Alessandro Gavazza, Simon Mongey, and Giovanni L. Violante

Aggregate recruiting intensity

Aggregate Recruiting Intensity*

Alessandro Gavazza[‡] *Simon Mongey*[§] *and* *Giovanni L. Violante*[¶]

August 22, 2017

Abstract

We develop an equilibrium model of firm dynamics with random search in the labor market where hiring firms exert recruiting effort by spending resources to fill vacancies faster. Consistent with microevidence, fast-growing firms invest more in recruiting activities and achieve higher job-filling rates. These hiring decisions of firms aggregate into an index of economy-wide recruiting intensity. We study how aggregate shocks transmit to recruiting intensity, and whether this channel can account for the dynamics of aggregate matching efficiency during the Great Recession. Productivity and financial shocks lead to sizable pro-cyclical fluctuations in matching efficiency through recruiting effort. Quantitatively, the main mechanism is that firms attain their employment targets by adjusting their recruiting effort in response to movements in labor market slackness.

Keywords: Aggregate Matching Efficiency, Firm Dynamics, Macroeconomic Shocks, Recruiting Intensity, Unemployment, Vacancies.

*We thank Steve Davis, Jason Faberman, Mark Gertler, Bob Hall, Leo Kaas, Ricardo Lagos, and Giuseppe Moscarini for helpful suggestions at various stages of this project, and our discussants Russell Cooper, Kyle Harkenoff, William Hawkins, Jeremy Lise, and Nicolas Petrosky-Nadeau for many useful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

[‡]London School of Economics and CEPR

[§]Federal Reserve Bank of Minneapolis

[¶]Princeton University, CEPR, IFS, IZA, and NBER

1 Introduction

A large literature documents cyclical changes in the rate at which the US macroeconomy matches job seekers and employers with vacant positions. Aggregate matching efficiency, measured as the residual of an aggregate matching function that generates hires from inputs of job seekers and vacancies, epitomizes this crucial role of the labor market. In fact, matching efficiency is a key determinant, over and above market tightness, of the aggregate job-finding rate, i.e. the speed at which idle workers are hired. Swings in the job-finding rate account for the bulk of unemployment fluctuations (Shimer, 2012). Identifying the deep determinants of aggregate matching efficiency is therefore necessary to fully understand labor market dynamics.

The Great Recession represents a particularly stark episode of deterioration in aggregate matching efficiency. Our reading of the data, displayed in Figure 1, is that this decline contributed to a depressed vacancy yield, to a collapse in the job-finding rate, and to persistently higher unemployment following the crisis.

A number of explanations have been offered for the decline in aggregate matching efficiency during the recession, virtually all of which have emphasized the worker side.[1] A shift in the composition of the pool of job seekers toward the long-term unemployed, by itself, goes a long way toward explaining the drop (Hall and Schulhofer-Wohl, 2015); however, as documented by Mukoyama, Patterson, and Şahin (2014), workers' job search effort is countercyclical and tends to compensate for compositional changes. Hornstein and Kudlyak (2016) include both margins in their rich measurement exercise and conclude that they offset each other almost exactly, leaving the entire drop in match efficiency from unadjusted data to be explained. A rise in occupational mismatch shows more promise, but it can account for at most one-third of the drop and for very little of its persistence (Şahin, Song, Topa, and Violante, 2014).

The alternative view we set forth in this paper is that fluctuations in the effort with which firms try to fill their open positions affect aggregate matching efficiency. When aggregated over firms, we call this factor *aggregate recruiting intensity*. Our goal is to investigate whether this factor is an important source of the dynamics of aggregate matching efficiency, and to study the economic forces that shape how it responds to macroeconomic shocks.

[1] A notable exception is the model in Sedláček (2014) that generates endogenous fluctuations in match efficiency through firms' time-varying hiring standards.

