Lost Generations of Firms and Aggregate Labor Market Dynamics[☆]

Petr Sedláček

September 22, 2015

(first version November 2011)

Abstract

Can the unprecedented lack of startups during the Great Recession in the U.S. negatively impact the economy in future years? A structural model of firm dynamics and a frictional labor market suggests that, despite strong general equilibrium effects in the short run, a lost generation of firms creates a persistent dent in the employment potential of the economy. Estimating the model using aggregated firm data shows that had firm entry remained constant during the Great Recession, output would have recovered 4-6 years earlier and unemployment would have been 0.5 percentage points lower even 10 years after the crisis.

Keywords: firm age, firm dynamics, heterogeneous firms, unemployment

JEL Classification: E24, E32, J64, M13

[☆]This paper is a substantially revised part of my Ph.D. dissertation written at the University of Amsterdam. I thank my advisor Wouter den Haan for his support. I am also grateful for useful comments from Eric Bartelsman, Christian Bayer, Michael Elsby, Bart Hobijn, Philip Jung, Leo Kaas, Philip Kircher, Keith Kuester, Matija Lozej, Fabien Postel-Vinay, Richard Rogerson, Nicolas Roys, Vincent Sterk, Christian Stoltenberg, Emily and Robert Swift, Antonella Trigari, Sweder van Wijnbergen and seminar participants at CERGE, the DNB, the ECB, the Universities of Amsterdam, Autonoma de Barcelona, Bonn, Carlos III, Edinburgh and Louvain.

Address: University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany. Tel.: +49 228 739 236. Email: sedlacek@uni-bonn.de

1. Introduction

The number of startups in the U.S. hit an all time low at the end of the Great Recession in 2010 (31 percent below its pre-crisis level). At the same time, the unemployment rate peaked at 10 percent and remained close to this level even two years after the official end of the downturn. This paper asks to what extent these two phenomena are interlinked and what the lost generation of firms in the Great Recession implies for the U.S. economy in the medium- to long-run.

As a first step, this paper highlights that changes in firm entry impact the economy not only directly, but also indirectly in later years as affected cohorts of startups age. In particular, using Business Dynamics Statistics (BDS) data it is shown that the pro-cyclical nature of firm entry creates a ripple effect resulting in pro-cyclical movements of the share of young firms (not older than five years). Moreover, it is shown that young firms account for 40 percent of aggregate employment fluctuations (even though they employ only 16 percent of all workers). Cyclical changes in the firm age distribution therefore help shape aggregate fluctuations. These findings complement the results in Haltiwanger, Jarmin, and Miranda (2013) regarding young firms' disproportionate contributions to aggregate job creation in the long-run.[1]

The above findings raise concerns about the medium- to long-run impact of the recently lost generation of firms. A simple simulation of an exogenous drop in the number of startups (of the magnitude observed in the Great Recession) together with fixed survival and growth rates of incumbent firms suggests that the impact may be severe: even 10 years after the shock subsides, the unemployment rate remains more than 1 percentage point above its initial level.

However, the simple simulation exercise abstracts from potential general equilibrium feedback effects that may dampen the unemployment rate response. Therefore, in the next

[1] A related paper is Pugsley and Sahin (2014) who analyze the effect of the secular decline in the share of startups on the aggregate economy.

step this paper uses a general equilibrium model of firm dynamics and a frictional labor market to demonstrate that while such effects are indeed important in the short-run, periods of subdued entry do negatively impact the economy in the long-run. The reason is that in the short run incumbent firms take advantage of the lack of job creation by startups and they almost fully compensate for the drop in employment. However, in future years the missing entrants generate fewer older firms (which on average account for the bulk of aggregate employment). This creates a very persistent dent in the employment potential of the economy essentially raising the "natural" rate of unemployment.

The general equilibrium effects dampening the short-run impact of a drop in firm entry operate mainly through the frictional labor market. In particular, because new (young) firms account for a large chunk of overall hiring, an exogenous drop in startups leads to a fall in aggregate vacancies. The tighter labor market makes it easier for incumbent firms to hire. Moreover, employees become more reluctant to leave their current jobs because their outside options worsen. This resembles the "insulation effect" of recessions pointed out by Caballero and Hammour (1994). These factors are also reflected in the drop of wages further promoting new hiring. Finally, lower wages raise profits and thus induce an endogenous *increase* in the number of startups following the initial (exogenous) drop.

In order to quantify the impact of the lost generation of firms in the Great Recession, the structural model is *estimated* using BDS data. This enables us to determine to what extent fluctuations in firm entry were driven by forces specific to startups (i.e. not directly related to incumbents) and to what extent they were an endogenous reaction to a shock common to all firms. The results suggest that about 40 percent of the Great Recession drop in the number of startups can be attributed to factors specific to new firms, while the remainder is an endogenous response to a particularly strong recession.

Moreover, the estimated time-paths of the structural shocks and the model variables allow us to conduct counterfactual scenarios. The results reveal that had the number of startups remained at its pre-crisis level during and in the aftermath of the crisis, the immediate impact on the aggregate economy would have been relatively small. However, the larger share of young firms would have helped to speed up the recovery in later years. Specifically, output would have reverted back to its trend 4-6 years earlier and the unemployment rate would have been 0.5 percentage points lower even 10 years after the end of the crisis.

By documenting the cyclicality of the firm-age distribution and its effect on aggregate employment, this paper extends the results in Haltiwanger, Jarmin, and Miranda (2013) and Fort, Haltiwanger, Jarmin, and Miranda (2013) regarding the job creation prowess of young firms in the long-run and the cyclicality of young and small businesses. The structural model relates to several recent studies extending versions of the Mortensen and Pissarides (1994) model to include multi-worker firms and firm dynamics.[2] The focus on firm entry is related to Samaniego (2008), Lee and Mukoyama (2013) and Clementi and Palazzo (2013) who study entry and exit patterns within an extended version of the Hopenhayn and Rogerson (1993) model and to Siemer (2014) who investigates the impact of a drop in the number of startups after a financial shock within a partial equilibrium framework. In contrast to the above studies, the presented model focuses on firm entry in a search and matching framework highlighting the importance of general equilibrium feedback effects operating via the frictional labor market which go beyond an endogenous response of wages.

The next section provides empirical evidence on young firms' contributions to aggregate employment variation and uses a simple exercise to show that a missing generation of firms has a potentially very severe negative impact on the aggregate labor market. Section 3 builds a structural model with firm dynamics and a frictional labor market and Section 4 describes the solution method, calibration and the quantitative properties of the model. Section 5 presents the main results and the last section offers some concluding remarks.

[2] See e.g. Elsby and Michaels (2013), Gertler and Trigari (2009), Kaas and Kircher (2014) and Schaal (2012).

2. Lost generations of firms and aggregate employment in the data

The Great Recession in the US was accompanied by a unprecedentedly large drop in the number of startups. According to the Business Dynamics Statistics data, the number of startups hit an all time low in 2010, 31 percent below the pre-crisis level.[3] Even in 2012, three years after the official end of the recession, entry of new businesses was still 26 percent below its level in 2006.

This section first documents how the cyclicality of firm entry and subsequent changes in the number of young firms contribute to aggregate employment variation. Next, it shows that while the Great Recession was not an exception qualitatively, young firms (and especially the lack of their numbers) contributed strongly to the observed drop in aggregate employment during and in the aftermath of the crisis. Finally, a simple counterfactual exercise suggests that the observed lack of firm entry during the latest downturn may have a substantial and very persistent negative impact on the aggregate labor market.

This section uses firm data taken from the Business Dynamics Statistics of the Census Bureau.[4] The information is based on administrative records and covers almost all US private employment (98 percent). The annual data represents a snapshot taken in March of each year and it runs from 1977 until 2012. As is common in the literature, young firms are defined as those younger than six years and old firms account for the rest.

2.1. Contribution of young firms to aggregate employment variation on average

When it comes to firm entry, the latest downturn was not an exception qualitatively. It is well known that plant entry is pro-cyclical in the manufacturing sector (see e.g. Campbell, 1999; Lee and Mukoyama, 2013). This pattern also holds for the aggregate economy. De-

[3] Also the entry rate fell by almost 30% during this period (from 10.8% in 2006 to 7.7% in 2010). Part of this fall can, however, be attribute to the secular decline in the entry rate (see e.g. Pugsley and Sahin, 2014, for an analysis of the secular trends in the entry rate and the firm age distribution).

[4] A firm is a business organization consisting of one or more establishments that were specified under common ownership or control. An establishment is defined as a single physical location where business is conducted or where services or industrial operations are performed. Results based on establishment data are very similar.

Figure 1: Young firm share and real GDP; cyclical components

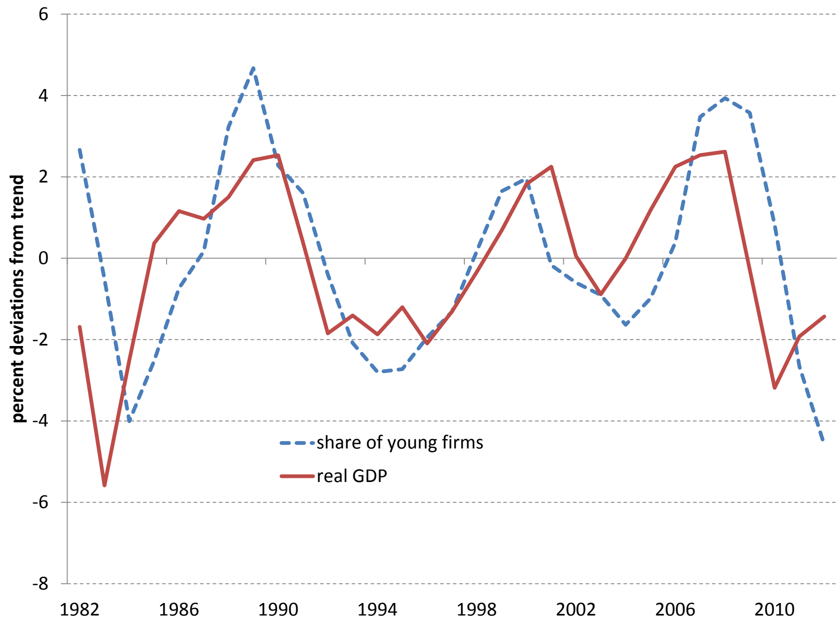

Notes: The figure plots the cyclical components of the (log of) the share of young firms (0 to 5 years of age) and real GDP. The correlation coefficient between the two series is 0.73.

pending on the detrending method, the correlation of the number of startups and real GDP varies between 0.30 (linear trend) and 0.66 (HP-filter).[5]

Importantly, fluctuations in firm entry affect aggregate employment not only directly, but also indirectly in later years as cohorts of startups age. For instance, only about 2% of the variation in the number of young firms is driven by changes in their survival rates, while the rest is accounted for by fluctuations in (past) firm entry.[6] This ripple effect of firm entry is then reflected in the pro-cyclical nature of the share of young firms in the economy (Figure 1).

The business cycle variation in the firm-age distribution apparent in Figure 1 may be inconsequential for aggregate labor market fluctuations as long as young and old firms contribute in similar magnitudes to overall employment variation. Let us decompose the variance

[5]Throughout the paper the smoothing coefficient in the HP-filter is set to 1600 (100) for quarterly (annual) data.

[6]Specifically, constructing a counterfactual young firms' employment series based on survival rates being fixed to their sample averages results in a series which is 2% less volatile than the data.

of aggregate employment into contributions of young and old firms as follows

$$var\left(\frac{\widehat{E}_t}{\overline{E}_t}\right) = cov\left(\frac{\widehat{E}_t^y}{\overline{E}_t}, \frac{\widehat{E}_t}{\overline{E}_t}\right) + cov\left(\frac{\widehat{E}_t^o}{\overline{E}_t}, \frac{\widehat{E}_t}{\overline{E}_t}\right) + cov\left(\epsilon_t, \frac{\widehat{E}_t}{\overline{E}_t}\right),$$

where a bar stands for an HP-filter trend and a hat indicates deviations from this trend.[7] E_t denotes aggregate employment in period t, E_t^y and E_t^o denote employment by young and old firms, respectively, and ϵ is a residual term coming from the detrending method.

Expressing the relative contributions in percentage terms of total variance, young firms account for 38% of all fluctuations in aggregate employment on average. This number is rather striking when compared to the employment share of young firms which amounts to only 16%.[8] This result, related to business cycle frequencies, complements recent findings that young firms on average account for a disproportionately large fraction of overall net job creation (Haltiwanger, Jarmin, and Miranda, 2013). Combined, these facts suggest that a lost generation of young firms may not only lead to a sluggish recovery in the short run, but possibly also to a persistent drag on aggregate employment.

Further decomposing employment changes of young firms reveals that about half of the variation is accounted for by fluctuations in the number of young firms and the rest is due to changes in their average size.[9]. Recalling that variation in firm survival rates accounts for very little of the changes in the number of young firms suggests that firm entry is responsible for about $15-20\%$ of aggregate employment fluctuations (even though the employment share of startups is only 3%). Moreover, this constitutes a lower bound because part of the

[7]The components combine the effect of employment growth rates in young and old firms together with the changes in their shares: $\frac{\widehat{E}_t^y}{\overline{E}_t} = \frac{\widehat{E}_t^j}{\overline{E}_t^j}\frac{\overline{E}_t^j}{\overline{E}_t}$ for $j = y, o$.

[8]This is consistent with Fort, Haltiwanger, Jarmin, and Miranda (2013) who find that young/small businesses are more volatile than old/large ones. It is also not inconsistent with Moscarini and Postel-Vinay (2012) who document that the differential growth rate between small and large businesses is pro-cyclical. As Fort, Haltiwanger, Jarmin, and Miranda (2013) point out, small-old businesses, of which there are many in the BDS data (e.g. 60 percent of firms aged 25 and older has less than 10 employees), are much less cyclically sensitive than small-young firms.

[9]Conducting the decomposition on startups shows that 58% of employment variation is driven by changes in the number of startups. For young firms the contribution is 42%.

Figure 2: Percentage changes in the number of startups and the unemployment rate

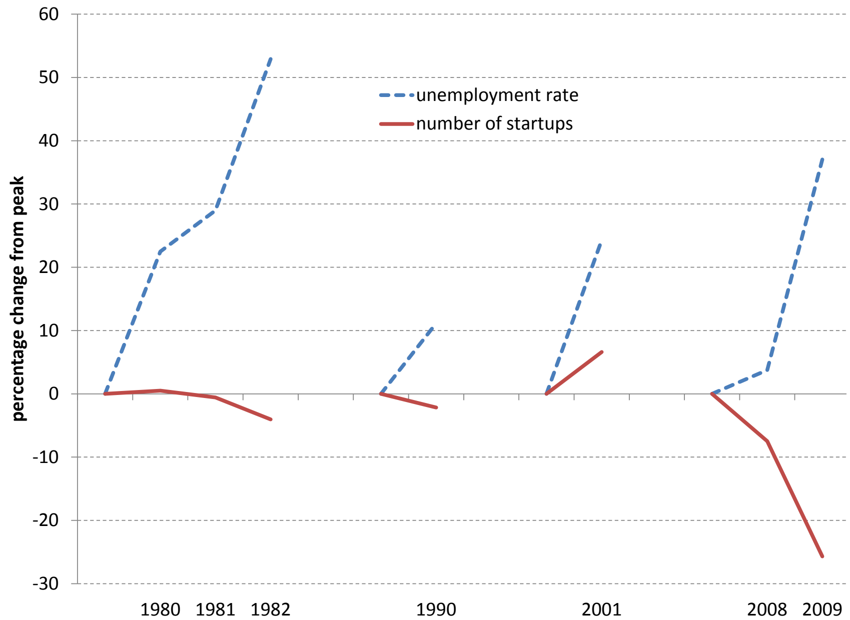

Notes: The figure plots the log changes in the number of startups and the unemployment rate both from the respective values prior to the start of each recession.

contribution of old firms to aggregate employment fluctuations could also be traced back to past changes in firm entry.

2.2. Lost generation of firms during the Great Recession and the potential long-run impact

Let us now focus specifically on the latest downturn for which the common view is that startups (and young firms) were hit particularly hard (see e.g. Fort, Haltiwanger, Jarmin, and Miranda, 2013; Sahin, Kitao, Cororaton, and Laiu, 2011). The exceptional nature of the Great Recession is also highlighted in Figure 2 which shows cumulative percentage drops (increases) in the number of startups (unemployment rate) during different recessions in the sample.[10]

Figure 2 clearly shows that during the Great Recession startup activity dropped dramatically even when taking the severity of the recession into account. In particular, between 2007 and 2009 44% of the cumulative drop in aggregate employment was driven by young

[10] A similar picture is painted if one considers percentage point changes in the unemployment rate, or changes in the growth rate of real GDP.

Figure 3: Unemployment rate; actual and counterfactual

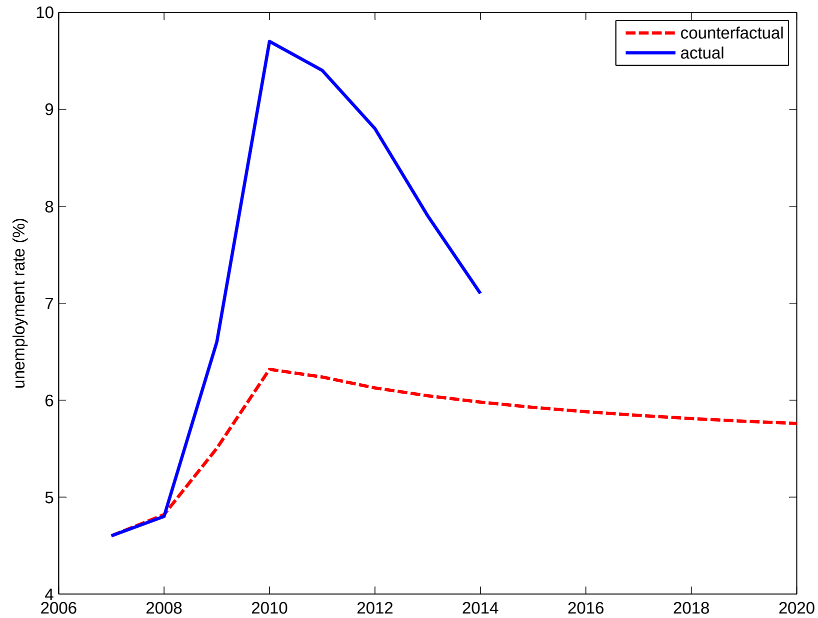

Notes: The figure plots the actual unemployment rate together with a counterfactual one based on a fixed firm life-cycle as observed in the BDS and the observed drop in startups during the Great Recession.

firms. In the aftermath of the crisis (between 2010 and 2012), when aggregate employment was still 5.1 percent below its pre-crisis level, this contribution was 70 percent. Moreover, three quarters of young firms' cumulative employment drop during and in the aftermath of the Great Recession was driven by a decline in their numbers, rather than a fall in their average size.

