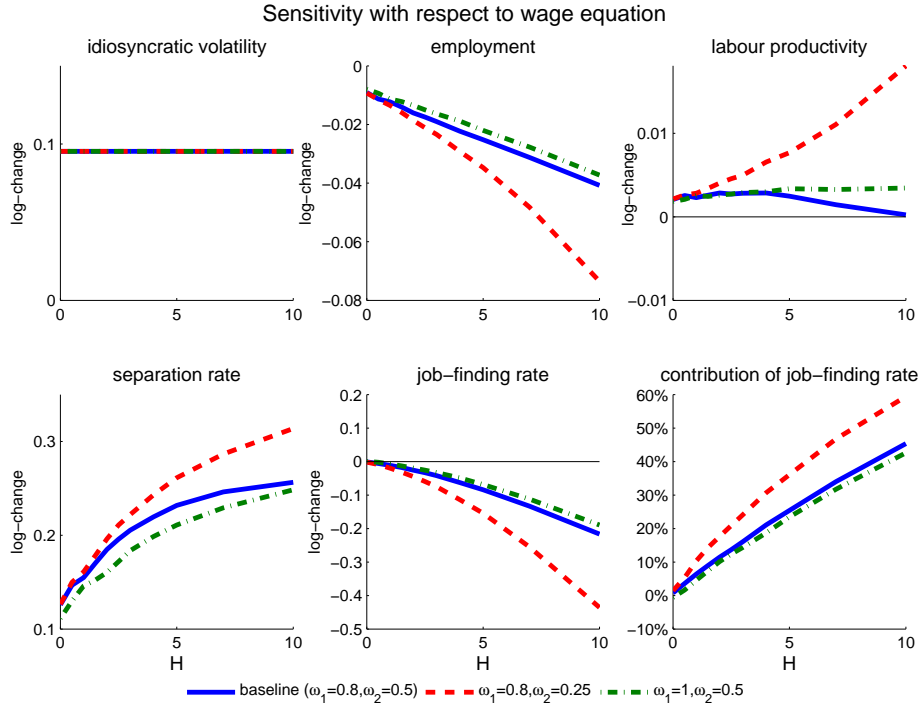


Figure 27: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution uses a wage equation that is less sensitive to changes in market tightness, and the green dash-dotted solution makes the wage independent of firm size.

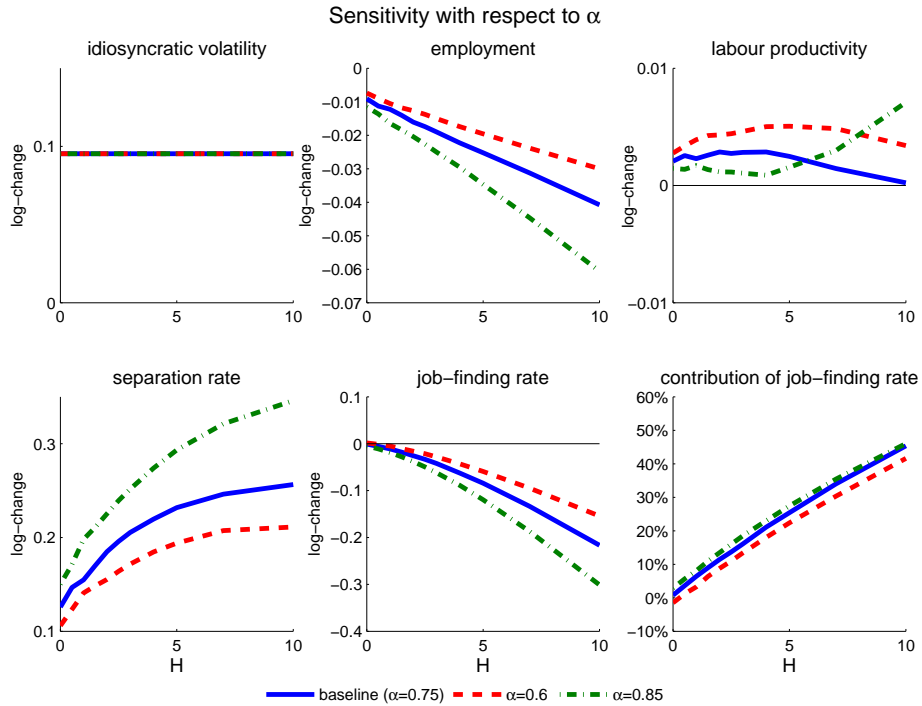


Figure 28: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The alternative solutions use different values of curvature in the production function.