Figure 1: Labor market dynamics during the Great Recession (2008:01 - 2014:01)

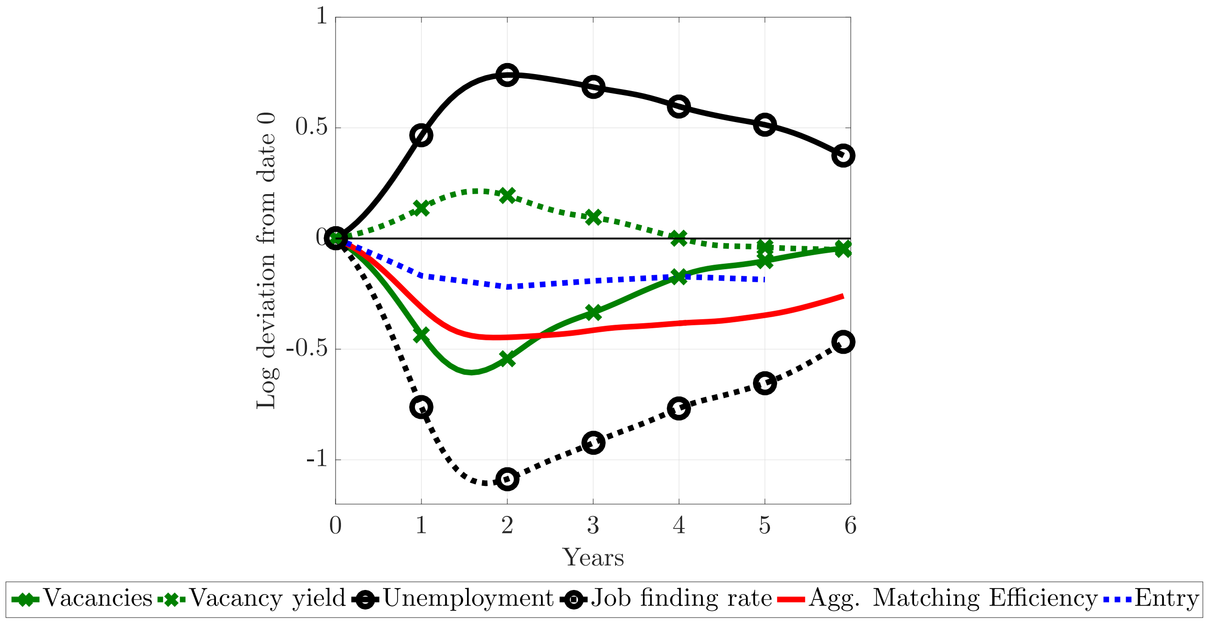

Notes: *(i)* Vacancies V_t and hires H_t (used to compute vacancy yield H_t/V_t) taken from monthly Job Openings and Labor Turnover Survey (JOLTS) data. Hires exclude recalls. *(ii)* Unemployment U_t is from the Bureau of Labor Statistics (BLS) and excludes workers on temporary layoffs. *(iii)* The job-finding rate is H_t/U_t. *(iv)* Aggregate matching efficiency is equal to $H_t/\left(V_t^{\alpha}U_t^{1-\alpha}\right)$ with $\alpha = 0.5$. *(v)* These first five series are measured from January 2001 to January 2014, expressed in logs and then HP-filtered. We plot level differences of these series from January 2008. *(vi)* (Firm) entry is taken from the Census Bureau's Business Dynamics Statistics and computed annually as the number of firms aged less than or equal to one year at the time of survey and is available from 1977 to 2007. To this we fit and remove a linear trend. We plot log differences of this series from 2007.

Our main motivation is the empirical analysis of recruiting intensity at the micro level in Davis, Faberman, and Haltiwanger (2013) (henceforth DFH)—the first paper to rigorously use JOLTS data to examine which factors are correlated with vacancy yields at the establishment level. The robust finding of DFH is that establishments with a larger hiring rate (total hires per employment) fill their vacancies at a faster rate.[2]

One would therefore expect that, if an aggregate negative shock depresses firm growth rates, aggregate recruiting intensity—and, thus, aggregate match efficiency—declines since hiring firms use lower recruiting effort to fill their posted vacancies. We call this transmission channel, whereby the macro shock affects the growth rate distribution of hiring firms, the *composition effect*. Macro shocks also induce movements in equilibrium labor market tightness. When a negative shock hits the economy, job seekers become more abundant relative to vacancies, so firms

[2]The numerous exercises in DFH show that this finding is not in any way spurious. For example, by definition, an establishment that luckily fills a large amount of its vacancies will have both a higher vacancy yield and a higher growth rate. The authors show that luck does not drive their main result.

meet workers more easily and can therefore exert less recruiting effort to reach a given hiring target. We call this second transmission channel the *slackness effect*, in reference to aggregate labor market conditions.