What is the potential medium- to long-run impact of the lost generation of firms coming from the Great Recession? Let us use a simple simulation to gauge how the unemployment rate would have responded to the observed firm entry drop but holding all else unchanged. Similar to Gourio, Messer, and Siemer (2014), Figure 3 shows a counterfactual unemployment rate based on an exogenous drop in the number of startups of the magnitude observed in 2008-2009, but where survival and growth rates of all other firms are fixed to their sample averages.

Given the relatively small (3 percent) employment share of startups, the short-term impact of the drop in firm entry is mild (pushing unemployment up to 6.3% in 2010). Therefore, the lack of startups alone cannot explain the observed unemployment rate increase during the crisis. More strikingly, however, the negative effect is extremely persistent. Even ten years after entry reverts back the unemployment rate remains more than 1 percentage point above its pre-crisis level.

While this simple exercise gives an indication of the potential importance of lost generations of firms for aggregate labor market dynamics, the assumption of unchanged behavior of incumbent firms is unrealistic. The rest of the paper is therefore devoted to analyzing the effects of lost generations of firms within a structural model of firm dynamics in which businesses optimally hire and fire workers on a frictional labor market.

3. Labor market model with firm dynamics

This section builds a structural general equilibrium model of a frictional labor market with endogenous firm dynamics. This framework is particularly suitable for the question at hand because it allows for the possibility that a drop in firm entry feeds back into the employment behavior of incumbent firms through several channels. In particular, changes in the number of entrants will not only be reflected in bargained wages, but also in the chances of incumbent firms hiring new workers and via their effect on workers' outside options also in the probabilities of the employed separating from their current jobs.

Among other things, the model will be able to quantify to what extent the drop in firm entry was driven by a shock common to all firms or a shock specific to startups. A key prerequisite of such a decomposition is a model which can deliver realistic fluctuations in firm entry. Many existing models of firm dynamics have trouble replicating entry patterns with only aggregate productivity shocks (see e.g. Samaniego, 2008; Kaas and Kircher, 2011). The setup of firm entry in this paper, explained in detail in Subsection 3.4, does not suffer from such a drawback.

The following paragraphs describe the model economy. To facilitate the exposition of the model, aggregate variables are denoted by upper case letters, firm- or worker-specific variables are denoted by lower case letters and next period values are indicated by a prime.

3.1. Matching in the labor market

In each period, unemployed workers search for jobs and firms that wish to hire employees post vacancies. The total number of unemployed (U) and the total number of posted vacancies (V) engage in (random) matching on the aggregate labor market. The number of new hires is determined by an aggregate matching function

$$M = mU^{\mu}V^{1-\mu}, \tag{1}$$

where I follow the majority of the literature and assume a Cobb-Douglas functional form with m being the (constant) matching efficiency and μ being the elasticity of matches with respect to the number of unemployed. Given the above, the probability with which an unemployed worker finds a firm is given by $F = M/U$ and the probability with which a firm hires a worker is given by $Q = M/V$.

Already at this point it is possible to highlight one of the general equilibrium effects of a fall in the number of startups. In particular, the total number of vacancies is a sum over vacancies posted by individual firms. Therefore, a drop in firm entry will, ceteris paribus, decrease the total number of vacancies. This in turn increases the probability of hiring a worker (Q) making it easier for existing firms to compensate for the lack of job creation by startups.

3.2. Household behavior

The economy is populated by a representative household consisting of a continuum of risk-neutral and ex-ante homogeneous workers. Workers can find themselves either unemployed (and searching for jobs) or employed by one of the heterogeneous firms. At the beginning of each period, employed workers obtain an iid draw of (worker-specific) productivity z from a distribution $H(z)$. Particularly bad draws will result in the employment relationships that are not profitable and such workers will be fired. Let us denote the cutoff value of worker-specific productivity below which employment relationships are severed by $\widetilde{z}$.

Household members pool their income from work and non-work activities and spend it on the consumption good. Formally, the household maximizes the present value of life-time utility (i.e. the present discounted value of consumption), subject to a budget constraint

$$C = bU + W + P,$$

where C is aggregate consumption, b is the value of home production, W is aggregate wage income and P are aggregate profits (the latter two are defined at the end of the section).

3.3. Firm behavior

This subsection describes the behavior of incumbent firms of which there is an endogenous mass. After the realization of aggregate shocks, but prior to observing worker-specific shocks, firms bargain with their workers over wages (w) that will be paid out in the current period. Thereafter, worker-specific productivity shocks are realized and firms decide to fire a fraction of their workforce for which the idiosyncratic productivity shocks were particularly bad, i.e. those with productivity draws below $\widetilde{z}$. Firms pay out the bargained wages to the remaining employees and produce output.

While it is assumed that all firms operate the same decreasing returns to scale production technology, they differ in the efficiency with which they operate it. In particular, there is a finite number of (permanent) technology types, indexed by $i = 1, 2, ..., I$, which differ in the level of total factor productivity (ϵ). The production function, which uses labor as its only input, is assumed to take on the following form

$$y = A\epsilon\widehat{z}n^{\alpha},$$

where A is the level of aggregate productivity, $\widehat{z} = \int_{\widetilde{z}} z \frac{h(z)}{1-H(\widetilde{z})} dz$ is average worker-specific productivity of employees who remain in the firm, n is the number of workers in production and α is the parameter of decreasing returns to scale.

After observing all shocks, firms post vacancies on a frictional labor market to attract new workers for production in the next period. The costs of hiring are assumed to take on the following form

$$\frac{\kappa}{\gamma}x^{\gamma}n,$$

where $x = v/n$ is the vacancy rate (vacancies over employment) and $\gamma > 1$ and $\kappa > 0$ are parameters. This functional form is borrowed from Merz and Yashiv (2007) and states that the costs of hiring are proportional to the size of the firm and that they are convex in the hiring rate (see e.g. Kaas and Kircher, 2014; Gertler and Trigari, 2009, for models using such a functional form).

At the end of each period, firms face an exogenous but age-dependent probability of shutting down. The assumption of an exogenous exit probability is made not only for greater tractability, but is also justified by the BDS data which indicate that job destruction by exiting firms accounts for only one percent of the variation in aggregate employment.[11]

Notice that under the above assumptions all firms of the same type and age make the same decisions. This property greatly increases the tractability of the model and simplifies the solution method. In what follows I will therefore index individual firms by their type (i) and age (a).

Formally, an incumbent firm maximizes expected firm value ($\Pi_{i,a}$) by choosing employment available at the beginning of the next period ($\widetilde{n}'_{i,a+1}$), the vacancy rate ($x_{i,a}$) and the cutoff for worker-specific productivity ($\widetilde{z}_{i,a}$) below which workers get fired, subject to the law of motion for firm-specific employment:

$$\Pi_{i,a} = \max_{\widetilde{n}'_{i,a+1}, x_{i,a}, \widetilde{z}_{i,a}} \left[y_{i,a} - w_{i,a}n_{i,a} - \frac{\kappa}{\gamma}x_{i,a}^{\gamma}n_{i,a} + \beta(1-\delta_a)\mathbb{E}\Pi'_{i,a+1} \right] \quad \text{s.t.}$$

$$\widetilde{n}'_{i,a+1} = (1 - H(\widetilde{z}_{i,a}))(1 + Qx_{i,a})\widetilde{n}_{i,a},$$

[11]Nevertheless, as a robustness check the Appendix shows that the results change little when considering realistic variation in firm survival rates.

$$y_{i,a} = A\epsilon_i \widehat{z}_{i,a} n_{i,a}^{\alpha},$$

$$n_{i,a} = (1 - H(\widetilde{z}_{i,a}))\widetilde{n}_{i,a},$$

where δ_a is the age-dependent exogenous rate of firm exit and $H(\widetilde{z}_{i,a})$ is the endogenous separation rate defined as the fraction of workers who obtain a productivity draw below $\widetilde{z}$. The resulting first-order conditions can be combined into a "optimal hiring" condition describing the optimal hiring behavior and a condition implicitly defining the worker-specific productivity cutoff value. The optimality condition is given by

$$\frac{\kappa x_{i,a}^{\gamma-1}}{Q} = \beta(1 - \delta_a)\mathbb{E}\mathcal{J}'_{i,a+1}, \tag{2}$$

where $\mathcal{J}_{i,a} = \frac{\partial \Pi_{i,a}}{\partial \widetilde{n}_{i,a}}$ is the beginning-of-period marginal value of a job for the firm. The optimal hiring condition therefore takes on a familiar form where the effective marginal costs of posting a vacancy are equal to the expected marginal benefits.

The condition defining the worker-specific productivity cutoff also balances the benefits and costs of changing the separation cutoff ($\widetilde{z}_{i,a}$). Intuitively, raising the separation cutoff will increase average worker-specific productivity of the remaining employees, but at the same time it will reduce the size of the workforce in the firm. This is formalized in the following condition which implicitly defines the separation cutoff and which anticipates that wages depend on the number of workers in the firm:

$$\frac{\partial y_{i,a}}{\partial \widetilde{z}_{i,a}} - \frac{\partial w_{i,a}}{\partial \widetilde{z}_{i,a}} n_{i,a} = -\frac{\partial n_{i,a}}{\partial \widetilde{z}_{i,a}} \frac{\mathcal{J}_{i,a}}{1 - H(\widetilde{z}_{i,a})}. \tag{3}$$

3.4. Firm entry and subsequent survival

As was mentioned at the beginning of this section, a key feature of the model is the ability to replicate entry patterns observed in the data. The setup in this paper follows Sedláček and Sterk (2014) who treat entry and the selection of a particular technology type as an endogenous choice of potential startups rather than it being assigned to them based on a random draw from an exogenous distribution (as in e.g. Hopenhayn and Rogerson, 1993). As will become clear, this setup delivers realistic firm entry dynamics even when the only exogenous driving forces are shocks to aggregate productivity.

Starting up a firm requires the sacrifice of a (potentially time-varying) cost $X > 0$ which captures initial costs of doing market research, formulating a business plan etc. Upon paying this cost, a potential entrant chooses one business opportunity from a (fixed) finite measure of possibilities given by ψ_i. Each business opportunity allows for at most one successful startup.[12]

It is assumed that potential entrants cannot coordinate on which business opportunities to select. That is, not all individual opportunities are seized whereas others are pursued by several aspiring startups. This results in the number of startups within type i ($\omega_{i,0}$) being strictly smaller than both the number of business opportunities and the number of startup attempts (e_i). It follows that an attempted startup of a business type i is successful only with probability $\frac{\omega_{i,0,t}}{e_i}$. Unsuccessful startups exit before production takes place. This way of modeling firm entry is similar in spirit to models of innovation and research and development (see e.g. Klette and Kortum, 2004; Saint-Paul, 2002).

The coordination friction among aspiring startups is concisely summarized by an entry matching function, borrowed from the search and matching literature. This function relates the number of startups within each type to the respective number of startup attempts and business opportunities. It is assumed to be increasing in both arguments and to display constant returns to scale. In particular, $\omega_{i,0} = e_i^{\phi}\psi_i^{1-\phi}$, where $\phi \in (0, 1)$ is the elasticity with respect to the number of startup attempts.[13]

Free entry implies that in equilibrium the costs of starting up a firm of a specific tech-

[12] At a deeper level, the exclusivity of business opportunities could arise from patents claimed by individual firms. Alternatively, exclusivity could be generated by market size limitations coupled with fixed costs in production. For tractability, we do not model these factors explicitly.

[13] See Saint-Paul (2002) for a similar specification in the context of firms' research and development.

nology type equal the expected benefits

$$X = \frac{\omega_{i,0}}{e_i}\Pi_{i,0}, \text{ for } i = 1, 2, .., I, \tag{4}$$

The free entry conditions (4) imply that entry happens in all technology types. This is because technology types associated with high firm values will attract many potential new firms which in turn lowers the probability of successfully starting up. This, in turn, encourages entry of firms into technology types with lower firm values. In equilibrium, aspiring entrants are indifferent between all of the business opportunities, akin to models of directed search. Furthermore, the free entry conditions (4) show that the entry elasticity parameter ϕ essentially controls the sensitivity of the number of firms with respect to fluctuations in firm values. This enables the model to match firm entry patterns even with aggregate productivity being the only source of exogenous variation, a property that other firm dynamics models struggle with (see e.g. Samaniego, 2008).

Finally, firm exit is governed by an exogenous, but age-dependent probability δ_a. The evolution of the mass of firms of type i and age a is thus given by

$$\omega'_{i,a+1} = (1 - \delta_a)\omega_{i,a} \text{ for i=1,2,...,I and a} \geq 0. \tag{5}$$

3.5. Wage setting

The presence of a frictional labor market results in employment relationships being characterized by positive surplus values over which workers and firms bargain. The assumption of decreasing returns to scale in firms' production functions implies that these surplus values will depend on the number of workers employed in a given firm.

In this paper, wage setting is assumed to be conducted under the bargaining solution of Stole and Zwiebel (1996) which generalizes the Nash bargaining solution to a setting with decreasing returns to scale. Under this bargaining setup, the resulting wage is the same as under Nash bargaining, but where the bargaining happens over the *marginal* surplus. The resulting bargained wage takes on the following form

$$w_{i,a} = \eta \left(\frac{\alpha y_{i,a}/n_{i,a}}{1-\eta(1-\alpha)} + \frac{\kappa(\gamma-1)}{\gamma} x_{i,a}^{\gamma} + \Phi \right) + (1-\eta)b, \tag{6}$$

where $\Phi = \kappa \frac{V}{U} \sum_i \sum_a \frac{\omega_{i,a} v_{i,a}}{V} x_{i,a}^{\gamma-1}$.[14] The bargained wage has the familiar interpretation of being a weighted average of the marginal product, savings on hiring costs and the flow income in unemployment. Moreover, for linear hiring costs and constant returns to scale ($\gamma = 1$ and $\alpha = 1$) the above expression collapses to the standard Nash wage.

Before moving on, let us anticipate the quantitative properties of the model. As has been pointed out in other models with frictional labor markets and heterogeneous firms, heterogeneity per se does not necessarily lead to greater amplification of shocks (see e.g. Kaas and Kircher, 2014; Hawkins, 2011). Therefore, the model will be allowed to display a certain degree of wage rigidity, which is known to be one of the solutions to the labor market volatility puzzle. In particular, wages in individual firms will be a weighted average of the Stole-Zwiebel wage derived in (6) and its steady state counterpart. The weight (ζ) given to the steady state wage then governs the degree wage rigidity in the model.[15]

3.6. Aggregate shocks, market clearing and equilibrium

There are two aggregate shocks present in the model. The first is a shock to aggregate productivity (A) and the second is a shock to the entry cost (X). Time variation in entry costs is meant to capture driving forces particularly affecting startups, e.g. changes in the tightness of credit conditions, house prices or uncertainty.[16] The latter is meant to capture driving forces specific to startups, e.g. tightened financial conditions for new firms as in

[14] The derivation of the wage can be found in the Appendix.

[15] In the quantitative exercises the resulting wages lie within the bargaining sets of workers and firms.

[16] For examples of studies linking these factors to startup activity see Drautzburg (2014), Siemer (2014) and Schmalz, Sraer, and Thesmar (forthcoming).

Siemer (2014). Both are assumed to follow an AR(1) process in logs

$$\ln A' = \rho_A \ln A + \eta^A, \quad \eta^A \sim N(0, \sigma_A^2), \tag{7}$$

$$\ln X' = (1-\rho_X)\overline{X} + \rho_X \ln X + \eta^X, \quad \eta^X \sim N(0, \sigma_X^2), \tag{8}$$

where $e^{\overline{X}}$ is the steady state entry cost, ρ_j is the autocorrelation coefficient and η_t^j is the shock innovation assumed to be distributed identically and independently according to a normal distribution with zero mean and standard deviation σ_j, with $j = A, X$.

Given all the above the aggregate resource constraint can be written as

$$Y + bU = C + \sum_i \sum_a \omega_{i,a} \frac{\kappa}{\gamma} x_{i,a}^{\gamma} n_{i,a} + X \sum_i e_i, \tag{9}$$

where aggregate firm output $Y = \sum_i \sum_a \omega_{i,a} y_{i,a}$ together with home production are spent on consumption, the costs of posting vacancies and the cost of starting up firms. The elements of the household's budget constraint left to define are total wage income $W = \sum_i \sum_a \omega_{i,a} w_{i,a} n_{i,a}$ and aggregate profits $P = \sum_i \sum_a \omega_{i,a}(y_{i,a} - w_{i,a} n_{i,a} - \frac{\kappa}{\gamma} x_{i,a}^{\gamma} n_{i,a}) - X \sum_i e_i$. Finally, the number of unemployed is given by $U = L - \sum_i \sum_a \omega_{i,a} n_{i,a}$ with L being the size of the labor force.

Let $\mathcal{S} = \{A, X, \omega_{i,a}, \widetilde{n}_{i,a}\}_{i=1,2,\ldots,I, a \in \mathbb{N}}$ be the aggregate state consisting of both aggregate shocks, but also of the entire firm distribution (the mass and beginning-of-period employment levels of all firms). The entire firm distribution is a state variable because individual firms need to know and be able to predict the evolution of the matching probabilities on the frictional labor market in order to make their optimal decisions. The matching probabilities, however, depend on aggregate vacancies and unemployment which are in turn determined by the sum of employment levels and vacancies at individual firms.[17] The following lines define the equilibrium of the economy.

Definition 1. *A recursive equilibrium is defined by*

[17] Posted vacancies in turn depend only on firm type (productivity), age and size.

- *individual firm policy rules for next periods' available employment* $\widetilde{n}_{i,a}(S)$*, vacancy posting* $x_{i,a}(S)$*, firm value* $\Pi_{i,a}(S)$*,*
- *the representative household's consumption choice* $C(S)$*,*
- *wages* $w_{i,a}(S)$*,*
- *a measure of firm startups* $\omega_{i,0}(S)$ *for each firm type* $i = 1, 2..., I$*,*

that solve the household's problem, clear the labor market, satisfy the entry condition, satisfy the aggregate resource constraint, solve the firm's problem and are consistent with the evolution of individual firm masses and employment levels for all firm ages $a \in \mathbb{N}$ *and types* $i = 1, 2, ..., I$ *and are consistent with the evolution of the aggregate productivity and entry cost shocks.*

4. Quantitative implementation

This section first describes the adopted solution method and the parametrization procedure. Before moving on to the results, the end of this section is devoted to inspecting the model's predictions along cross-sectional and business cycle dimensions and comparing them to the observed patterns in the data.