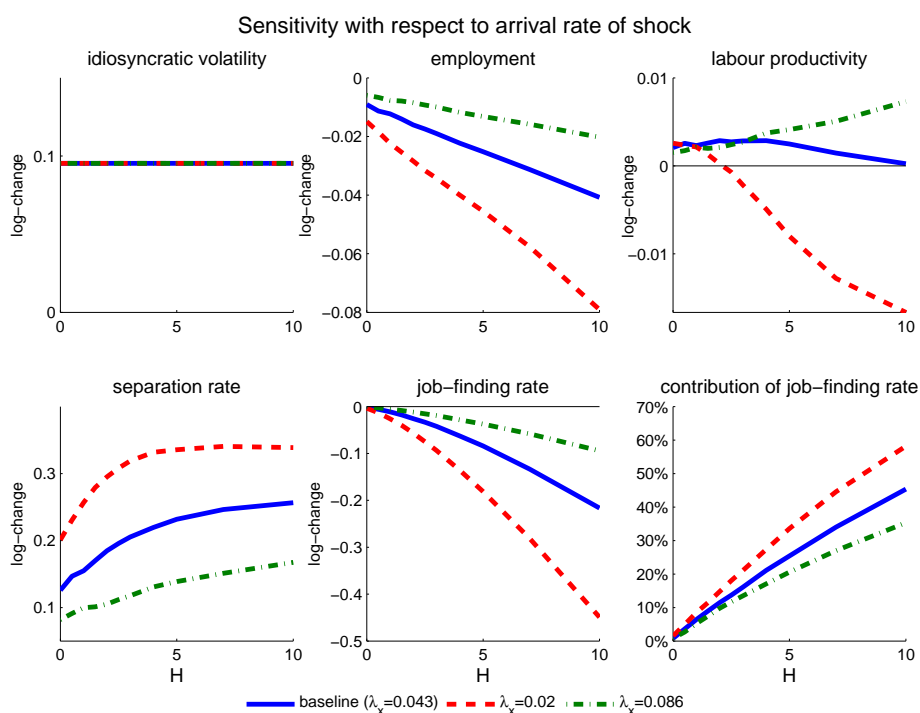


Figure 29: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . In the alternative solutions, the arrival rate of the idiosyncratic shock is varied.