Both mechanisms seem potentially relevant in the context of the Great Recession. As evident from Figure 1, the data display a collapse in market tightness indicating the potential for a strong slackness effect. The figure also shows that the rate at which firms entered the economy fell dramatically in the aftermath of the recession. The dominant narrative is that the crisis was associated with a sharp reduction in borrowing capacity, and start-up creation as well as young firm growth are particularly sensitive to financial shocks (Chodorow-Reich, 2014; Siemer, 2014; Davis and Haltiwanger, 2015; Mehrotra and Sergeyev, 2015). Combining this observation with the fact that much of job creation (and an even larger share of gross hires) are generated by young firms (Haltiwanger, Jarmin, and Miranda, 2010) paves the way for a sizable composition effect.

Our approach is to develop a model of firm dynamics in frictional labor markets that can guide us to inspect the transmission mechanism of two common macroeconomic impulses—productivity and financial shocks—on aggregate recruiting intensity. The model is consistent with the stylized facts that are salient to an investigation of the interaction between macro shocks and recruiting activities: (i) it matches the DFH finding that increases in hiring rates are realized chiefly through increases in vacancy yields rather than increases in vacancy rates; (ii) it allows for credit constraints that hinder the birth of start-ups and slow the expansion of young firms; and (iii) it is set in general equilibrium, since the recruiting behavior of hiring firms depends on labor market tightness, which fluctuates strongly in the data (Shimer, 2005).

Our model is a version of the canonical Diamond-Mortensen-Pissarides random matching framework with decreasing returns in production and nonconvex hiring costs (Cooper, Haltiwanger, and Willis, 2007; Elsby and Michaels, 2013; Acemoglu and Hawkins, 2014). The model simultaneously features a realistic firm life cycle, consistent with its classic competitive setting counterparts (Jovanovic, 1982; Hopenhayn, 1992), and a frictional labor market with slack on both demand and supply sides. We augment this environment in three dimensions.

First, we allow for endogenous entry and exit of firms. This is a key element for understanding the effects of macroeconomic shocks on the growth rates of hiring firms, since it is

Figure 2: Breakdown of spending on recruiting activities. Source: Bersin and Associates (O'Leonard, 2011)

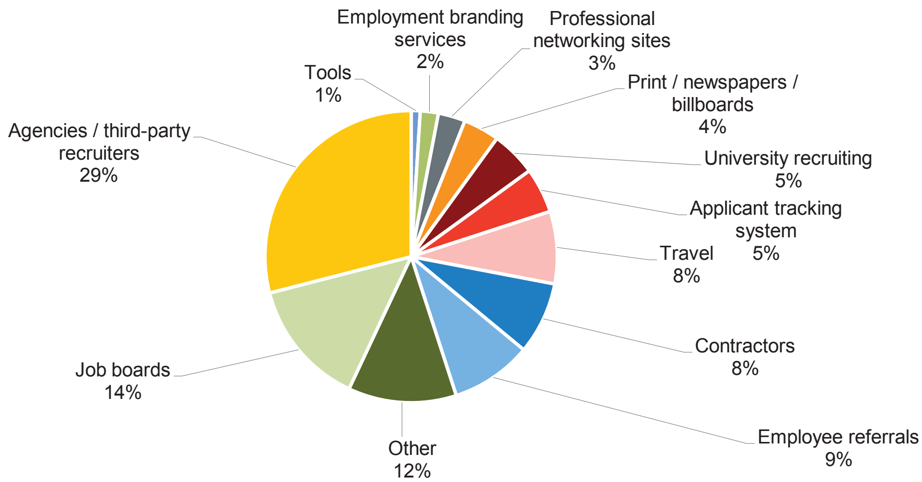

well documented that young firms account for a disproportionately large fraction of job creation, grow faster than old firms, and are more sensitive to financial conditions.

Second, we introduce a recruiting intensity decision at the firm level: besides the number of open positions that they are willing to fill in each period, hiring firms choose the amount of resources that they devote to recruitment activities. This endogenous recruiting intensity margin generates heterogeneous job-filling rates across firms. In turn, the sum of all individual firms' recruitment efforts, weighted by their vacancy share, aggregates to the economy's measured matching efficiency.