4.1. Solution method

As is clear from the previous section, the entire firm distribution is a state variable. Moreover, in the presence of aggregate uncertainty, this distribution is time-varying. A popular solution was proposed by Krusell and Smith (1998) who assume that agents forecast only a limited number of moments of the true distribution. The disadvantage of such a procedure is that it typically requires many simulation steps making it relatively time-consuming.

This paper uses the solution method of Sedláček and Sterk (2014) which is based on a truncation of firm age at a certain maximum value K and which uses the fact that all firms of the same type and age make the same choices. This reduces the number of state variables to a finite number: two aggregate shocks and $2 \times I \times K$ firm-specific states describing the firm distribution (the mass and beginning-of-period employment for each type and age). This approach therefore enables the exact tracking of the *entire* distribution (given the assumed maximum firm age), rather than approximating it with simulation.[18] Policy rules are then obtained using first-order perturbation for each firm-age type, i.e. along the stationary growth-path of firms of all types. Further details on the exact numerical implementation are given in the Appendix.

4.2. Parametrization

All model parameters, except for those pertaining to the two aggregate shocks, are calibrated using aggregate and BDS data for the available sample period 1977-2012. The parameters, and time-paths, of the two structural shocks are then *estimated* using Maximum Likelihood on data for the unemployment rate and the number of startups.[19] To facilitate the exposition, let us first discuss the parametrization strategy as a whole, then move on to parameters found in standard search and matching models and finally also parameters pertaining to firm dynamics. Table 1 summarizes all the parameter values, their targets and the respective model predictions. In the benchmark specification of the model, the number of firm types is set to $I = 5$ and the maximum firm age is set to $K = 100$.[20]

The parametrization strategy is based on the interpretation of the entry cost shock as "residual" variation in startups not driven by fluctuations in firm values. In other words, it is meant to capture forces particular to startups and not directly affecting incumbent firms. The entrant elasticity parameter (ϕ) plays a crucial role in this respect. In particular, using the firm entry conditions (4) it can be shown that $\phi/(1-\phi)$ is the elasticity of the number of

[18]This faster methodology therefore allows for the estimation of the model which requires it to be solved many times.

[19]Both time-series are linearly detrended in order not to impose artificial autocorrelations into the estimation. Nevertheless, the results are similar when estimated on HP-filtered data. While the unemployment rate is available at a quarterly frequency, the BDS data offers only annual information. This, however, can be conveniently handled by the Kalman filter used to construct the likelihood function. More details on the estimation procedure are in the Appendix.

[20]The Appendix shows that increasing the maximum firm age does little to the results.

entrants with respect to changes in firm values.[21] Therefore, ϕ is chosen such that the model generates a relative volatility of startups with respect to output of 2.5 as in the BDS data *without* exogenous variation in entry costs.[22] The estimation procedure will then assign all other variation in startups to the entry cost shock.

The model period is assumed to be one quarter and the discount factor (β) is therefore set to 0.99 implying an annual interest rate of 4 percent. The size of the labor force (L) is set such that the resulting steady state unemployment rate is equal to 6.3 percent. The level of match efficiency (m) is set to target a job filling probability of 0.71 as in den Haan, Ramey, and Watson (2000). The matching elasticity (μ) takes on the value of 0.6, which is in the middle of values found in the literature (see Petrongolo and Pissarides, 2001, for an overview). The value of home production (b) is set to match the overall separation rate of 4.6 percent, taken from the Current Population Survey. The intuition for this target is that for a given parametrization of worker-specific productivity, a higher value of b reduces the surplus of the employment relationship implying a greater chance of separation.[23] The resulting value gives rise to an average replacement rate of 66 percent. Following most of the existing literature, it is assumed that workers and firms have equal bargaining power, i.e. $\eta = 0.5$. The distribution of worker-specific productivity shocks is assumed to be logistic with mean 1 and scaling parameter σ_H which is set such that the volatility of separations relative to output volatility is 3.4 as in the data. Given all the model parameters, the costs of posting vacancies are determined via the optimal hiring condition. The resulting value implies overall costs of less than 1 percent of output. Last, as anticipated already in Subsection 3.5, the model requires an additional (exogenous) source of wage rigidity in order to be able to match the volatility of labor market variables. As explained earlier, wages are a weighted average of the Stole-Zwiebel bargained wage and the respective steady state

[21]Specifically, using the firm entry condition we can write $d \ln \omega_{i,a} = \frac{\phi}{1-\phi}(d \ln \Pi_{i,0} - d \ln X)$.

[22]In a similar fashion, the calibration of parameters related to the volatility of separations and wages (discussed below) is also based on a model without entry cost shock variation.

[23]This way of calibrating the payoff from unemployment was also suggested by Elsby and Michaels (2013).

counterpart. The weight (ζ) is chosen such that the resulting wage elasticity with respect to labor productivity is 0.5 as in Hagedorn and Manovskii (2008).

Let us now discuss firm dynamics parameters which are common to all firms. The returns-to-scale parameter is set to 0.85 as in Schaal (2012). This value lies in the middle of the estimates in Basu and Fernald (1995); Basu (1996); Basu and Kimball (1997). Vacancy posting costs are assumed to be quadratic in the vacancy rate, i.e. $\gamma = 2$, as in Gertler and Trigari (2009).[24] In order to capture the age-dependent character of firm exit, it is assumed that the exit probabilities are given by $\delta_a = \overline{\delta}_0 + \overline{\delta}_1 / a$ for $a < K$ and $\delta_K = 1$. The coefficients $\overline{\delta}_0$ and $\overline{\delta}_1$ are then chosen to match the observed exit rates in the BDS data, conditional on firm age. All firms are assumed to start with an employment level of n_0 which is set to match average firm size of startups.[25] Finally, the mass of entrants is normalized to 1. Given all the above choices, the level of the entry cost does not affect any variables in the model, except for the number of startup attempts, or equivalently the startup probabilities. Interpreting the startup probability as the survival rate during the first period, the entry cost is set such that the model matches a probability of success of 34 percent implied by the calibrated firm exit function.[26]

Finally, turning to firm-specific parameters, the productivity levels (ϵ_i) and associated masses of business opportunities (ψ_i) are pinned down by requiring the model to have a realistic firm size distribution. Specifically, the productivity parameters are set such that the model matches the employment shares within five size brackets found in the BDS data: $1-49$, $50-249$, $250-999$, $1,000-9,999$ and $> 10,000$.[27] Similarly, the firm mass parameters ψ_i are chosen to match the firm shares within the above five size brackets.

[24] The Appendix investigates the sensitivity of the results to alternative values of γ.

[25] This helps the model to match the employment shares by age but otherwise changes very little.

[26] In particular, assuming that the survival rate in the first period is constant, one can use the calibrated δ_a function and evaluate it at age 0.5 which delivers a survival rate of 0.34.

[27] The level of TFP of the least productive firm type is normalized to 1.

Table 1: Calibrated parameters

	parameter	value	target/source	model
β	discount factor	0.99	4% annual interest rate	
L	labor force	822.7	$u = 6.3\%$, BLS	6.3%
m	matching efficiency	0.631	$Q = 71\%$, den Haan, Ramey, and Watson (2000)	71%
μ	match elasticity	0.600	Petrongolo and Pissarides (2001)	
b	value of non-work activities	0.504	$\rho_T = 4.6\%$, CPS	4.6%
η	worker bargaining power	0.500	symmetric bargaining	
κ	vacancy posting cost level	1.75	optimal hiring condition	
ζ	wage rigidity parameter	0.87	$\epsilon_{w,z} = 0.5$, Hagedorn and Manovskii (2008)	0.5
μ_H	idio-syncartic productivity shock, mean	1.000	normalization	
σ_H	idio-syncratic productivity shock, shape parameter	0.220	$\sigma(ln(\rho_T))/\sigma(\ln(Y)) = 3.3$, CPS	3.1
α	returns-to-scale parameter	0.850	Schaal (2012)	
γ	convexity of hiring costs	2.000	Gertler and Trigari (2009)	
$\overline{\delta}_0$	firm exit level parameter	0.005	exit rate age profile, BDS	
$\overline{\delta}_1$	firm exit age parameter	0.330	exit rate age profile, BDS	
n_0	initial employment stock	6.000	average entrant size of 6.2, BDS	6.3
$\overline{X}$	firm entry cost level	2.351	success probability 34%, BDS approximation	34%
ϕ	entry elasticity	0.690	$\sigma(\ln(\sum_i \omega_{i,0}))/\sigma(\ln(Y)) = 2.5$, BDS	2.3
σ_Z	standard deviation, aggregate productivity	0.008 (0.001)	estimated	
ρ_Z	AR coefficient, aggregate productivity	0.925 (0.021)	estimated	
σ_Z	standard deviation, startup cost	0.019 (0.003)	estimated	
ρ_Z	AR coefficient, startup cost	0.864 (0.058)	estimated	
ϵ_i	firm productivity	1	1.47 1.82 2.23	2.95
$100\frac{\omega_{i,0}}{\sum_i \omega_{i,0}}$	fraction of startup types (%)	92.7	5.9 0.9 0.4	0.1

Notes: The table shows the calibrated parameters (based on a model without entry cost shocks), their respective targets or sources and the last column shows the value of the statistic as predicted by the calibrated model. ρ_T is the total separation rate, Y is aggregate output, $\overline{w}$ denotes the average wage, $\epsilon_{w,z}$ is the elasticity of wages with respect to productivity and u is the unemployment rate.

Table 2: Firm age and size distributions

shares (in %)	new	young	old	1–49	50–249	250–999	1,000–9,999	> 10,000
				BDS data				
firm	10.6	31.7	57.7	95.04	4.15	0.57	0.22	0.02
employment	3.0	13.2	83.8	37.8	21.1	12.5	17.4	11.2
				model				
firm	11.3	31.0	57.7	95.10	4.05	0.58	0.24	0.03
employment	3.0	15.5	81.5	35.9	19.3	11.6	20.8	12.4

Notes: The table reports employment and firm shares within firm age (first 3 columns) and size (last five columns) bins. "New" firms are less than one year old, "young" firms are between 1 and 5 years of age, and "old" firms are defined as more than five years of age. All values are in percent.

4.3. Cross-sectional and business cycle properties

Before moving on the the main results of the paper, this subsection documents that the presented structural model is consistent with the data along several cross-sectional and business cycle dimensions.

Firm size and age distributions. The previous subsection described that the firm-specific parameters of the model are set to match the firm size distribution observed in the BDS data. Table 2 shows these targets, together with firm and employment shares of new, young and old firms. The table documents that the structural model does well not only in capturing the firm size distribution but it is also consistent with the empirical firm age distribution.

Net job creation in young and old firms. Because the model will be used to evaluate the labor market implications of a lost generation of (young) firms, it is important that the model replicates net job creation rates conditional on age in the cross-section. Figure 4 shows that the model predicts a negative relation between firm age and the net job creation rate

Figure 4: Net job creation by age

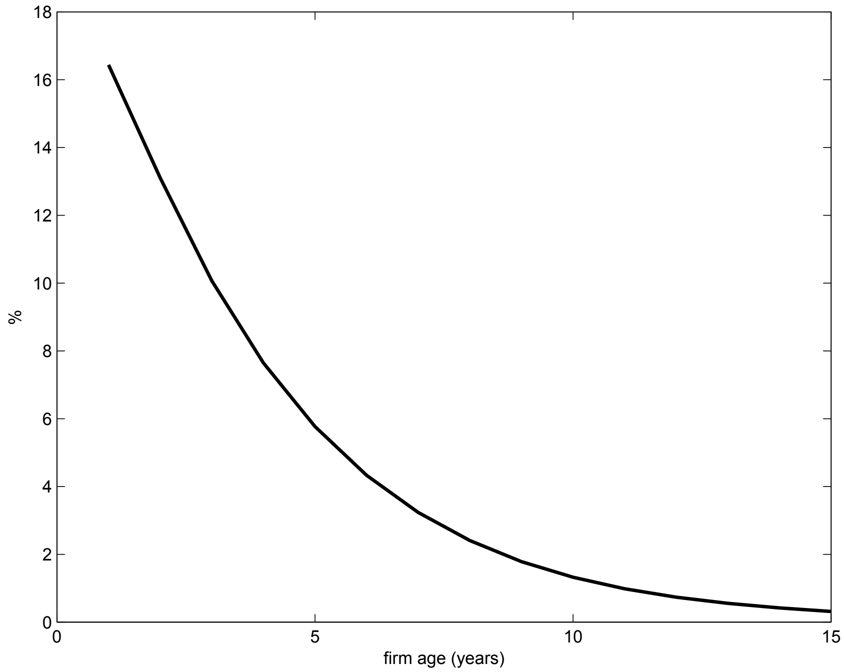

Notes: Model-predicted net job creation of continuing firms by age.

of continuing firms.[28] Moreover, quantitatively the growth rates are also close to their data counterparts. Haltiwanger, Jarmin, and Miranda (2013) report growth rates of one year old firms of about 15% which drop relatively quickly with age and stay below 5% after the age of about four.

Time-paths of variables not in estimation. Finally, before moving on to the results it is worth inspecting how the model performs in terms of capturing dynamics of variables not directly used in the estimation. Figure 5 shows the actual and model-predicted time-paths for several variables. Overall the model does well in capturing the dynamics of both firm-level and aggregate variables.

Since the focal point of the paper is the relative (employment) behavior of startups and incumbent firms it is crucial that the model performs well in this dimension. The top row of

[28] Following Haltiwanger, Jarmin, and Miranda (2013), the net job creation rate is defined as $g_{i,a} = (n_{i,a,t} - n_{i,a-1,t-1})/(0.5(n_{i,a,t} + n_{i,a-1,t-1}))$.

Figure 5: Actual and model-predicted time-paths

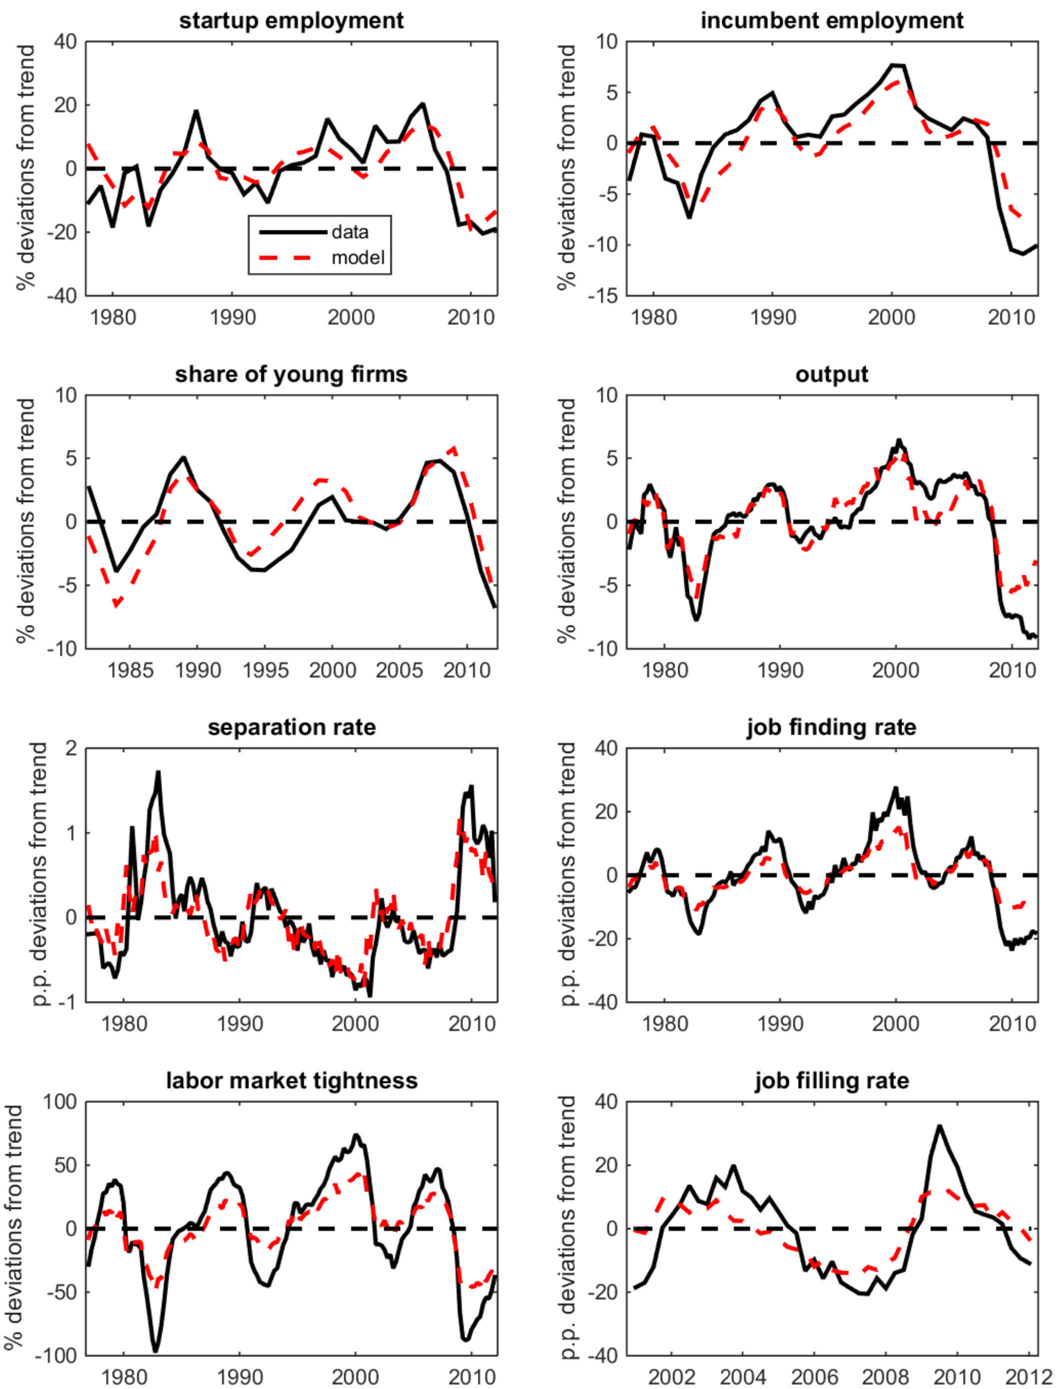

Notes: actual and model-predicted time-paths of variables (based on an estimation using the unemployment rate and the number of entrants as data inputs). The employment in new and incumbent firms as well as the share of young firms is based on BDS data, the separation and job finding rates are taken from the CPS, labor market tightness is defined as vacancies (using the time series constructed by Barnichon (2010)) over unemployment and finally the job filling rate is the vacancy yield from Job Openings and Labor Turnover Survey. All data are linearly detrended.