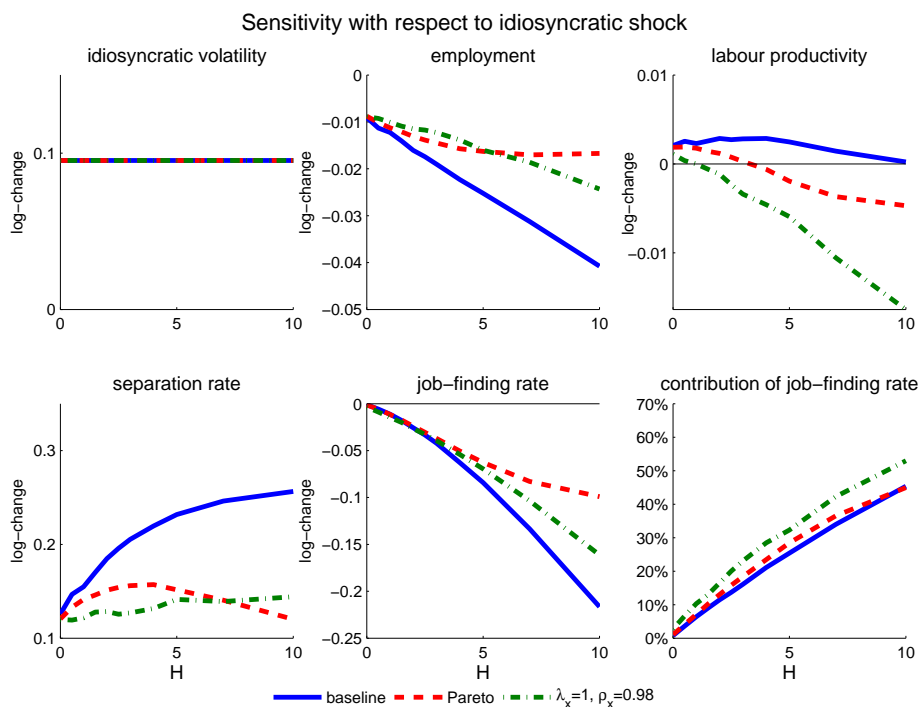


Figure 30: Responses after a permanent 10% increase of idiosyncratic volatility for varying H . The red dashed solution uses a Pareto distribution instead of the log-normal distribution for the idiosyncratic productivity draw. The green dash-dotted solution assumes that shocks arrive in each period but that productivity is autocorrelated.

C Decomposition exercise

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 (11) simplifies, and the number of positions is determined by

$$\omega_1 \alpha z x n^{\alpha-1} - (\omega_2 \theta + \omega_3) + \beta (1 - \delta) H + \beta (1 - \delta) (1 - \lambda) \mathbb{E} \left[\frac{c}{q(\theta(s'))} | s \right] = H + \frac{c}{q(\theta(s))}. \quad (45)$$

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:

$$\log \tilde{n}_{LB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log \left\{ \omega_2 \theta + \omega_3 + [1 - \beta(1 - \delta)] H + [1 - \beta(1 - \delta)(1 - \lambda)] \frac{c}{q} \right\}}{1 - \alpha}. \quad (46)$$

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 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.

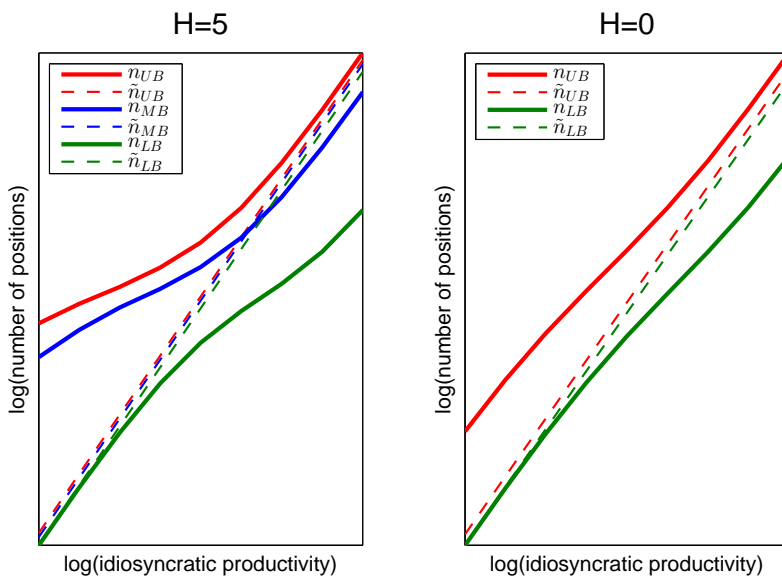


Figure 31: Optimal cutoffs (solid) and comparison cutoffs (dashed), that assume a fixed marginal continuation value.

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 (10), the middle cutoff is given by the condition

$$J(n_{MB}(x), x, s) = \omega_1 \alpha z x n_{MB}(x)^{\alpha-1} - (\omega_2 \theta + \omega_3) - \lambda \frac{c}{q(\theta(s))} + \dots \quad (47)$$

$$\begin{aligned} & \beta(1-\delta) \mathbb{E} [J((1-\delta)n_{MB}(x, s), x', s') | x, s] \\ & = (1-\lambda) \frac{c}{q(\theta(s))}. \end{aligned} \quad (48)$$

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:⁴⁵

$$\mathbb{E} [J((1 - \delta) \tilde{n}_{MB}(x), x', s') | x, s] = J(\tilde{n}_{MB}(x), x, s). \quad (49)$$

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

$$\log \tilde{n}_{MB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log \left\{ \omega_2 \theta + \omega_3 + [1 - \beta(1 - \delta)(1 - \lambda)] \frac{c}{q(\theta(s))} \right\}}{1 - \alpha}. \quad (50)$$

A comparison of this cutoff with the perfect insurance cutoff in equation (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 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 31.