Third, we introduce financial frictions: incumbent firms cannot issue equity, and a constraint on borrowing restricts leverage to a multiple of collateralizable assets, as in Evans and Jovanovic (1989).[3]

We parameterize our model to match a rich set of aggregate labor market statistics and firm-level cross-sectional moments. In choosing the recruiting cost function, we reverse-engineer a specification that allows the model to replicate DFH's empirical relation between the job-filling rate and the hiring rate at the establishment level from the JOLTS microdata. Our parameteri-

[3]Other papers that consider various forms of financial constraints in frictional labor market models include Wasmer and Weil (2004), Petrosky-Nadeau and Wasmer (2013), Eckstein, Setty, and Weiss (2014), and Buera, Jaef, and Shin (2015), though none of these models displays endogenous fluctuations in match efficiency. An exception is Mehrotra and Sergeyev (2013), where a financial shock has a differential impact across industries and induces sectoral mismatch between jobseekers and vacancies.

zation of this cost function is based on a novel source of data, a survey of recruitment cost and practices based on over 400 firms that are representative of the US economy. Figure 2 gives a breakdown of spending on all recruitment activities in which firms engage in order to attract workers and quickly fill their open positions, as reported by the survey. Our hiring cost function is meant to summarize all such components.

We find that both productivity and financial shocks—modeled as shifts in the collateral parameter—generate substantial procyclical fluctuations in aggregate recruiting intensity. However, the financial shock generates movements in firm entry, labor productivity, and borrowing that are consistent with those observed during the 2008 recession, whereas the productivity shock does not. The credit tightening accounts for approximately half of the drop in aggregate matching efficiency observed during the Great Recession through a decline in aggregate recruiting intensity. Notably, our model is consistent with a key cross-sectional fact documented by Moscarini and Postel-Vinay (2016): the vacancy yield of small establishments spiked up as the economy entered the downturn, whereas that of large establishments was much flatter. The reason is that the financial shock impedes the growth of a segment of very productive, already large, but relatively young, firms with much of their growth potential still unrealized. These firms drastically cut their hiring effort.

Our examination of the transmission mechanism indicates that the slackness effect is the dominant force: aggregate recruiting intensity falls mainly because the number of available job seekers per vacancy increases, allowing firms to attain their recruitment targets even by spending less on hiring costs. Surprisingly, the impact of the shock through the shift in the distribution of firm growth rates (and, in particular, the decline in firm entry and young-firm expansion) on aggregate recruiting intensity is quantitatively small. Two counteracting forces weaken this composition effect: (i) hiring firms are selected, and thus are relatively more productive than they are in steady state; and (ii) the rise in the abundance of job seekers, relative to open positions, allows productive firms—especially those that are financially unconstrained—to grow faster.

In an extension of the model, we augment the composition effect with a sectoral component by allowing permanent heterogeneity in recruiting technologies across industries. As Davis, Faberman, and Haltiwanger (2013) document, construction and a few other sectors stand out in terms of their frictional characteristics by systematically displaying higher than average vacancy filling rates. In addition, these are the industries that were hit hardest by the crisis. In agreement with Davis, Faberman, and Haltiwanger (2012b), our measurement exercise concludes that, in the context of the Great Recession, the shift in the composition of labor demand away from these high-yield sectors played a nontrivial role in the decline of aggregate recruiting intensity.

We argue that our taxonomy of slackness and composition channels is useful for three reasons. First, it offers a useful heuristic lens for thinking about the complex—sometimes offsetting—firm-level forces that determine the dynamics of aggregate hiring in response to a macroeconomic shock. Second, when the slackness effect is dominant, as we conclude, firms' recruiting efforts are very responsive to the availability of job seekers relative to vacant positions in the labor market. This, in turn, implies that any aggregate impulse that reduces labor market tightness will have an amplified impact on the aggregate job-finding rate through firms' recruiting intensity decisions. Third, a strong slackness effect also implies that any policy intervention directed at raising unemployed workers' search effort with the aim of accelerating their reentry into the employment ranks will, by lowering aggregate tightness, reduce firms' recruiting effort, thus mitigating the original intent of the policy. By the same logic, the endogenous response of recruiting effort to labor market tightness reinforces the direct impact of subsidies targeted to hiring firms.

To the best of our knowledge, only two other papers have developed models of recruiting intensity. Leduc and Liu (2017) extend a standard Diamond-Mortensen-Pissarides model to one in which a representative firm chooses search intensity per vacancy. Without firm heterogeneity, they are unable to speak to the cross-sectional empirical evidence that recruiting intensity is tightly linked to firm growth rates, a key observation that we use to discipline our framework and assess the magnitude of the composition effect. Kaas and Kircher (2015) is the only other paper that focuses on heterogeneous job-filling rates across firms. In their directed search environment, different firms post distinct wages that attract job seekers at differential rates, whereas we study how firms' costly recruiting activities determine differential job-filling rates. One would expect both factors to be important determinants of the ability of firms to grow rapidly. For example, from Austrian data, Kettemann, Mueller, and Zweimuller (2016) document that job-filling rates are higher at high-paying firms. However, after controlling for the firm component of wages, they remain increasing in firms' growth rates, implying that wages are not the whole story: employers use other instruments besides wages to hire quickly.