Figure 5 shows that the model does well in capturing the employment dynamics of new and incumbent firms. While both time-series are slightly less volatile in the model compared to the data, the model-predicted volatility of employment in new relative to incumbent firms is virtually identical to that in the data (2.35 in the data compared to 2.30 in the model).

At the aggregate level, labor market tightness is slightly less volatile than in the data and this spills over to the job finding and filling rates. This is due to a lower volatility of vacancies predicted by the model. The reason is that a large fraction of vacancies is posted by young firms which are characterized by being far away from their optimal size, and thus having relatively large marginal products of labor. This in turn, increases their surplus resulting in relatively lower sensitivity to aggregate shocks.[29]

Finally, it is worth noting that the strong Beveridge curve relation in the data is preserved by the model. While in the data the correlation between (cyclical components of) unemployment and vacancies is -0.87, the model predicts a correlation of -0.91.

5. Aggregate dynamics and lost generations of firms

The purpose of this section is to understand the quantitative effects of a lost generation of firms within a general equilibrium model in which existing firms are free to adjust their hiring behavior in response to the lack of jobs created by young firms. Towards this end, the next subsection presents a set of impulse response functions (IRFs) to a one-time shock to the entry cost which results in a 30 percent drop in the number of entrants (the magnitude observed in the Great Recession).

Then, the structural model is used to *estimate* the time-path of the two aggregate shocks (productivity and entry cost) during and in the aftermath of the Great Recession. This enables the model to quantify the relative contribution of a strong but "standard" recession

[29] This contrasts the results of Elsby and Michaels (2013) who find that a labor market model with heterogeneous firms does well in replicating the volatility of labor market variables (even without exogenous wage rigidity). The key difference is that in their framework there is no notion of a firm's life-cycle (i.e. they abstract from firm entry and exit) and thus all firms (including small ones) are close to their optimal sizes.

Figure 6: Impulse response of the firm-age distribution

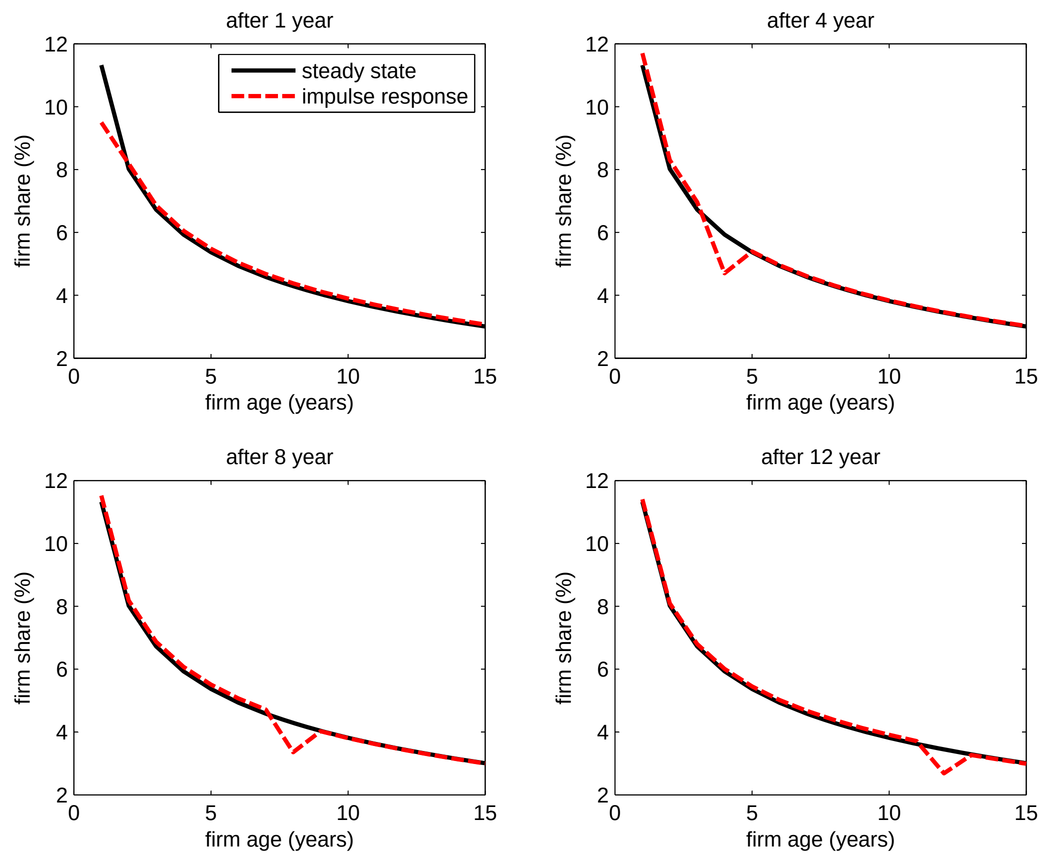

Notes: the steady state firm-age distribution (up until the age of 15 for ease of exposition) and the impulse response of the distribution 1, 4, 8 and 12 years after the shock to the entry cost subsides. The magnitude of the entry cost shock is such that the resulting drop in the number of startups resembles the one observed during the Great Recession.

shock and a shock that specifically affected entrants. Moreover, it is possible to use the model and the estimated shocks to conduct counterfactual scenarios and ask how the economy would have developed had the number of startups not fallen as much.

5.1. Impulse response functions

Firm-age distribution. Let us start by visualizing how an exogenous increase in the startup cost resulting in a drop in the number of entrants reverberates through the firm-age distribution. The magnitude of the shock to the entry cost is such that the resulting drop in the number of startups resembles that observed during the Great Recession. Figure 6 shows the steady state firm-age distribution and the impulse response of the distribution 1, 4, 8

Figure 7: Impulse response functions of aggregate variables

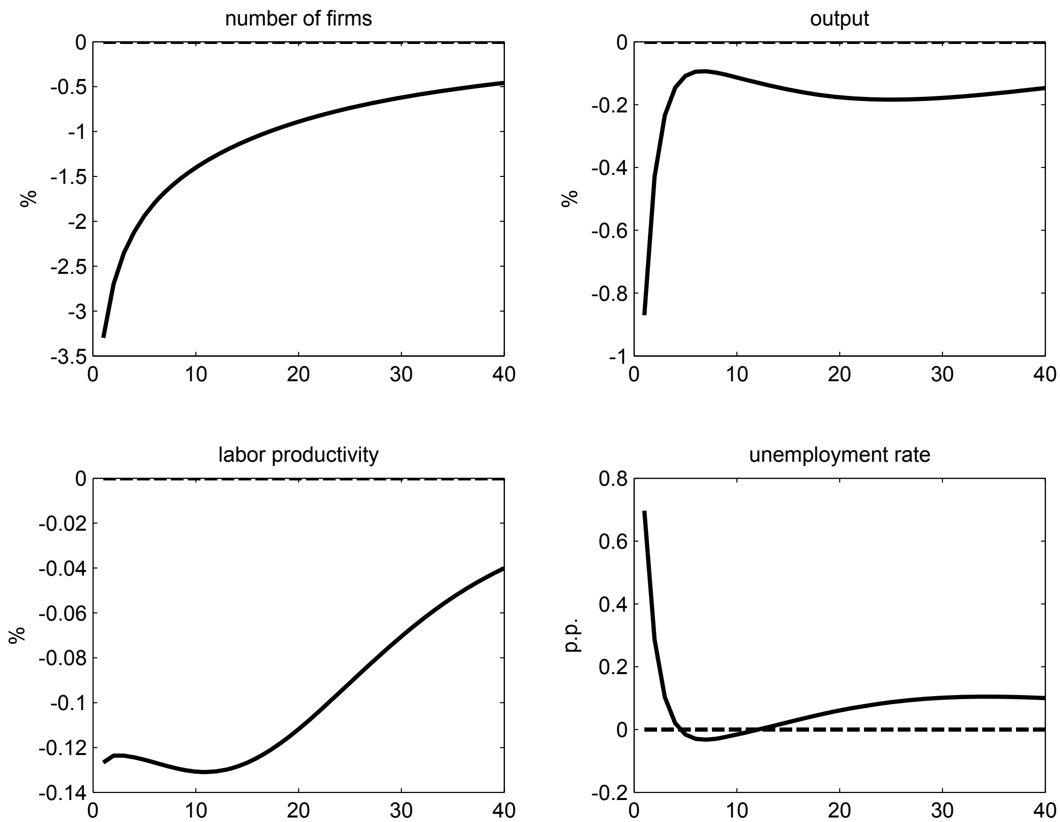

Notes: impulse response functions of the number of firms, output, labor productivity and the unemployment rate to a one-time shock to the entry cost. The magnitude of the entry cost shock is such that the resulting drop in the number of startups resembles the one observed during the Great Recession.

and 12 years after the shock subsides. The figure makes clear that even a one-time shock to entry may affect the economy many years after it subsides as the affected cohort of firms grows old and works its way through the firm-age distribution.

Output, productivity and unemployment. Figure 7 plots the IRFs (over a period of ten years) of the number of firms, output, labor productivity and unemployment to the one-time shock to the entry cost. The top left panel depicts the response of the number of firms which remains persistently below its steady state level owing to the fact that the affected cohort of firms dies out only slowly over time. The top right panel shows the response of output which after an initial almost 1 percent drop remains to be 0.2 percent below its steady state even ten years after the initial shock subsides. The output response stems both from a persistent decline in labor productivity (bottom left panel) and a very persistent increase in the unemployment rate. The reason for the former is that the economy shifts away from young firms which are more productive than older businesses (because of the decreasing returns to scale).

While the impact of the lost generation of firms is as persistent as in the simple counterfactual exercise in Section 2, the magnitude of the effect is substantially dampened. The reason for this stark difference is the presence of several general equilibrium effects present in the structural model. The following paragraphs inspect these general equilibrium (GE) effects in more detail.

General equilibrium effects. After an entry cost shock, the number of entrants falls and this reverberates through the firm-age distribution as a lost generation of firms (Figure 6). This brings with it a shift in the composition of firms away from young businesses lowering the number of posted vacancies and thus creating a tighter labor market ($\theta = V/U$ falls). This effect serves to dampen the initial increase in the unemployment rate through four distinct channels depicted in Figure 8.

First, because of the tighter labor market workers' outside options worsen and they settle for lower wages (top left panel).[30] This by itself increases the incentives for existing firms to hire more workers.

Second, because of the fall in wages, firm profits increase. The higher profitability increases the incentives to start up new businesses. The top right panel shows that after the initial (exogenous) drop in firm entry, the number of startups immediately jumps back and remains persistently *above* its steady state level. The higher number of startups in subsequent years helps offset some of the initial lack of job creation by young firms.

Third, while the tight labor market makes it harder for the unemployed to find jobs,

[30]The drop in the average wage also reflects the shift away from younger more productive firms (because of the decreasing returns to scale).

Figure 8: Impulse response functions: general equilibrium effects

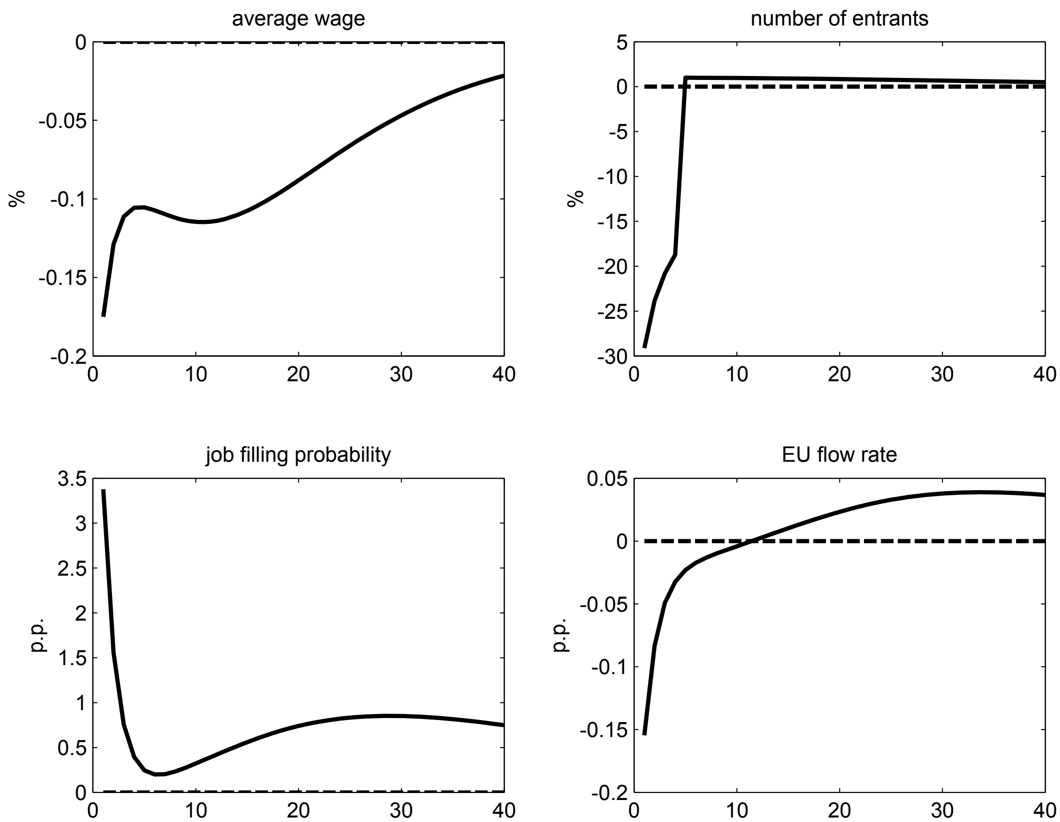

Notes: impulse response functions of the number of entrants, average wages, the probability of hiring a worker and the separation rate to a one-time shock to the entry cost. The magnitude of the entry cost shock is such that the resulting drop in the number of startups resembles the one observed during the Great Recession.

it makes it relatively easier for incumbent firms to find workers (bottom left panel). This again increases hiring incentives as the effective costs of posting vacancies (which include the expected duration of a vacant position) fall.

Fourth, it is easier to retain workers because their outside option (of finding a new job) worsens. This channel resembles the "insulation effect" of lower job creation put forward in Caballero and Hammour (1994). This insulation effect is, however, quantitatively relatively weak. The majority of the resulting separation rate decline (bottom right panel) is because of a composition shift away from young firms which have on average higher rates of shutting down.[31] All the above channels serve to dampen the increase in the unemployment rate

[31]The subsequent increase in the separation rate is also driven by composition changes as the share of

following the lack of job creation by startups (and subsequently young firms).

Decomposing the unemployment rate response. To quantify the relative strength of the above channels, it is possible to construct a counterfactual unemployment rate response where the general equilibrium effects are "shut down". The solid black lines in Figure 9 indicate, respectively, the benchmark IRF ("flexible") and the IRF without the GE effects ("fixed"). Without the GE effects unemployment is almost 1.5 percentage points above its steady state even ten years after the shock subsides. This is reassuringly similar to the unemployment response in the simple counterfactual (which also does not allow for GE effects) presented in Section 2.

The remaining lines characterize the relative strengths of the four general equilibrium channels. Specifically, "flex entry" allows for firm entry to vary (i.e. to overshoot following the exogenous drop), "flex entry & wages" lets not only firm entry but also wages to respond endogenously and finally "flex entry, wages & Q" assumes only a fixed separation rate while the remaining three channels are left to endogenously adjust.

Figure 9 therefore quantifies to what extent each channel contributes to the considerably milder response of unemployment compared to the simple counterfactual which does not allow for GE effects. The closer the given counterfactual response is to the "flexible" response, the stronger the given general equilibrium channel. Therefore, the figure suggests that while the endogenous response of firm entry and wages (red dashed line) is an important channel that mutes the increase in the unemployment rate, the GE effects operating through the frictional labor market are equally strong. The dominant channel in the labor market is the changing probability of hiring workers.

young firms rises above its steady state after about 3 years. Recall that firm entry overshoots from the second period onwards.

Figure 9: Unemployment rate response: benchmark and counterfactuals

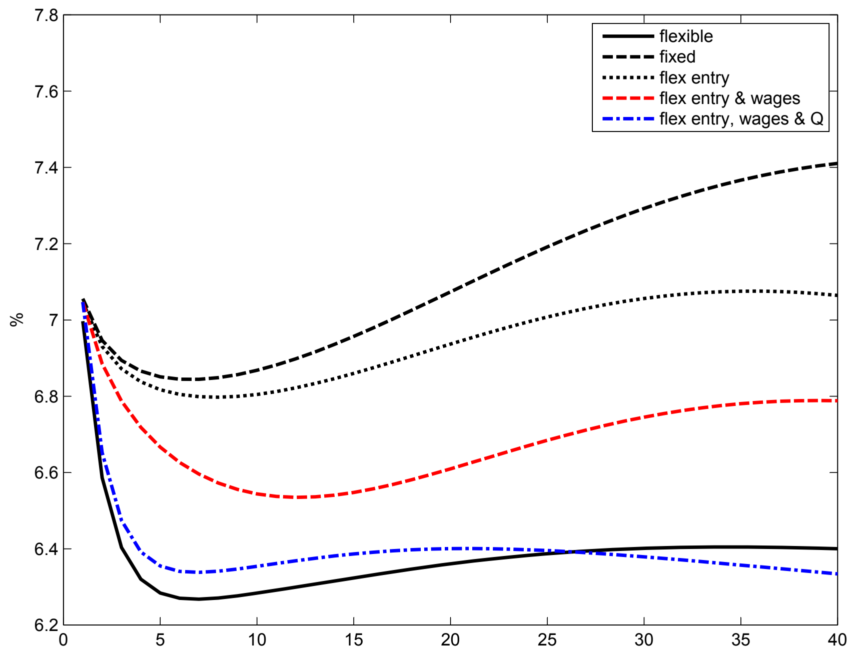

Notes: impulse response of the unemployment rate in the benchmark specification ("flexible") and four counterfactuals. "fixed" refers to the case when the job filling probability, wages and separation rates are fixed and when entry reverts back to steady state after the initial shock. "flex entry" is like the "fixed" case, but entry responds endogenously. "flex entry & wages" is like the "flex entry" case, but in addition firms' wages respond endogenously. Finally, "flex entry, wages & Q" is like the "flex entry & wages" case, but in addition the job filling probability responds endogenously. The magnitude of the entry cost shock is such that the resulting drop in the number of startups resembles the one observed during the Great Recession.