The upper cutoff is implicitly defined by equation (6)

$$\begin{aligned} J(n_{UB}(x, s), x, s) &= (1 - \lambda) \left\{ \begin{array}{l} \omega_1 \alpha z x [(1 - \lambda) n_{UB}(x, s)]^{\alpha-1} - (\omega_2 \theta + \omega_3) + \dots \\ \beta(1 - \delta) \mathbb{E} [J((1 - \delta)(1 - \lambda) n_{UB}(x, s), x', s') | x, s] \end{array} \right\} \\ &= 0. \end{aligned} \quad (51)$$

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

$$\mathbb{E} [J((1 - \delta)(1 - \lambda) \tilde{n}_{UB}(x), x', s') | x, s] = 0. \quad (52)$$

Then the cutoff is given by

$$\log \tilde{n}_{UB}(x) = \frac{\log x + \log(\omega_1 \alpha z) - \log(\omega_2 \theta + \omega_3)}{1 - \alpha} - \log(1 - \lambda). \quad (53)$$

This cutoff is parallel to $\tilde{n}_{MB}(x)$ and $\tilde{n}_{LB}(x)$ in the $(\log x, \log n)$ -space⁴⁶. 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

⁴⁵Note 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.

⁴⁶In 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.

smaller for high levels of productivity because it becomes less likely for the firm to become even more productive in the future⁴⁷.

Let $d_j^\sigma(p)$ be the log-difference between cutoffs n_j and \tilde{n}_j , evaluated at the $100p^{th}$ -percentile, when the standard deviation of idiosyncratic productivity is σ .⁴⁸ Denoting the inverse cumulative distribution function of the idiosyncratic shock by $F_{\sigma_x}^{-1}(\cdot)$, the three differences can be written as

$$d_{LB}^\sigma(p) = \log n_{LB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{LB}(F_{\sigma_x}^{-1}(p)) \leq 0, \quad (54)$$

$$d_{MB}^\sigma(p) = \log n_{MB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{MB}(F_{\sigma_x}^{-1}(p)), \quad (55)$$

$$d_{UB}^\sigma(p) = \log n_{UB}(F_{\sigma_x}^{-1}(p)) - \log \tilde{n}_{UB}(F_{\sigma_x}^{-1}(p)) \geq 0. \quad (56)$$

C.2 Asymmetry of wait-and-see effect

The shifts of the cutoffs for the wait-and-see effect in Figure 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 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

⁴⁷In 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 $x' \leq x_{\max}$.

⁴⁸To keep notation simpler, I use subscript j , whenever an expression applies to all three cutoffs LB , MB , and UB .

wait-and-see effect becomes more relevant when H is large.

D Comparison with US time series

This appendix provides additional graphs for the exercise matching the US time series in Section 6.