Moreover, while they (and Leduc and Liu, 2017) study aggregate productivity shocks—as we do as well—we further analyze financial shocks, a more natural choice if one's attention is on the Great Recession. Finally, while aggregate recruiting intensity drops after a negative aggregate shock in both our model and theirs, the reasons for the drop fundamentally differ. Kaas and Kircher (2015) argue that the drop depends on recruiting intensity being a concave function of firms' hiring policies, whose dispersion across firms increases after a negative shock. Our decomposition of the transmission mechanism linking macroeconomic shocks and aggregate recruiting intensity allows us to infer that the main source of the drop is the increase in the number of available job seekers per vacancy, which allows firms to scale back their recruiting effort.

The rest of the paper is organized as follows. Section 2 formalizes the nexus between firm-level recruiting intensity and aggregate match efficiency. Section 3 outlines the model economy and the stationary equilibrium. Section 4 describes the parameterization of the model and highlights some cross-sectional features of the economy. Section 5 describes the dynamic response of the economy to macroeconomic shocks, explains the transmission mechanism, and outlines the main results of the paper. Section 6 examines the robustness of our main findings. Section 7 concludes.

2 Recruiting Intensity and Aggregate Matching Efficiency

We briefly describe how we can aggregate hiring decisions at the firm level into an economy-wide matching function with an efficiency factor that has the interpretation of average recruiting intensity. This derivation follows DFH. For much of the paper, we abstract from quits and search on the job, and thus in our baseline model there is no role for replacement hiring: gross hires always equal the net growth of expanding firms. We discuss the implications of this assumption in Section 6.

At date t, any given hiring firm i chooses v_{it}, the number of open positions ready to be staffed and costly to create, as well as e_{it}, an indicator of recruiting intensity. Let $v^*_{it} = e_{it}v_{it}$ be the number of *effective* vacancies in firm i. Integrating over all firms, we obtain

$$V_t^* = \int e_{it} v_{it} di, \tag{1}$$

the aggregate number of effective vacancies. Under our maintained assumption of a constant returns to scale Cobb-Douglas matching function, aggregate hires equal

$$H_t = (V_t^*)^\alpha U_t^{1-\alpha} = \Phi_t V_t^\alpha U_t^{1-\alpha}, \text{ with } \Phi_t = \left(\frac{V_t^*}{V_t}\right)^\alpha = \left[\int e_{it} \left(\frac{v_{it}}{V_t}\right) di\right]^\alpha, \tag{2}$$

which corresponds to DFH's generalized matching function. Therefore, measured aggregate matching efficiency Φ_t is an average of firm-level recruiting intensity weighted by individual vacancy shares, raised to the power of α, the economy-wide elasticity of hires to vacancies. Finally, consistency requires that each firm i faces hiring frictions, implying that

$$h_{it} = q\left(\theta_t^*\right) e_{it} v_{it}, \tag{3}$$

where $\theta_t^* = V_t^*/U_t$ is effective market tightness.[4] Thus, $q\left(\theta_t^*\right) = H_t/V_t^* = \left(\theta_t^*\right)^{\alpha-1}$ is the aggregate job-filling rate per effective vacancy, constant across all firms at date t.

3 Model

Our starting point is an equilibrium random-matching model of the labor market in which firms are heterogeneous in productivity and size, and the hiring process occurs through an aggregate matching function. As discussed in the introduction, we augment this model in three dimensions—all of which are essential to developing a framework that can address our question. First, our framework features endogenous firm entry and exit. Second, beyond the number of positions to open (vacancies), hiring firms optimally choose their recruiting intensity: by spending more on recruitment resources, they can increase the rate at which they meet job seekers. Third, once in existence, firms face financial constraints.

In what follows, we present the economic environment in detail, outline the model timing,

[4] Throughout, we are faithful to the notation in this literature and denote *measured* labor market tightness V_t/U_t as θ_t.

and then describe the firm, bank, and household problems. Finally, we define a stationary equilibrium for the aggregate economy. Since our experiments will consist of perfect foresight transition dynamics, we do not make reference to aggregate state variables in agents' problems. We use a recursive formulation throughout.