5.2. The lost generation of firms in the Great Recession

This subsection uses the estimated time-paths of the model variables to inspect the relative contribution of the two structural shocks during the Great Recession and to conduct counterfactual scenarios determining how the U.S. economy would have evolved had it not been for the strong decline in firm entry in the past years.

First, let us discuss to what extent the drop in firm entry was caused by forces specific to startups and to what extent it was a reaction to a "standard" but large recession. Towards this end, I simulate the benchmark economy using the estimated aggregate productivity shock alone, i.e. without exogenous variation in the entry cost. In this case, firm entry would have still been about 10 percent below its trend in 2010. This means that roughly 60 percent of the subdued entry during the Great Recession was driven by a "standard"

Figure 10: Time-paths of variables under alternative scenarios

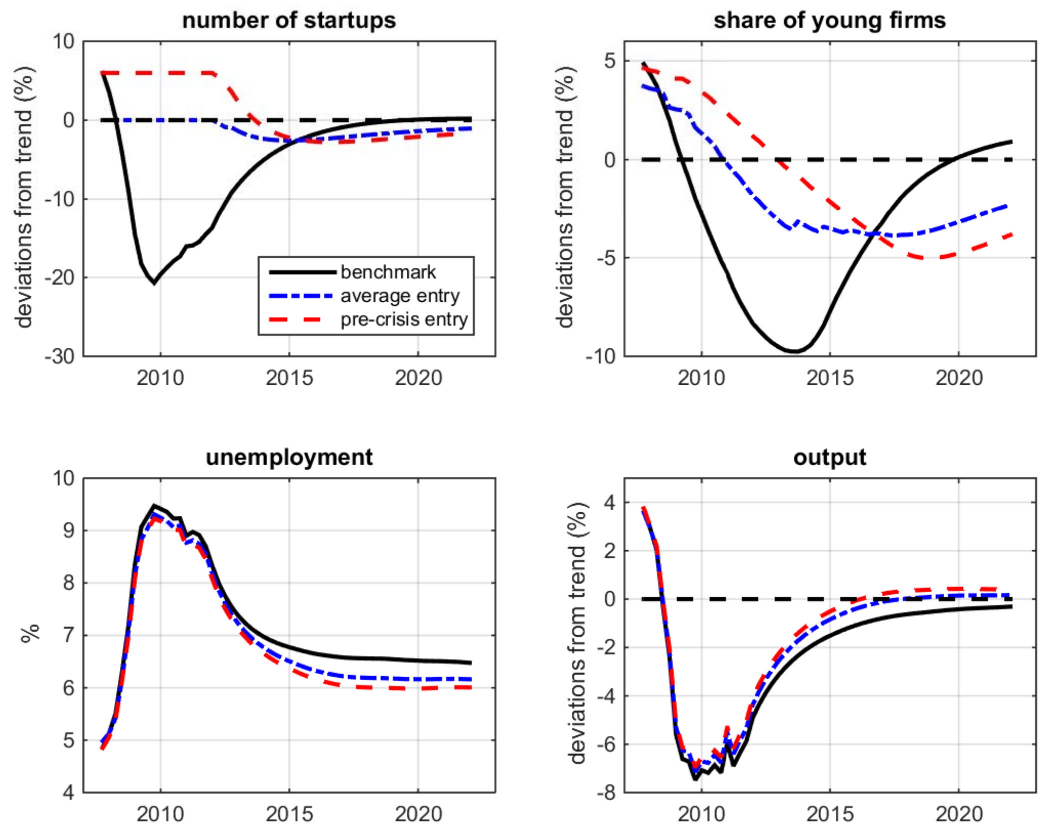

Notes: actual and couterfactual time-paths of the number of startups, the share of young firms, output and the unemployment rate using estimated shocks to aggregate productivity and the entry cost. "Average entry" ("pre-crisis entr") refers to the case when the entry costs are such that the model generates firm entry equal to its trend (pre-crisis) level during and in the aftermath of the Great Recession. The data used in the estimation (unemployment rate and the number of entrants) runs between 1977Q1 and 2012Q1. Thereafter the model is left to converge to its steady state (i.e. both aggregate shocks are set to zero).

response to a particularly strong recession while the rest can be attributed to factors specific to startups.

Next, let us consider two counterfactual scenarios: one in which the entry cost shocks are such that firm entry during and in the aftermath of the Great Recession remains at its trend level ("average entry") and one in which it remains at its pre-crisis level ("pre-crisis entry"). After the end of the sample period, the economies are left to converge to the steady state for 10 years (i.e. both shocks are fixed to their unconditional means until 2022Q1). Figure 10 plots the actual and counterfactual time-paths for the number of startups, the share of young firms, unemployment and output.

Both counterfactual scenarios have similar quantitative effects. The immediate consequences of a drop in firm entry on the aggregate economy are small as the general equilibrium effects allow incumbent firms to compensate for the loss of job creation by startups. Specifically, the unemployment rate peak would have been only about 0.2 percentage points lower.

However, in later years the negative effect of the lost generation of firms strengthens. The reason is that the missing entrants generate fewer older firms in the future (apparent from Figure 6) which account for the bulk of aggregate employment. Even though firm entry recovers somewhat faster in the benchmark economy compared to the counterfactuals (as explained in Figure 8) it is not enough to compensate for the depressed employment in mature firms. The reason is that the employment gains of the additional startups kick in only in later years as these new firms grow older and larger.

Therefore, the missing generation effect creates a very persistent dent in the employment potential of the economy essentially raising the "natural" rate of unemployment. Specifically, had the number of startups remained constant during and in the aftermath of the Great Recession then the unemployment rate would have been 0.5 percentage points lower even 10 years after the crisis. Similarly, output would have reverted back to trend 4-6 years earlier compared to the benchmark economy.

5.3. Discussion

This paper asks whether periods of subdued firm entry can have long-lasting negative effects on the aggregate labor market. The results suggest that it can but, as long as incumbent firms play a similar role as new businesses (in this case job creation), the short-run aggregate impact of a lost generation of firms is limited. However, one can imagine other channels which make young firms specific in certain dimensions and which may influence the impact firm entry has on the aggregate economy. This subsection briefly discusses several such channels hinting at possible directions for future research.

Composition of startups. Hurst and Pugsley (2011) document that many potential startups have little or no ambitions to grow or innovate. One may think that the composition of firms with respect to their potential to grow large changes over the business cycle. The direction of these compositional changes is, however, unclear. On the one hand, recessions may be times when only the relatively more productive firms can enter the economy. On the other hand, recessions may give rise to "necessity" entrepreneurs who start businesses as a means of escaping unemployment without ambitions (or the ability) to grow large.

Sedláček and Sterk (2014) investigate the potential of startups to grow large over the business cycle in the BDS data. They find that recessionary periods give rise to firms that are smaller and that remain small even several years into the future. Using an estimated heterogeneous firm model they conclude that changes in the composition of startups (with respect to their potential to grow large) are extremely important for cohort-level employment variation, but they also help shape the medium- to long-run fluctuations in aggregate employment. This suggests that while changes in the number of new firms alone are unlikely to have persistent effects, a greater concern may lie in the selection effects among startups over the business cycle.

Mismatch. In the current model all workers are ex-ante identical and therefore young and old firms play similar roles in job creation. In reality, there are many dimensions along which workers differ ex-ante (e.g. education, sector of employment, age, experience etc.). If new (young) firms disproportionately employ a particular type of workers, then the general equilibrium effects mentioned in this paper may be weakened because older firms would not be tempted to hire from the larger unemployment pool.[32] Even though such increased mismatch could strengthen the impact a lost generation of firms has on the aggregate labor market, current evidence suggests only a relatively modest role for mismatch in explaining

[32] Ouimet and Zarutskie (2013) show that young firms disproportionately employ (and hire) young workers even when controlling for several observables.

unemployment dynamics (see e.g. Sahin, Song, Topa, and Violante, 2014).[33]

Postponing of entry. It is possible that firms that do not enter during recessions are simply waiting until business conditions improve and they will enter in the subsequent recovery phase. In this sense startups are not lost, they are merely postponed. Interestingly, the presented model does indeed predict that an exogenous drop in firm entry is followed by an overshooting in the following periods. The magnitude of this effect is, however, rather weak. Also in the data there seems to be little evidence in support of a strong overshooting of startups following periods of subdued entry. Even in 2012, three years after the official end of the recession, the number of startups was still 28 percent below its pre-crisis level (and 16 percent below the sample average).

Labor force participation. The Great Recession was a time when the participation rate fell strongly (arguably contributing to some of the unemployment rate decline). It is beyond the scope of this paper to investigate the quantitative impact this additional margin of adjustment may have in the medium-run. However, given recent evidence that the participation margin can play an important role for unemployment rate fluctuations (see Elsby, Hobijn, and Sahin, forthcoming), incorporating this feature seems an interesting avenue for future research.

6. Conclusion

This paper documented that the firm-age distribution exhibits cyclical changes (shifts away from young firms during recessions) which help shape aggregate employment dynamics. A simple exercise would suggest that the lost generation of firms observed during the Great Recession may have substantial negative effects on the aggregate labor market for many

[33] Sedláček (2014), however, shows that the contribution of mismatch to unemployment rate fluctuations is substantially larger when job seekers from outside unemployment are taken into account.

years to come.

However, this conclusion is based on unchanged behavior of incumbent firms. A general equilibrium model of firm dynamics and a frictional labor market shows that accounting for feedback effects of firm entry into the employment behavior of incumbent firms is important for the magnitude of the aggregate response to the drop in firm entry. However, despite the strong general equilibrium effects, periods of subdued firm entry remain to negatively impact the economy in the medium- to long-run.

What is beyond the scope of this paper are statements about the efficiency and therefore policy recommendations. Given the ex-ante heterogeneity (and the imposed wage rigidity), it is unlikely that the simple Hosios condition would make the competitive equilibrium efficient (see e.g. Hawkins, 2014). Therefore an analysis of the efficiency conditions, together with the additional channels pointed out in the discussion seem to be interesting directions for future research.

References

Barnichon, R. (2010): "Building a Composite Help-Wanted Index," *Economics Letters*, 109(3), 175–178.

Basu, S. (1996): "Procyclical Productivity: Increasing Returns to Scale or Cyclical Utilization?," *Quarterly Journal of Economics*, 111(3), 719–751.

Basu, S., and J. Fernald (1995): "Are Apparent Productive Spillovers a Figment of Specification Error?," *Journal of Monetary Economics*, 36(1), 165–188.

Basu, S., and M. Kimball (1997): "Cyclical Productivity with Unobserved Input Variation," NBER Working Paper No. 5915.

Caballero, R., and M. Hammour (1994): "The Cleansing Effect of Recessions," *American Economic Review*, 84(5), 1350–1368.

Campbell, J. (1999): "Entry, Exit, Embodied Technology, and Business Cycles," *Review of Economic Dynamics*, 1(2), 371–408.

Clementi, G. L., and D. Palazzo (2013): "Entry, Exit, Firm Dynamics and Aggregate Fluctuations," mimeo.

den Haan, W. J., G. Ramey, and J. Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(3), 482–498.

Drautzburg, T. (2014): "Entrepreneurial Tail Risk: Implications for Employment Dynamics," mimeo.

Elsby, M., B. Hobijn, and A. Sahin (forthcoming): "On the Importance of the Participation Margin for Labor Market Fluctuations," *Journal of Monetary Economics*.

Elsby, M., and R. Michaels (2013): "Marginal Jobs, Heterogeneous Firms and Unemployment Flows," *American Economic Journal: Macroeconomics*, 5(1), 1–48.

Fort, T., J. Haltiwanger, R. Jarmin, and J. Miranda (2013): "How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size," NBER Working Paper no. 19134.

Gertler, M., and A. Trigari (2009): "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 117(1), 38–86.

Gourio, F., T. Messer, and M. Siemer (2014): "What is the Economic Impact of the Slowdown in New Business Formation?," Chicago FED Letter no. 326.

Hagedorn, M., and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98(4), 1692–1706.

Haltiwanger, J., R. Jarmin, and J. Miranda (2013): "Who Creates Jobs? Small vs. Large vs. Young," *The Review of Economics and Statistics*, 45(2), 347–361.

Hawkins, W. (2011): "Do Large-Firm Bargaining Models Amplify and Propagate Aggregate Productivity Shocks?," mimeo.

——— (2014): "Bargaining with Commitment between Workers and Large Firms," *Review of Economic Dynamics*, forthcoming.

Hopenhayn, H., and R. Rogerson (1993): "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy*, 101(5), 915–938.

Hurst, E., and B. Pugsley (2011): "What do Small Businesses Do?," *Brookings Papers on Economic Activity*, pp. 73–142, Fall.

Kaas, L., and P. Kircher (2011): "Efficient Firm Dynamics in a Frictional Labor Market," IZA discussion paper no. 5452.

——— (2014): "Efficient Firm Dynamics in a Frictional Labor Market," mimeo.

Klette, T. J., and S. Kortum (2004): "Innovating Firms and Aggregate Innovatino," *Journal of Political Economy*, 112.

Krusell, P., and A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(51), 867–896.

Lee, Y., and T. Mukoyama (2013): "Entry, Exit, and Plant-level Dynamics over the Business Cycle," mimeo.

Merz, M., and E. Yashiv (2007): “Labor and teh Market Value of the Firm,” *American Economic Review*, 97(4), 1419–1431.

Mortensen, D., and C. Pissarides (1994): “Job Creation and Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 61(3), 397–415.

Moscarini, G., and F. Postel-Vinay (2012): “The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment,” *American Economic Review*, 102(6), 2509–2539.

Ouimet, P., and R. Zarutskie (2013): “Who Works for Startups? The Relation between Firm Age, Employee Age and Growth,” Finance and Economics Discussion Series, Federal Reserve Board, 2013-75.

Petrongolo, B., and C. A. Pissarides (2001): “Looking into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390–431.

Pugsley, B., and A. Sahin (2014): “Grown-Up Business Cycles,” Federal Reserve Bank of New York Staff Reports, no. 707.

Sahin, A., S. Kitao, A. Cororaton, and S. Laiu (2011): “Why Small Businesses Were Hit Harder by the Recent Recession,” Federal Reserve Bank of New York, Current Issues in Economics and Finance.

Sahin, A., J. Song, G. Topa, and G. Violante (2014): “Mismatch Unemployment,” *American Economic Review*, 104(11), 3529–3564.

Saint-Paul, G. (2002): “Employment Protection, International Specialization, and Innovation,” *European Economic Review*, 46(2), 375–395.

Samaniego, R. (2008): “Entry, Exit, and Business Cycles in a General Equilibrium Model,” *Review of Economic Dynamics*, 11(3).

Schaal, E. (2012): “Uncertainty, Productivity and Unemployment in the Great Recession,” mimeo.

Schmalz, M., D. Sraer, and D. Thesmar (forthcoming): “Housing Collateral and Entrepreneurship,” *Journal of Fincnace*.

Sedláček, P. (2014): “The Aggregate Matching Function and Job Search Behavior from Employment and Out of the Labor Force,” mimeo.

Sedláček, P., and V. Sterk (2014): “The Growth Potential of Startups Over the Business Cycle,” mimeo.

Siemer, M. (2014): “Firm Entry and Employment Dynamics in the Great Recession,” mimeo.

Stole, L., and J. Zwiebel (1996): “Intra-firm Bargaining Under Non-binding Contracts,” *Review of Economic Studies*, 63(3), 375–410.

Appendix

A. Derivation of wages

As has been discussed in the main text, wages are being set at the beginning of the period, i.e. after observing all the aggregate shocks but before worker-specific productivity disturbances realize. The marginal surplus consists of the marginal value of a job for the firm and the worker less the value of unemployment for the worker (the value of a vacant job for the firm is driven down to zero because of free entry).

The beginning-of-period value of a job for a firm of type i and age a is given by

$$\mathcal{J}_{i,a} = (1 - H(\widetilde{z}_{i,a})) \left[\alpha \frac{y_{i,a}}{n_{i,a}} - w_{i,a} - \frac{\partial w_{i,a}}{\partial n_{i,a}} n_{i,a} - \frac{\kappa}{\gamma} x_{i,a}^{\gamma} + \beta(1 - \delta_a) \mathbb{E} \mathcal{J}'_{i,a+1} \right], \tag{10}$$

where the firm obtains the marginal product less the wage (where it has been anticipated that wages depend on the number of workers employed in the firm) and less the additional costs of hiring. If the firm survives into the next period, it obtains the marginal value of a job for a firm of age $a + 1$. The beginning-of-period value of a job for a worker employed in a firm of type i and age a is given by

$$\mathcal{W}_{i,a} = (1 - H(\widetilde{z}_{i,a})) \left[w_{i,a} + \beta(1 - \delta_a) \mathbb{E} \left(\mathcal{W}'_{i,a+1} - \mathcal{U}' \right) + \beta \mathbb{E} \mathcal{U}' \right] + H(\widetilde{z}_{i,a}) \mathcal{U}, \tag{11}$$

where the worker is fired with probability $H(\widetilde{z}_{i,a})$ in which case he/she obtains the value of unemployment ($\mathcal{U}$). With probability $(1 - H(\widetilde{z}_{i,a}))$ the worker remains in the employment relationship, obtains a wage $w_{i,a}$ and unless the firm shuts down at the end of the period, continues in the employment relationship until the next period. The value of unemployment is given by

$$\mathcal{U} = b + F\beta \sum_i \sum_a \frac{\omega_{i,a} v_{i,a}}{V} (1 - \delta_a) \mathbb{E} \left(\mathcal{W}'_{i,a+1} - \mathcal{U}' \right) + \beta \mathbb{E} \mathcal{U}', \tag{12}$$

where the unemployed worker obtains a flow income of b, with probability F he/she gets hired by a firm with an open vacancy and obtains next periods' surplus. The latter is a weighted average of worker surpluses in all existing firms which are hiring this period, where (because of the assumption of random search) the weights are the relative shares of vacancies of individual firms in the total number of open positions.