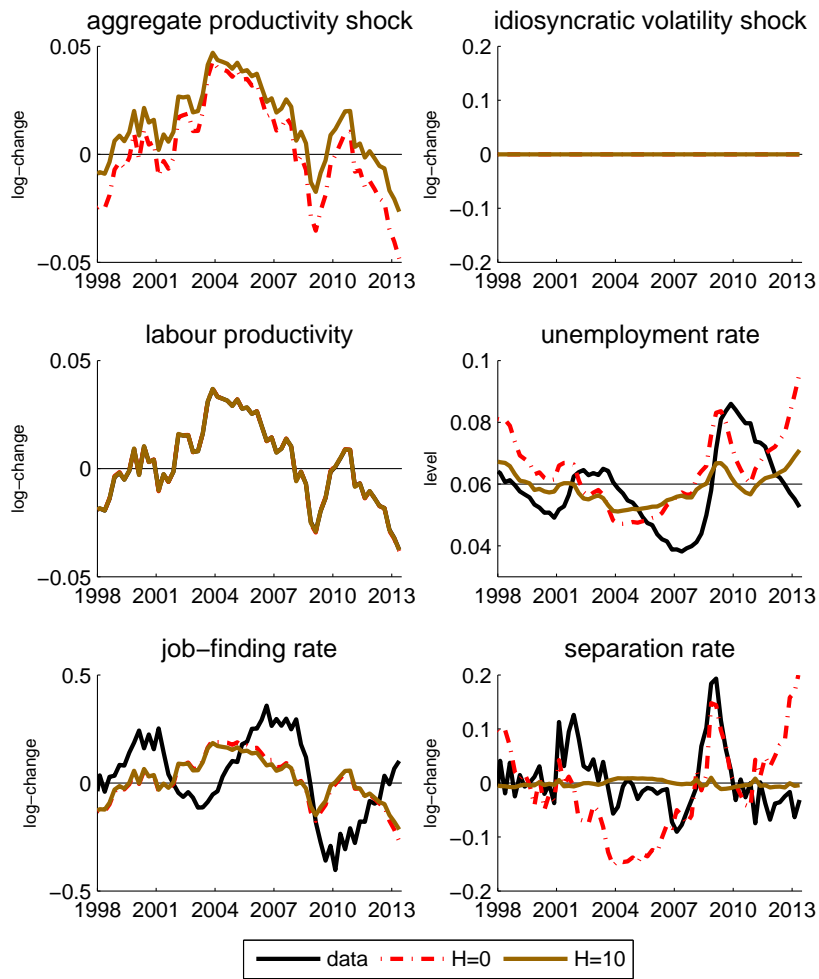


Figure 32: Model without uncertainty shocks: Aggregate productivity is estimated in each quarter to match labour productivity in the model to the data (linear trend removed). Relative to the baseline calibration, using $\omega_2 = 0.5$, a lower value of $\omega_2 = 0.2$ is used, which makes the wage less responsive to changes in market tightness.

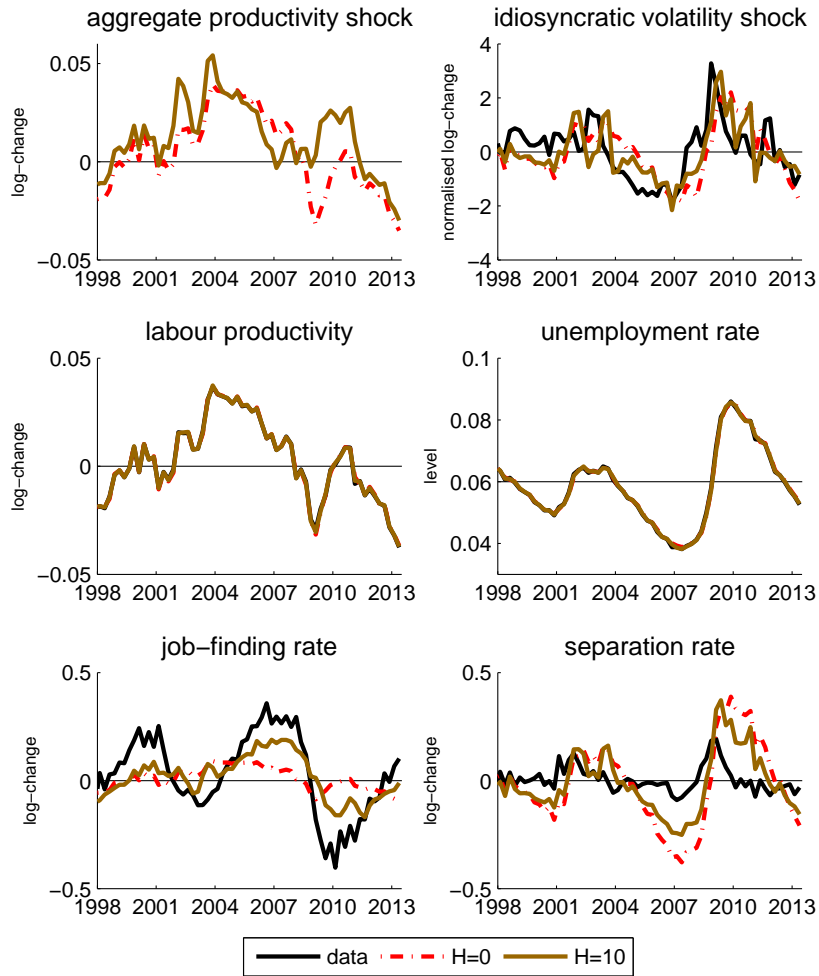


Figure 33: Model with aggregate productivity and uncertainty shocks: Aggregate productivity and idiosyncratic volatility are estimated in each quarter to match labour productivity and unemployment rate in the model to the data (linear trend removed). The top right panel compares the backed out volatility shocks with the empirical volatility index by normalising all series to have variance 1.