3.1 Environment

Time is discrete and the horizon is infinite. Three types of agents populate the economy: firms, banks, and households.

Firms. There is an exogenous measure λ_0 of potential entrants each period, and an endogenous measure λ of incumbent firms. Firms are heterogeneous in their productivity $z \in Z$, stochastic and i.i.d. across all firms, and operate a decreasing returns to scale (DRS) production technology $y(z, n', k)$ that uses inputs of labor $n' \in N$ and capital $k \in K$. The output of production is a homogeneous final good, whose competitive price is the numeraire of the economy.

All potential entrants receive an initial equity injection a_0 from households. Next, they draw a value of z from the initial distribution $\Gamma_0(z)$ and, conditional on this draw, decide whether to enter and become an incumbent by paying the setup cost χ_0. Those that do not enter return the initial equity to the households.[5] This is the only time when firms can obtain funds directly from households. Throughout the rest of their life cycle they must rely on debt issuance.

Incumbents can exit exogenously or endogenously. With probability ζ, a destruction shock hits an incumbent firm, forcing it to exit. Surviving firms observe their new value of z, drawn from the conditional distribution $\Gamma(dz', z)$, and choose whether to exit or continue production. Under either exogenous or endogenous exit, the firm pays out its positive net worth a to households. Those incumbents that decide to stay in the industry pay a per-period operating cost χ and then choose labor and capital inputs.

The labor decision involves either firing some existing employees or hiring new workers. Firing is frictionless, but hiring is not: a hiring firm chooses both vacancies v and recruitment effort e with associated hiring cost $\mathcal{C}(e, v, n)$, which also depends on initial employment. Given (e, v), the individual hiring function (3) determines current period employment n' used in pro-

[5]Without loss of generality, we could have assumed that a fraction of the initial equity is sunk to develop the blueprint, i.e., attain the draw of z, and in case of no entry, only the remaining fraction is returned to the financier.

duction. To simplify wage setting, we assume that firms' owners make take-it-or-leave-it offers to workers, so the wage rate equals ω, the individual flow value from nonemployment.

Firms face two financial constraints. First, the capital decision involves borrowing capital from financial intermediaries (banks) in intraperiod loans. Because of imperfect contractual enforcement frictions, firms can appropriate a fraction $1/\varphi$ of the capital received by banks, with $\varphi > 1$. To preempt this behavior, a firm renting k units of capital is required to deposit k/φ units of their net worth with the bank. This guarantees that, ex post, the firm does not have an incentive to abscond with the capital. Thus, a firm with current net worth a faces a collateral constraint $k \leq \varphi a$. This model of financial frictions is based on Evans and Jovanovic (1989). Second, as mentioned above, we assume that firms may only issue equity upon entry: an incumbent must keep nonnegative dividends.

The model requires both constraints. Without the equity (nonnegative dividend) constraint, firms can arbitrarily obtain funds from households. The collateral constraint will still impose a maximum ratio of k to a, but a can increase freely through raised equity $(d < 0)$, so k is in effect unconstrained. Without the collateral constraint, firms can arbitrarily increase k through debt while keeping dividends nonnegative. In both cases, the only limit is determined by the exit option (i.e., a negative continuation value).

Banks. The banking sector is perfectly competitive. Banks receive household deposits, freely transform them into capital, and rent it to firms. The one-period contract with households pays a risk-free interest rate of r. Capital depreciates at rate δ in production, and so the price of capital charged by banks to firms is $(r + \delta)$.

Households. We envision a representative household with $\bar{L}$ family members, U of which are unemployed. The household is risk-neutral with discount factor $\beta \in (0, 1)$. It trades shares M of a mutual fund comprising all firms in the economy and makes bank deposits T. It earns interest on deposits, the total wage payments that firms make to employed family members, and D dividends per share held in the mutual fund. Moreover, unemployed workers produce ω units of the final good at home. Household consumption is denoted by C.