Given the above definitions of marginal worker and firm values, wages are an outcome of a Nash bargaining between the firm and its workers over the marginal surplus. In particular, the wage solves the following condition $(1 - \eta)(\mathcal{W}_{i,a} - \mathcal{U}) = \eta \mathcal{J}$, where η is the bargaining power of workers. The bargained wage is then give by

$$w_{i,a} = \eta \left(\frac{\alpha y_{i,a} / n_{i,a}}{1 - \eta(1 - \alpha)} + \frac{\kappa(\gamma - 1)}{\gamma} x_{i,a}^{\gamma} + \Phi \right) + (1 - \eta) b, \tag{13}$$

where $\Phi = \kappa V/U \sum_i \sum_a \frac{\omega_{i,a} v_{i,a}}{V} x_{i,a}^{\gamma - 1}$.

B. Solution method

This part of the appendix provides details on the numerical solution procedure, which is the same as in Sedláček and Sterk (2014). To economize on notation, we can express the model compactly as:

$$\mathbb{E}_t f\left(y_{t+1}, y_t, x_{t+1}, x_t; \Upsilon, \zeta\right) = 0$$

where x_t is a vector containing the state variables (all variables in $\mathcal{S}_t$) and y_t is a vector containing the non-preditermined variables, Υ is a vector containing all parameters of the model and ζ is a scalar parameter pre-multiplying the covariance matrix of the shock innovations, as in Schmitt-Grohé and Uribe (2004). Importantly, the above is system of a finite number of expectational difference equations.

B.1. Solving for the steady state without aggregate uncertainty

One first solves for the equilibrium of a version of the model without aggregate uncertainty. That is, one find vectors $\overline{y}$ and $\overline{x}$ that solve $f\left(\overline{y}, \overline{y}, \overline{x}, \overline{x}; \Upsilon, 0\right) = 0$. As described in the main text, the calibration targets various parameters to match long-run statistics. The calibration procedure has the following steps:

1. given values for the technology types (ϵ_i), the vacancy posting cost (κ) and the distribution of worker-specific productivity shocks ($H(\mu_h, \sigma_H)$), one can calculate the growth paths of firm-level employment, firm values and the endogenous separation rates leading towards the firm size distribution targets. The firm size distribution targets are informative about the values of the firm-specific productivity levels.
2. given firm values of startups from (1.), and a value of the entry cost, one can back out the startup probabilities (and therefore the number of startup attempts) from the free entry conditions.
3. given the startup probabilities from (2.), and the firm shares of technology types taken from the assumed business opportunity distribution, one can back out the actual number of startups in each type.
4. given the mass startups in each type from (3.) and the exogenous survival rates, one can calculate the mass of firms in each age-type cell.
5. given firm masses from (4.), employment choices from (1.), one can calculate all the aggregates (total employment, output, consumption etc.).

B.2. Solving for the equilibrium with aggregate uncertainty

Next, one can solve for the dynamic equilibrium using first-order perturbation around the deterministic equilibrium (including the steady state *growth paths* of firms) found in the previous step. The first-order approximated solutions, denoted by hats, have the following form:

$$\begin{aligned}
\widehat{x}_{t+1} &= \overline{x} + \Theta\left(\widehat{x}_t - \overline{x}\right) \\
\widehat{y}_{t+1} &= \overline{y} + \Phi\left(\widehat{x}_t - \overline{x}\right)
\end{aligned}$$

where Θ and Φ are matrices containing the coefficients obtained from the approximation. The perturbation procedure is standard and carried out in one step.

Table 3: Estimation results

ρ_Z	AR coefficient, aggregate productivity shock	0.925	(0.021)
ρ_X	AR coefficient, entry cost shock	0.864	(0.058)
σ_Z	standard deviation, aggregate productivity shock	0.008	(0.001)
σ_X	standard deviation, entry cost shock	0.019	(0.003)

Notes: estimated parameter values of exogenous shocks (standard errors in brackets).

An advantage of perturbation methods is that the computational speed is relatively high and many state variables can be handled. An important prerequisite for perturbations to be accurate, however, is that deviations from the steady-state are not too large. For firm dynamics models like the one in this paper it may seem problematic because differences in employment levels across firms may be very large. The solution method adopted here, however, overcomes this problem since the steady state we perturb around contains the entire *growth paths* of firms. These growth paths, captured by the constants in the above equations, are themselves non-linear functions of age and type.

Hence, the fact that most newborn firms starts off much below their eventual sizes does not involve large accuracy losses since the same is true for the steady-state sizes of newborn firms. Similarly, the fact that the equilibrium features various firm types with very different optimal sizes does not reduce accuracy since we perturb around the growth path for each individual firm type.

C. Details on the model estimation

The structural model can be case into a state-space representation as

$$\mathcal{X}_t = \Phi \mathcal{X}_{t-1} + \Psi \epsilon_t, \tag{14}$$

$$\mathcal{Y}_t = \Theta \mathcal{X}_t, \tag{15}$$

where Φ and Ψ are coefficient matrices (functions of the structural parameters) ϵ_t is a vector of the two structural shocks (aggregate productivity and the entry cost shock) and Θ is a selection matrix mapping model variables to observables (note that it is assumed that there is no measurement error).

The estimation is conducted using 2 variables: unemployment and the number of startups. While unemployment is available on a quarterly basis (the model period), the number of startups is only available annually. Both variables are expressed in log-deviations from their respective linear trends (over the entire sample period between 1977 and 2012).

As is explained in the main text, all model parameters are calibrated using targets from the pre-crisis period (and assuming no entry cost shocks), except for those related to the shock processes. The latter parameters are estimated by maximum likelihood (ML). The following values are used to initialize the minimization routine: $\rho_Z = \rho_X = 0.95$, $\sigma_Z = \sigma_X = 0.01$. Moreover, the parameter values are bounded by 0 and 0.99 in the case of ρ_Z and ρ_X and by 0 and 0.1 (0.3) in the case of σ_Z (σ_X).

Table 3 reports the estimated parameter values and standard errors. The parameters

Figure C.1: Log-likelihood values for parameter deviations from point estimates

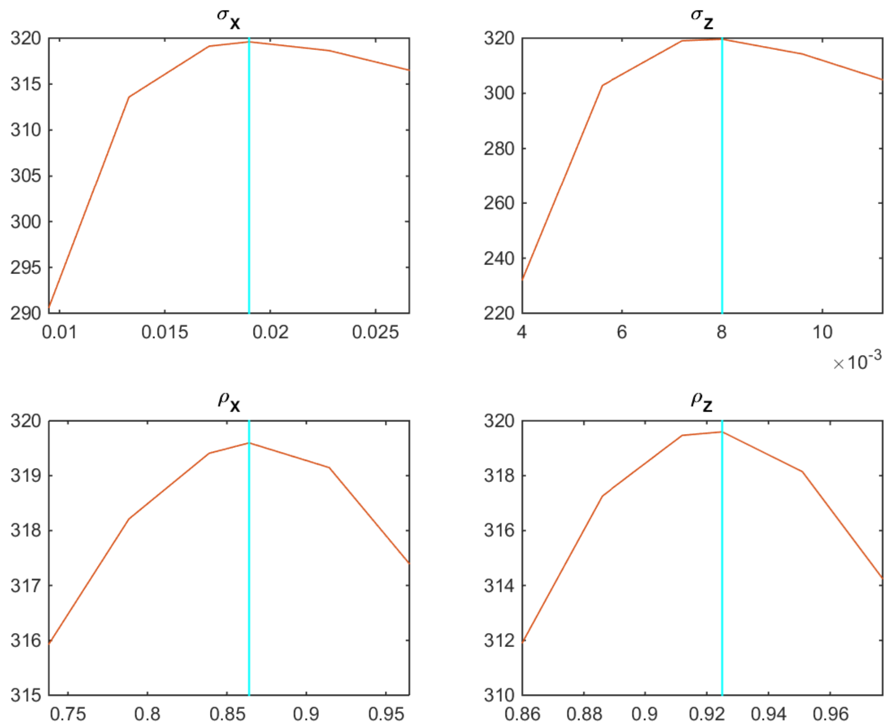

Notes: values of log-likelihood for deviations of individual parameters from their respective point estimates (keeping all other parameters fixed). Vertical lines indicate respective point estimates.

of the aggregate productivity shock are close to those found in the literature (and to those used in the calibration). The entry cost shock is somewhat less persistent than aggregate productivity and the standard deviation is larger. Parameters pertaining to the entry cost shock are estimated with considerably less accuracy than those associated with the aggregate productivity shock. This is because of the annual nature of the startup data (and hence missing values).

To inspect the convergence of the maximization, Figure C.1 plots the log-likelihood values for deviations of individual parameters from their point estimates (vertical lines). For all four parameters the point estimates are at the peaks of the partial likelihood functions.

D. Responses to an aggregate productivity shock

This Appendix shows the impulse response functions of several variables to a negative one-standard-deviation shock to aggregate productivity. Figure D.1 depicts the IRFs showing that both aggregate output and labor productivity fall following a negative productivity shock (top row). The drop in productivity also makes firm entry less attractive and the number of startups falls (second row, left panel). Even though the number of entrants starts reverting back to its steady state right after the first year, the number of startups is still 0.5 percent below its steady state even after 10 years. The resulting effect on the total number of firms is very persistent (second row, right panel).

The lower (labor) productivity is accompanied by an increase in the unemployment rate and the reduced incentives to hire new workers lead to a drop in vacancies (third row). The unemployment rate increase is a result of the drop in the number of entrants (firms), a spike in the separation rate (bottom right panel) and a fall in the probability with which unemployed workers find jobs (bottom left panel).

E. Robustness

This Appendix presents as robustness checks results for several alternative calibrations. First, the model is calibrated with almost linear vacancy posting costs and with considerably more convex costs than in the benchmark specification. Second, a larger maximum firm age is used to solve the model under the benchmark calibration.

E.1. Convexity of vacancy posting costs

The benchmark calibration assumes that the costs of posting vacancies are quadratic in the vacancy rate ($\gamma = 2$). This value is, however, chosen based on examples in the literature (see e.g. Gertler and Trigari, 2009) rather than direct evidence and therefore this section investigates how the results change when considering different values for γ. In particular, the benchmark results are compared with those from a calibration with almost linear costs ($\gamma = 1.1$) and a calibration with more convexity in the vacancy posting costs ($\gamma = 4$). In both alternative specifications, other parameter values are re-calibrated such that the model matches the same calibration targets as in the benchmark case.

Figure E.1 and E.2 show the impulse response functions of several variables to a negative (positive) aggregate productivity (entry cost) shock. The figures show that the model properties change little with alternative values for the convexity of hiring costs. The largest differences are in the dynamics of vacancies which become more responsive the larger the convexity of hiring costs. This then translates into a more responsive labor market tightness and thus also the job finding (filling) probability. For the same reason, the separation rate is also more responsive as labor market tightness enters workers' outside options.

E.2. Maximum firm age

The applied solution method relies on assuming a maximum firm age beyond firms cannot grow older. In the benchmark specification the maximum firm age is set at 25 years. This value strikes a compromise between the computational burden (the employment level and mass of each firm age-type is a state-variable) and allowing for a realistic firm-age distribution.

This subsection documents that the results change very little when considering a larger maximum firm age. In particular, $A = 160$, meaning that the maximum firm age is 40 years. Figures E.3 and E.4 show impulse response functions of several variables to a negative (positive) aggregate productivity (entry cost) shock. The figure depicts IRFs under the benchmark specification ($A = 100$) and compares them to the alternative one with a maximum age of $A = 160$. Increasing the maximum firm age does very little to the dynamics of the model.

Figure D.1: Impulse response functions to an aggregate productivity shock

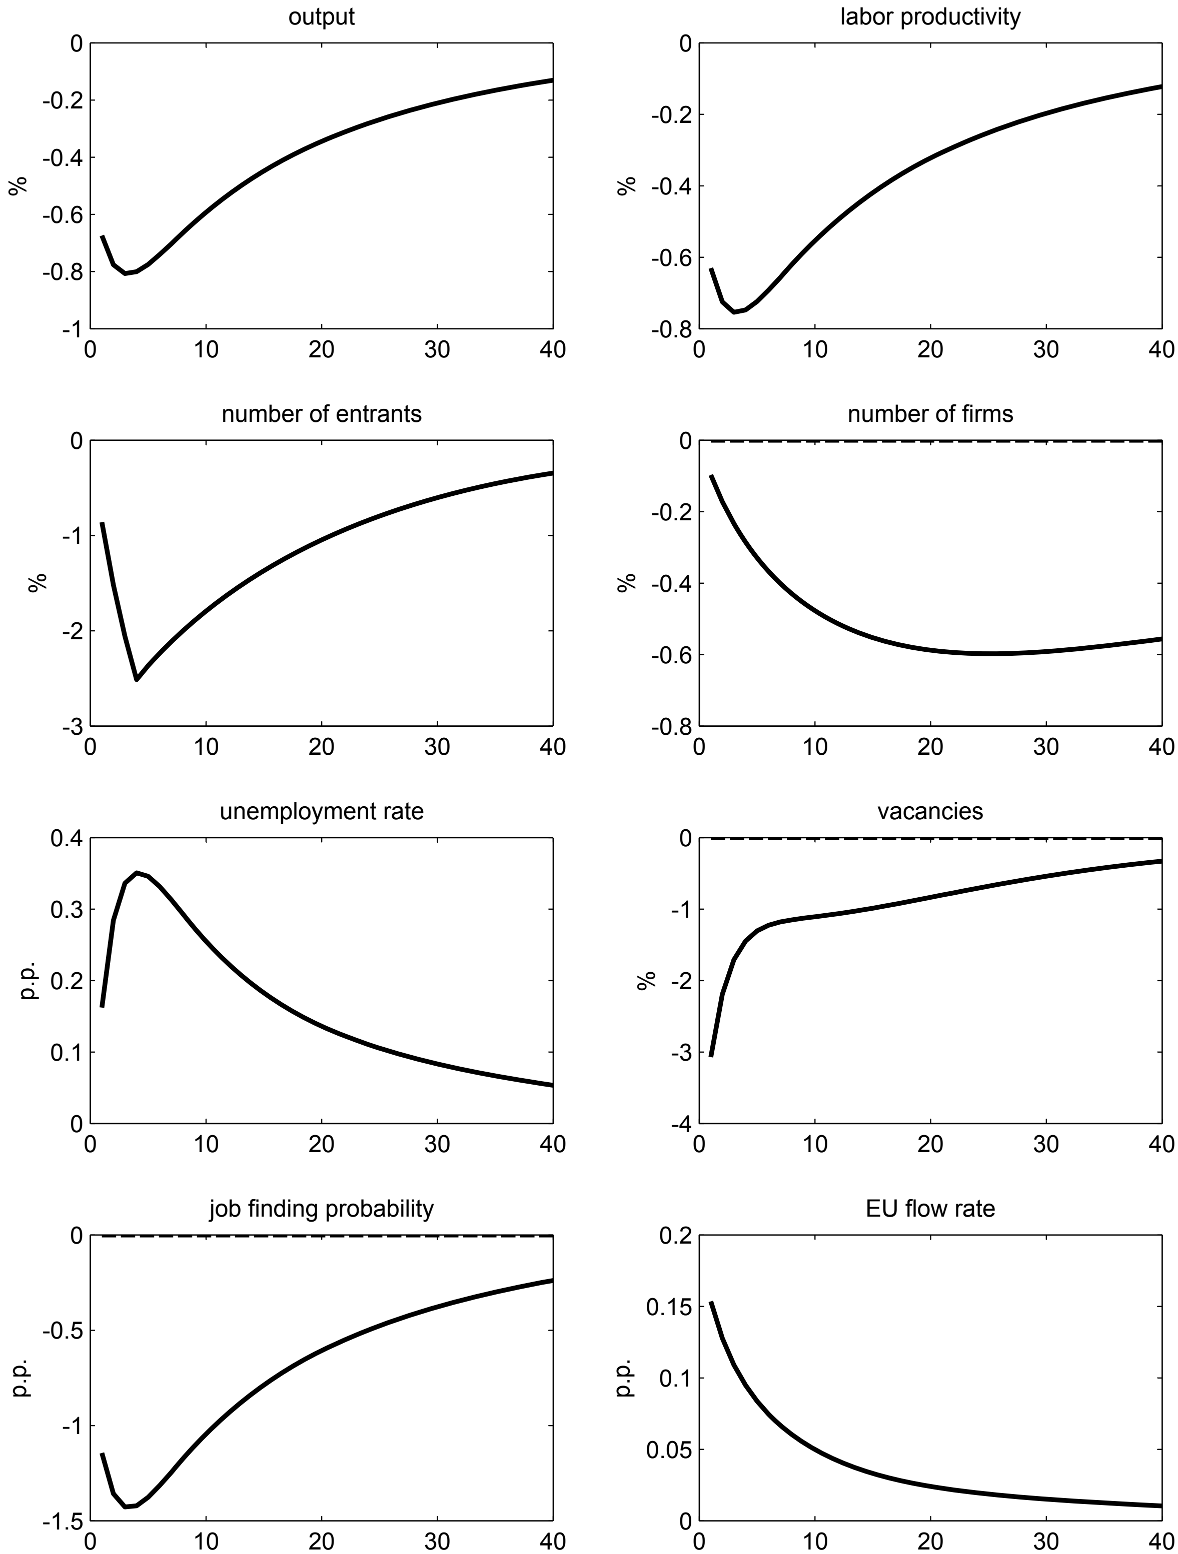

Notes: impulse response functions to a negative one-standard-deviation shock to aggregate productivity.