Before describing the firm's problem in detail, we outline the precise timing of the model, summarized in Figure 3. Within a period, the events unfold as follows: (i) realization of the productivity shocks for incumbent firms; (ii) endogenous and exogenous exit of incumbents;

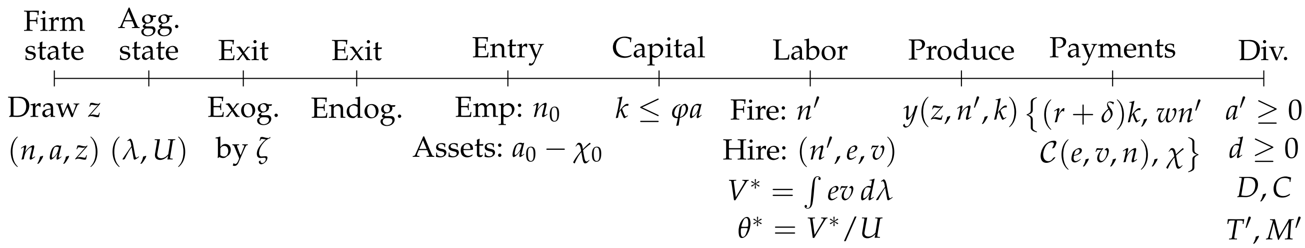

Figure 3: Timeline of the model

(iii) realization of initial productivity and entry decision of potential entrants; (iv) borrowing decisions by incumbents; (v) hiring/firing decisions and labor market matching; (vi) production and revenues from sales; (vii) payment of wage bill, costs of capital, hiring, and operation expenses, firm dividend payment/saving decisions, and household consumption/saving decisions.

To be consistent with our transition dynamics experiments in Section 5, it is useful to note that we record aggregate state variables—the measures of incumbent firms λ and unemployment U—at the beginning of the period, between stages (i) and (ii). Moreover, even though the labor market opens after firms exit or fire, workers who separate in the current period can only start searching next period.

3.2 Firm Problem

We first consider the entry and exit decisions, then analyze the problem of incumbent firms.

Entry. A potential entrant who has drawn z from $\Gamma_0(z)$ solves the following problem:

$$\max \left\{ a_0 \, , \, \mathbb{V}^i (n_0, a_0 - \chi_0, z) \right\}, \tag{4}$$

where $\mathbb{V}^i$ is the value of an incumbent firm, a function of (n, a, z). The firm enters if the value to the risk-neutral shareholder of becoming an incumbent with one employee ($n_0 = 1$), initial net worth equal to the household equity injection a_0 minus the entry cost χ_0, and productivity z exceeds the value of returning a_0 to the household. Let $i(z) \in \{0, 1\}$ denote the entry decision rule, which depends only on the initial productivity draw, since all potential entrants share the same entry cost, initial employment and ex ante equity injection. As $\mathbb{V}^i$ is increasing in z, there is an endogenous productivity cutoff z^* such that for all $z \geq z^*$, the firm chooses to enter. The measure of entrants is therefore

$$\lambda_e = \lambda_0 \int_Z i(z) d\Gamma_0 = \lambda_0 \Big[1 - \Gamma_0(z^*) \Big]. \tag{5}$$

Exit. Firms exit exogenously with probability ζ. Conditional on survival the firm then chooses to continue or exit. An exiting firm pays out its net worth a to shareholders. The firm's expected value $\mathbb{V}$ before the destruction shock equals

$$\mathbb{V}\big(n, a, z\big) = \zeta a + (1 - \zeta) \max \Big\{ \, \mathbb{V}^i\big(n, a, z\big), \, a \, \Big\}. \tag{6}$$

We denote by $x\,(n, a, z) \in \{0, 1\}$ the exit decision.

Hire or Fire. An incumbent firm i with employment, assets, and productivity equal to the triplet (n, a, z) chooses whether to hire or fire workers to solve

$$\mathbb{V}^i\big(n, a, z\big) = \max \Big\{ \mathbb{V}^h\big(n, a, z\big), \mathbb{V}^f\big(n, a, z\big) \Big\}. \tag{7}$$

The two value functions $\mathbb{V}^f$ and $\mathbb{V}^h$ associated with firing (f) and hiring (h) are described below.

The Firing Firm. A firm that has chosen to fire some of its workers (or to not adjust its work force) solves

$$\begin{aligned} \mathbb{V}^f\big(n, a, z\big) &= \max_{n', k, d} \; d + \beta \int_Z \mathbb{V}\big(n', a', z'\big) \Gamma(dz', z) \\ & s.t. \\ n' &\leq n, \\ d + a' &= y\big(n', k, z\big) + (1 + r)a - \omega n' - (r + \delta)k - \chi, \\ k &\leq \varphi a, \\ d &\geq 0. \end{aligned} \tag{8}$$

Firms maximize shareholder value and, because of risk neutrality, use β as their discount factor. The change in net worth $a' - a$ is given by revenues from production and interest on savings net of the wage bill, rental and operating costs, and dividend payouts d. The last two equations in (8) reiterate that firms face a collateral constraint on the maximum amount of capital they can rent and a nonnegativity constraint on dividends.