Figure E.1: Impulse response functions to an entry cost shock

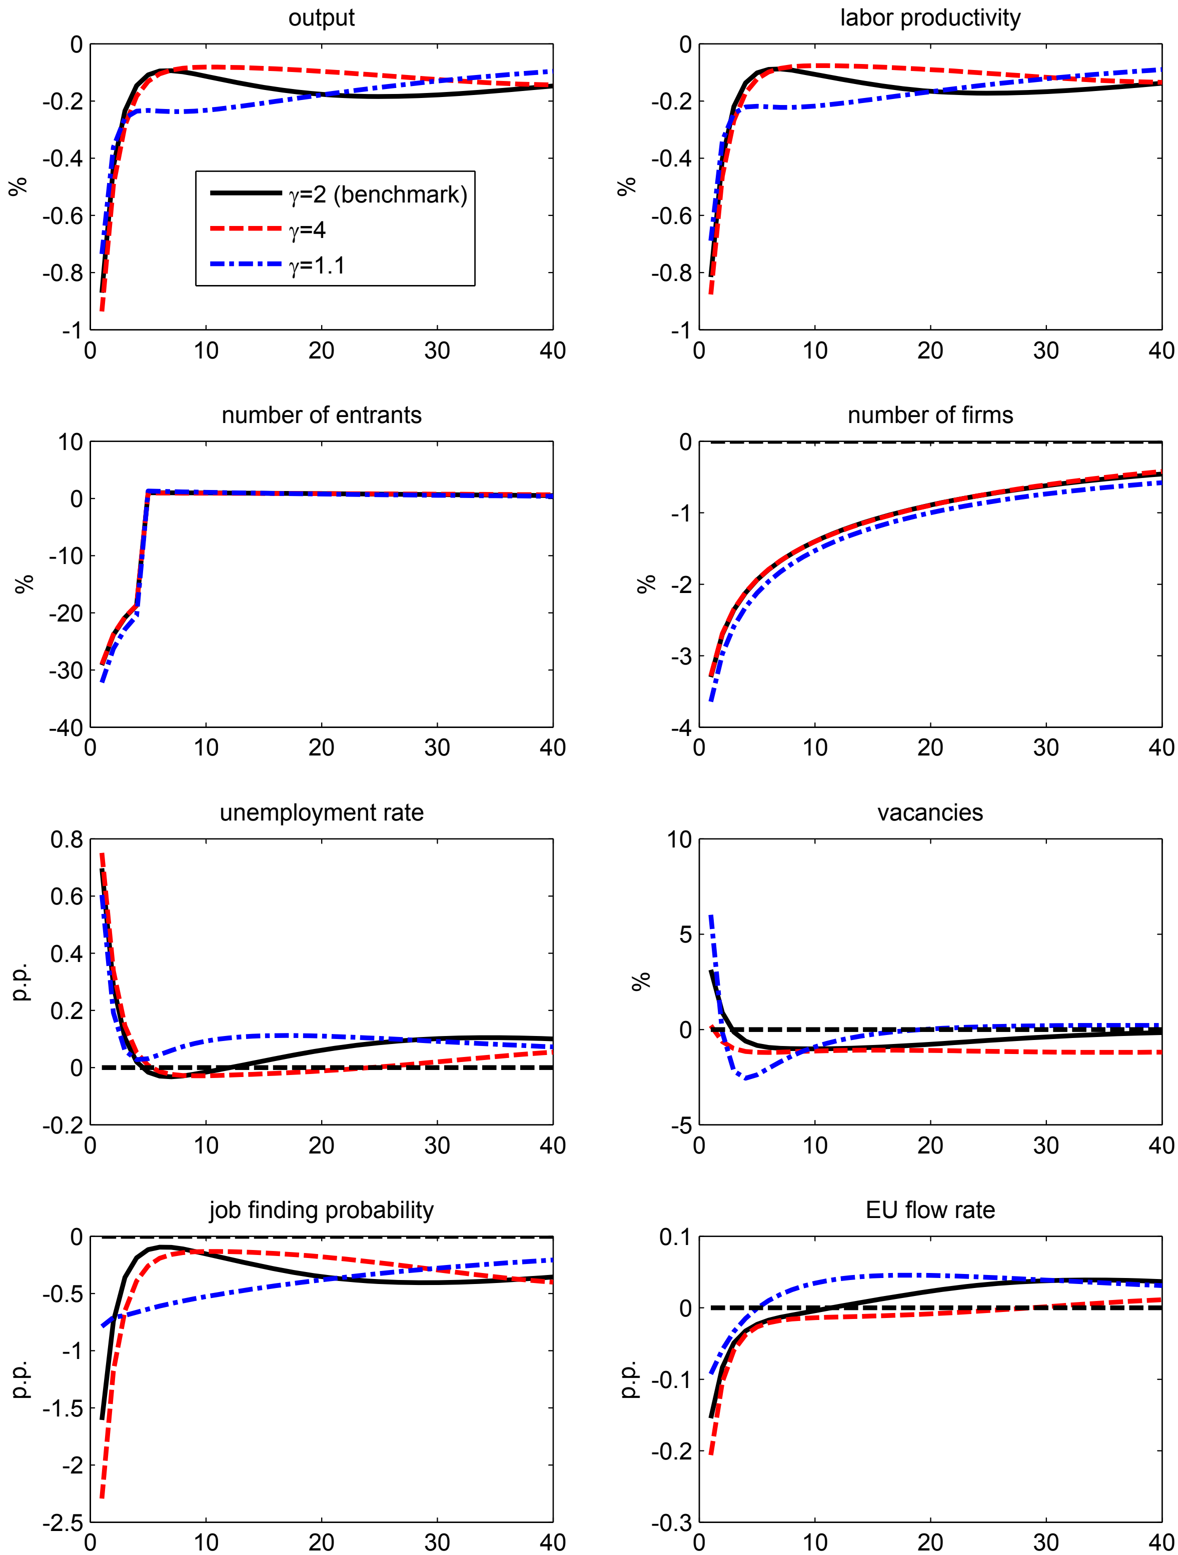

Notes:impulse response functions to a positive entry cost shock for different calibrations of the convexity of vacancy posting costs (γ). The magnitude of the shock is chosen such that entrants fall by about 30 percent in the benchmark specification.

Figure E.2: Impulse response functions to an aggregate productivity shock

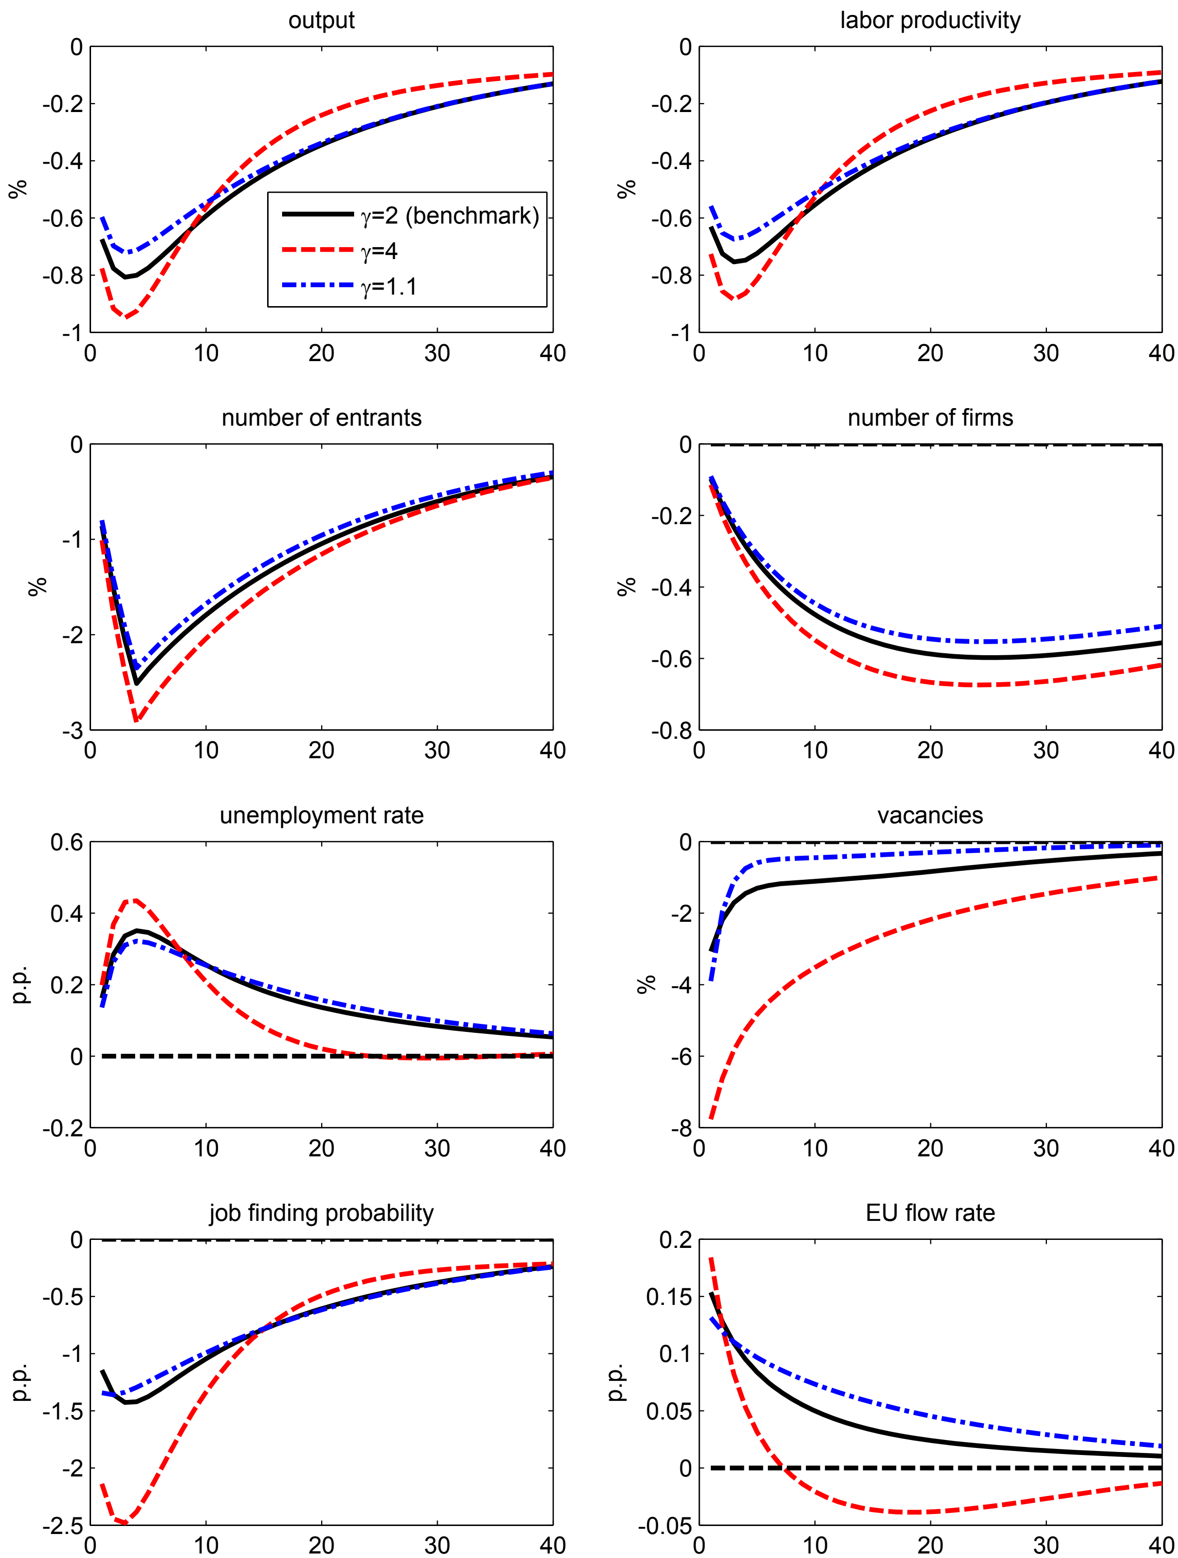

Notes: impulse response functions to a negative one-standard-deviation shock to aggregate productivity for different calibrations of the convexity of vacancy posting costs (γ).

Figure E.3: Impulse response functions to an entry cost shock

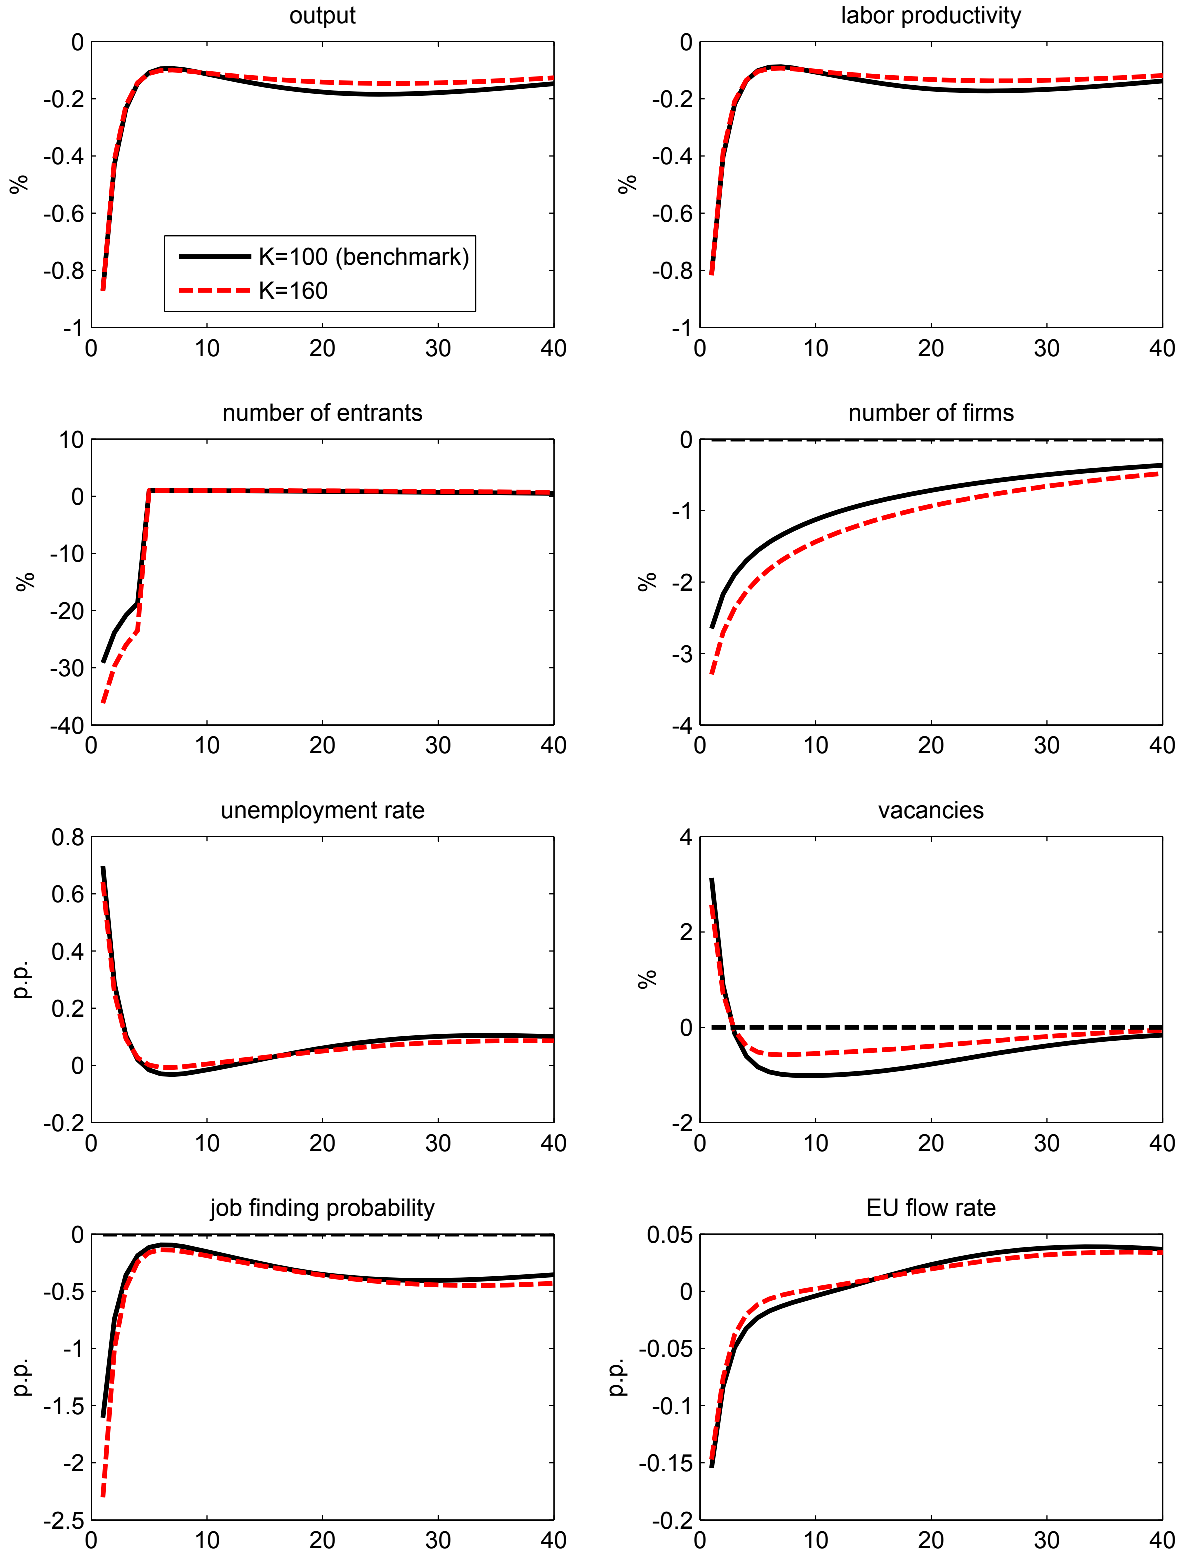

Notes: impulse response functions to a positive entry cost shock for different calibrations of the convexity of vacancy posting costs (γ). The magnitude of the shock is chosen such that entrants fall by about 30 percent in the benchmark specification.

Figure E.4: Impulse response functions to an aggregate productivity shock

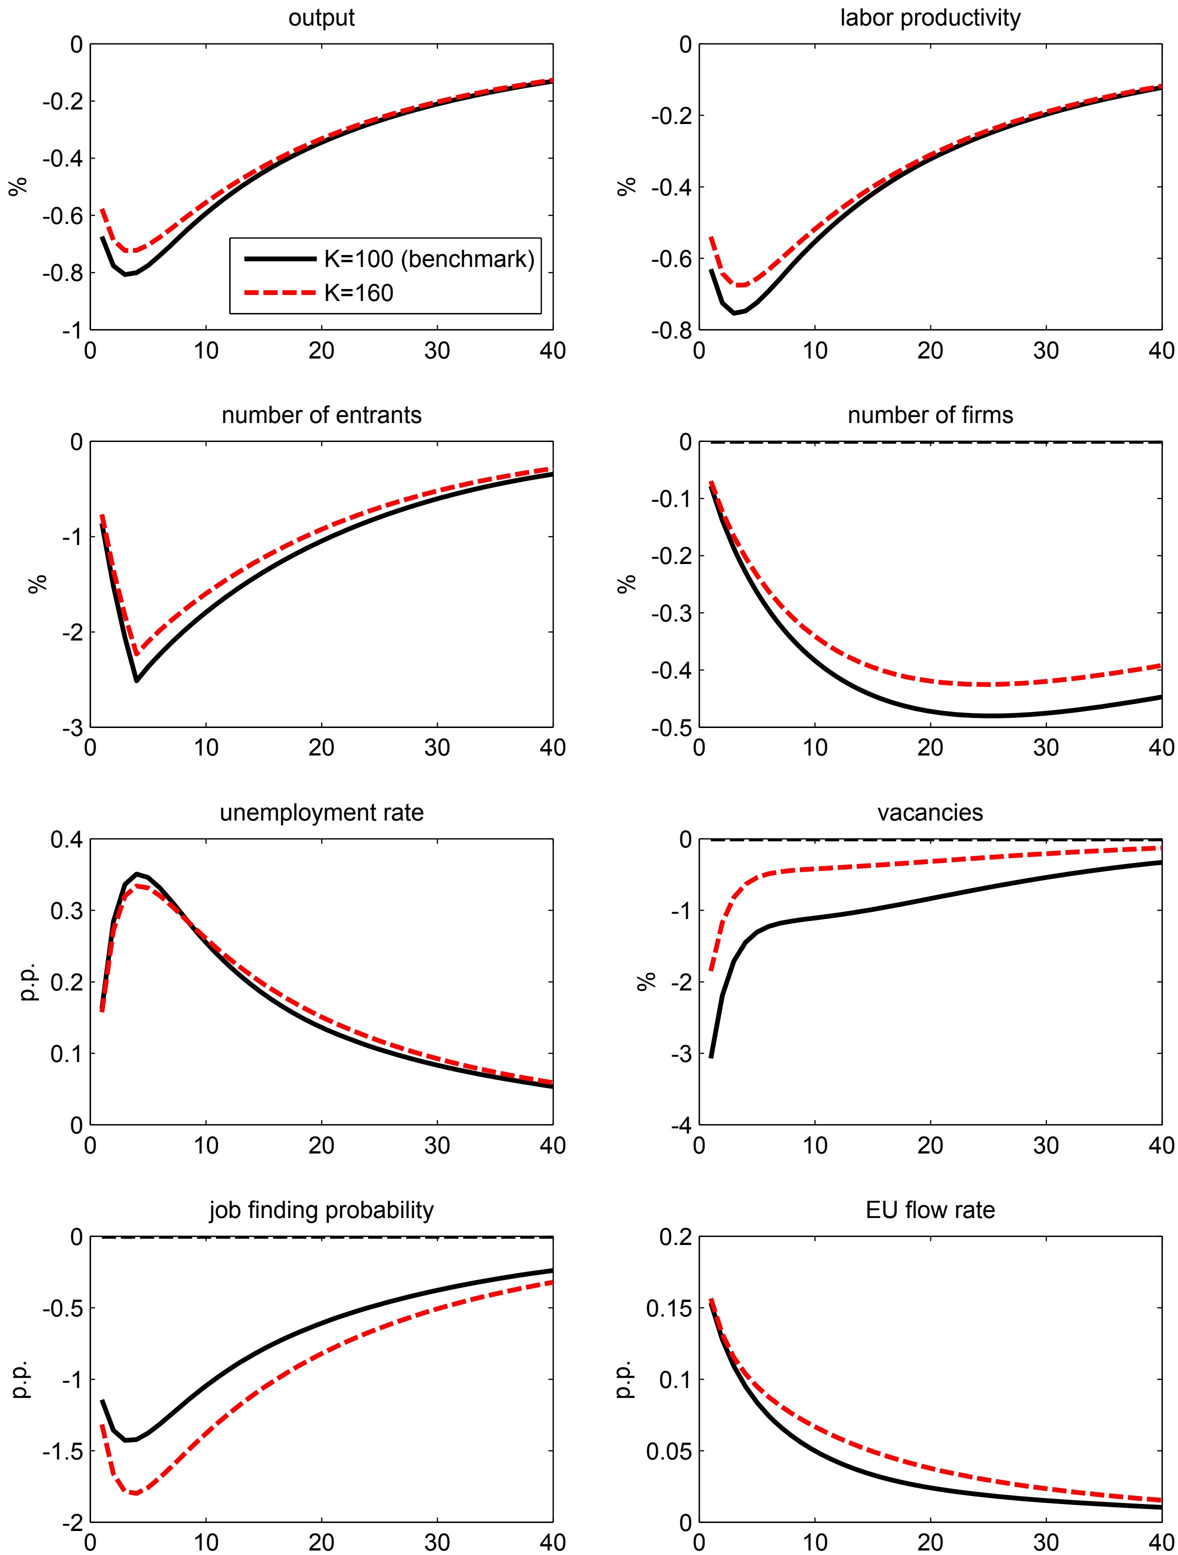

Notes: impulse response functions to a negative one-standard-deviation shock to aggregate productivity for different calibrations of the convexity of vacancy posting costs (γ).

E.3. Time-variation in firm exit rates

The benchmark model abstracts from variation in firm survival rates. This choice is partly made for tractability, but in part it is motivated by the BDS data itself. In particular, fixing job destruction at exiting firms to its sample average results in (a counterfactual) aggregate employment which is almost identical to the true time series (it varies more by 0.9%).

Nevertheless, this Appendix investigates how the properties of the model change when firm exit rates are allowed to vary. Specifically, the BDS data suggests that the volatility of firm exit rates is about 3.1 times that of output volatility. To allow for such variation in the model, I let the age-dependent firm exit rate depend on aggregate productivity, $\delta_{a,t} = \delta_a A_t^{-\overline{\delta}_A}$, where $\overline{\delta}_A$ is the elasticity of the exit rate with respect to aggregate productivity. Notice that this specification means that firm exit increases during recessions.[34] The magnitude of $\overline{\delta}_A$ is chosen such that the model replicates the relative volatility of the total firm exit rate with respect to output volatility. All other model parameters are recalibrated to match the original targets.

Figures E.5 and E.6 show the impulse response functions of several variables to an aggregate productivity and an entry cost shock, respectively, under the benchmark specification with fixed exit rates and compare them to those where exit rates vary. For the response to an aggregate productivity shock, the specification with varying exit rates creates somewhat more persistence. This is because the number of firms responds more sluggishly (it is not only lower entry, but also increased firm exit which pulls down the number of firms in recessions). Otherwise the results are very similar to the benchmark case.

The responses to an entry cost shock are almost identical to the benchmark, but not entirely. The reason is that even though firm exit responds only to aggregate productivity shocks, the model parameters are different owing to the recalibration under the case when exit rates are allowed to vary.

[34] The correlation of firm exit rates and labor productivity is actually only about -0.15 in the data owing to a relatively jumpy exit rate time-series.

Figure E.5: Impulse response functions to an entry cost shock

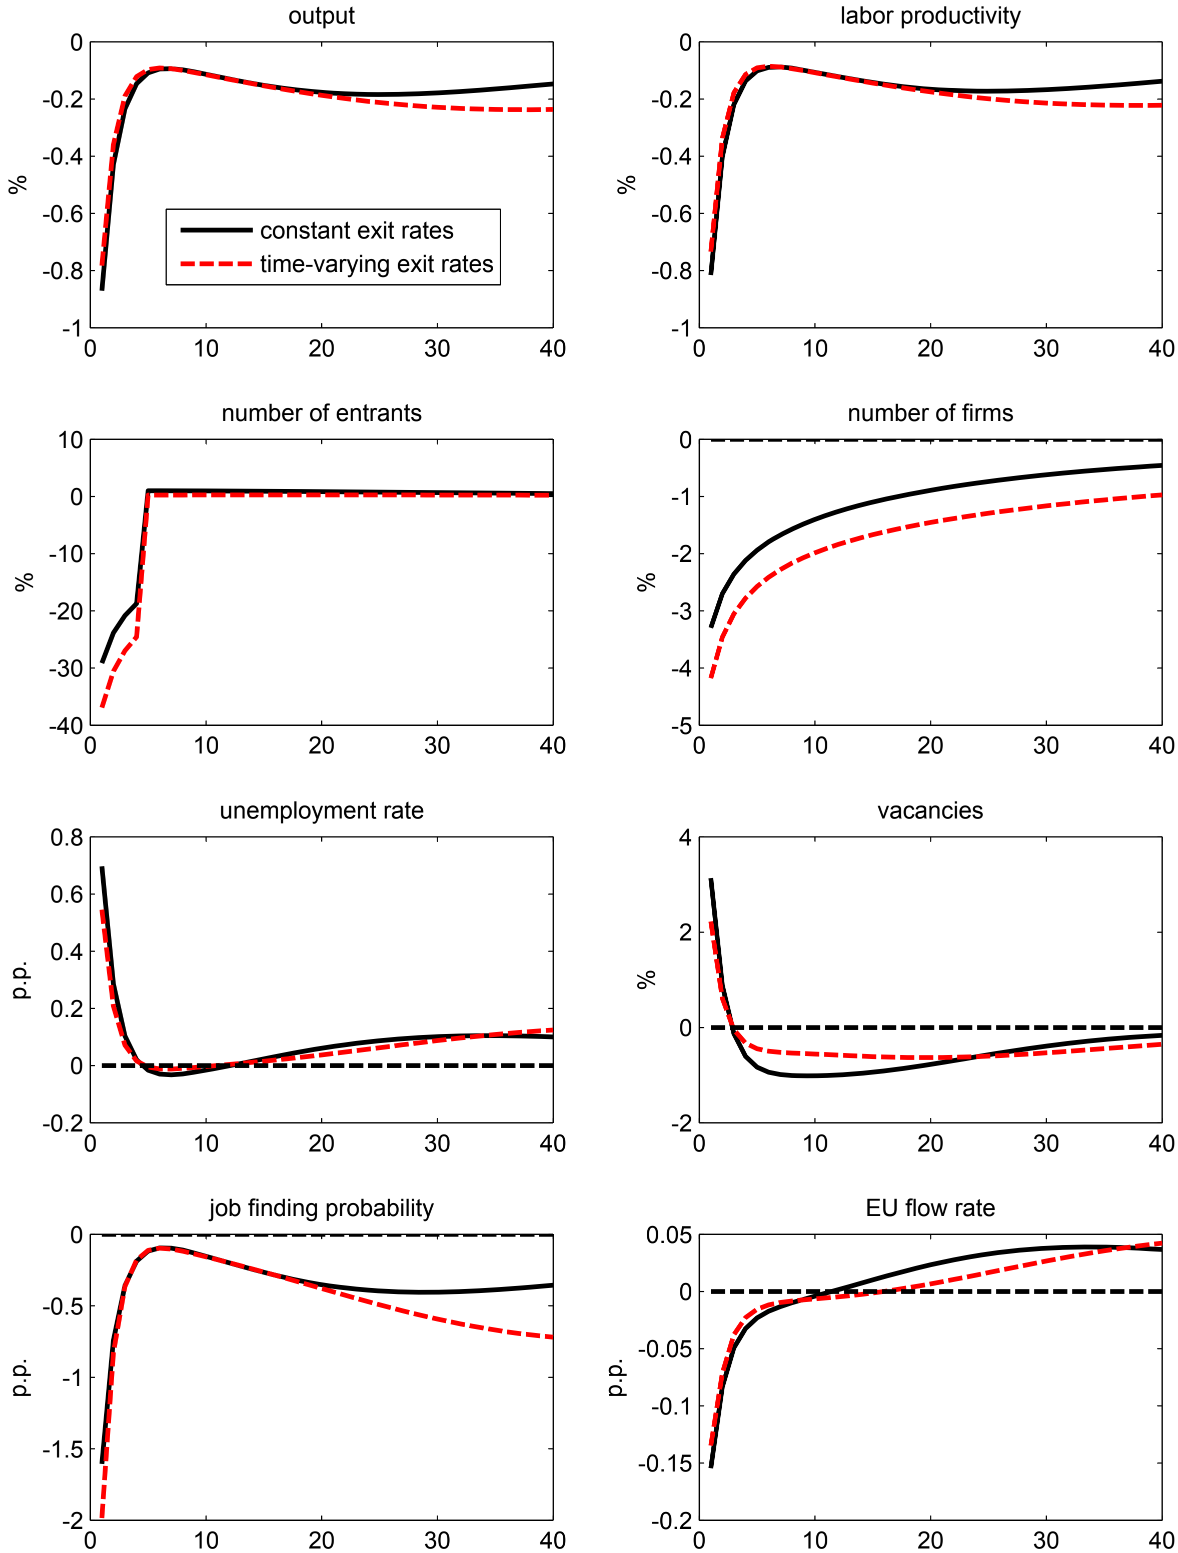

Notes:impulse response functions to a positive entry cost shock for the benchmark specification with fixed exit rates and for a calibration in which exit rates vary.

Figure E.6: Impulse response functions to an aggregate productivity shock

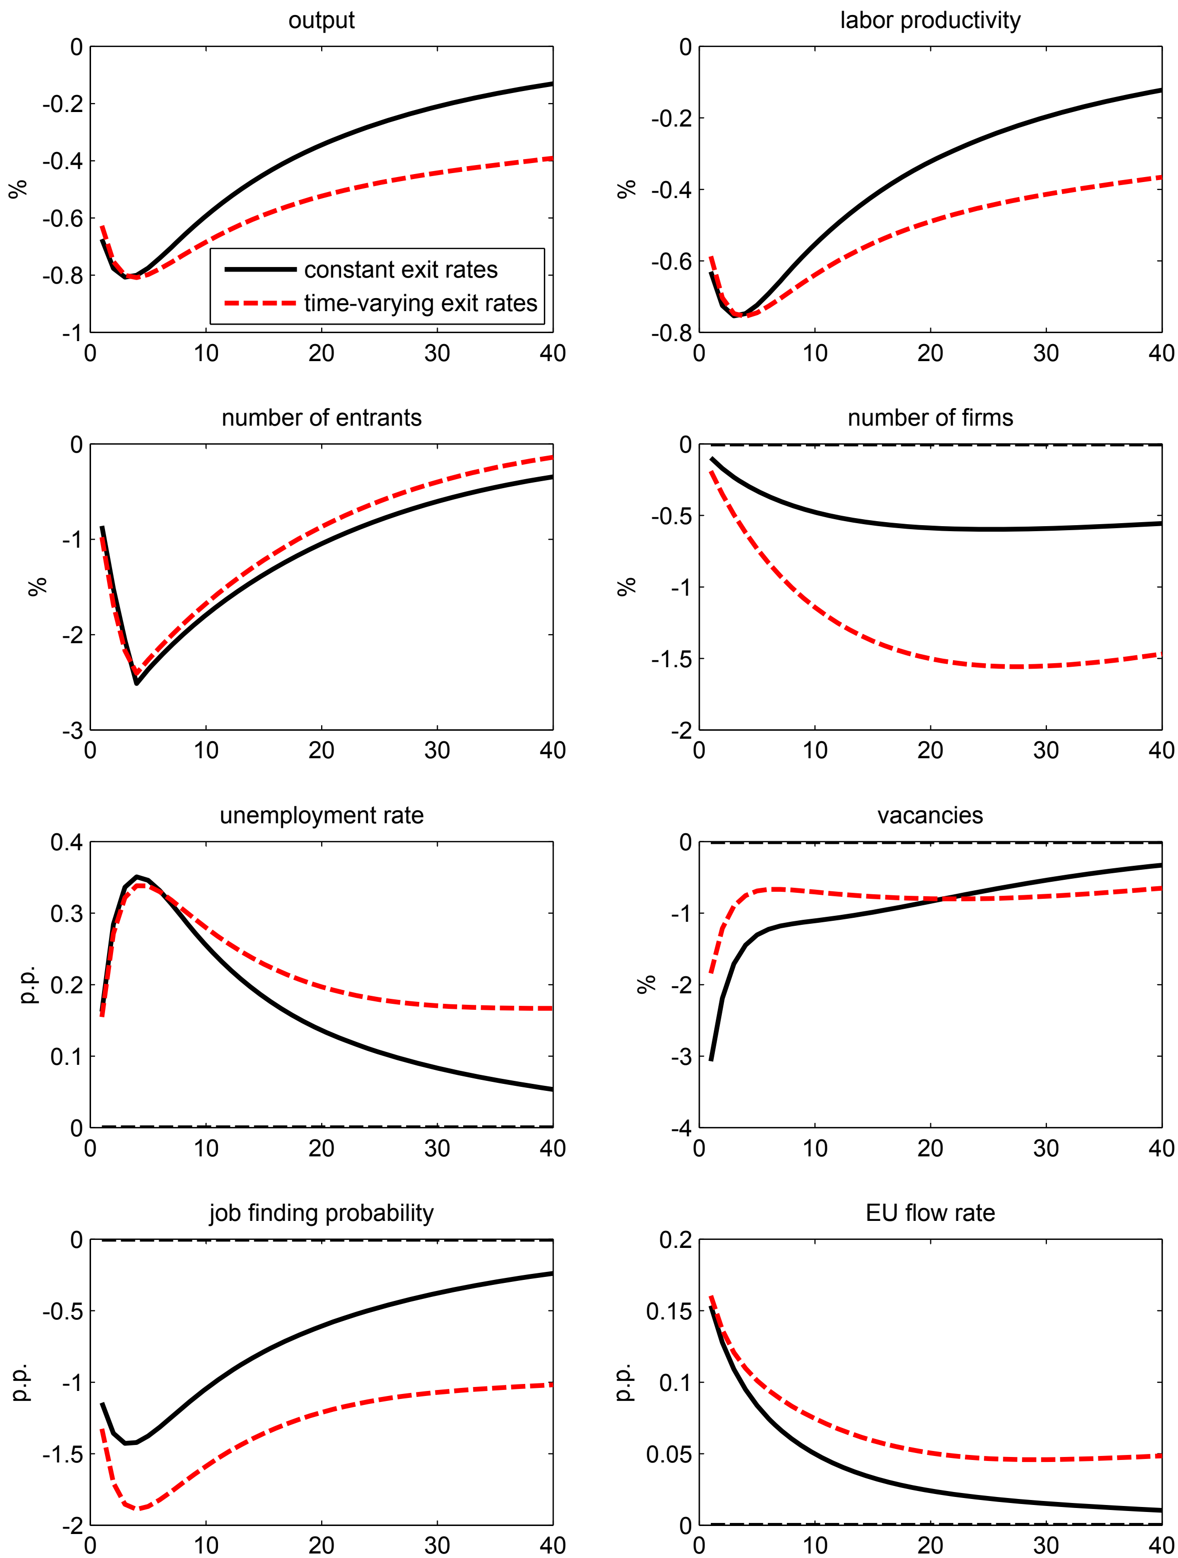

Notes: impulse response functions to a negative one-standard-deviation shock to aggregate productivity for the benchmark specification with fixed exit rates and for a calibration in which exit rates vary.