To help understand the budget constraint and preface how we take the model to the data, define firm debt by the identity $b \equiv k - a$, with the understanding that $b < 0$ denotes savings. Making this substitution reveals an alternative formulation of the model in which the firm owns its capital and faces a constraint on leverage. With state vector (n, k, b, z), the firm faces the following budget and collateral constraints:

$$\begin{aligned} d + \underbrace{k' - (1-\delta)k}_{\text{Investment}} &= \underbrace{y(n', k, z) - \omega n' - \chi - rb}_{\text{Operating Profit}} + \underbrace{b' - b}_{\Delta \text{ Borrowing}}, \\ b/k &\leq (\varphi - 1)/\varphi. \end{aligned}$$

This makes it clear that the firm can fund equity payouts and investment in capital through either operating profits or expanding borrowing/reducing saving.

The Hiring Firm. The hiring firm additionally chooses the number of vacancies to post $v \in \mathbb{R}_+$ and recruitment effort $e \in \mathbb{R}_+$, understanding that, by a law of large numbers, its new hires $n' - n$ equal the firm's job-filling rate qe of each of its vacancies times the number of vacancies v created: $n' - n = q(\theta^*)ev$.[6] Note that the individual firm job-filling rate depends on the aggregate meeting rate q, which is determined in equilibrium and the firm takes as given, as well as on its recruiting effort e. The firm faces a variable cost function $\mathcal{C}(e, v, n)$, increasing and convex in e and v.

A firm's continuation value depends on n', not on the mix of recruiting intensity e and vacant positions v that generates it. As a result, one can split the problem of the hiring firm into two stages. The first stage is the choice of n', k, and d. The second stage, given n', is the choice of the optimal combination of inputs (e, v). The latter reduces to a static cost minimization problem:

$$\begin{aligned} \mathcal{C}^*\left(n, n'\right) &= \min_{e,v} \mathcal{C}(e, v, n) \\ s.t. \quad e \geq 0, \quad v &\geq 0, \quad n' - n = q(\theta^*)ev, \end{aligned} \tag{9}$$

[6]The linearity of the individual hiring function in vacancies is one of the key empirical findings of DFH.

yielding the lowest cost combination $e(n, n')$ and $v(n, n')$ that delivers $h = n' - n$ hires to a firm of size n, and the implied cost function $\mathcal{C}^*(n, n')$.

The remaining choices of n', k, and d require solving the dynamic problem

$$
\begin{aligned}
\mathbb{V}^h(n, a, z) &= \max_{n', k, d} \; d + \beta \int_Z \mathbb{V}(n', a', z') \Gamma(dz', z) \\
&\text{s.t.} \\
n' &> n, \\
d + a' &= y(n', k, z) + (1+r)a - \omega n' - (r+\delta)k - \chi - \mathcal{C}^*(n, n'), \\
k &\leq \varphi a, \\
d &\geq 0.
\end{aligned}
\tag{10}
$$

The solution to this problem includes the decision rule $n'(n, a, z)$. Using this function in the solution to (9), we obtain decision rules $e(n, a, z)$ and $v(n, a, z)$ for recruitment effort and vacancies in terms of firm state variables.

Given the centrality of the hiring cost function $\mathcal{C}(e, v, n)$ to our analysis, we now discuss its specification. In what follows, we choose the functional form

$$
\mathcal{C}(e, v, n) = \left[\frac{\kappa_1}{\gamma_1} e^{\gamma_1} + \frac{\kappa_2}{\gamma_2 + 1} \left(\frac{v}{n} \right)^{\gamma_2} \right] v, \tag{11}
$$

with $\gamma_1 \geq 1$ and $\gamma_2 \geq 0$ being necessary conditions for the convexity of the maximization problem (9). This cost function implies that the average cost of a vacancy, $\mathcal{C}/v$, has two separate components. The first is increasing and convex in recruiting intensity per vacancy e. The idea is that, for any given open position, the firm can choose to spend resources on recruitment activities (recall Figure 2) to make the position more visible or the firm more attractive as a potential employer, or to assess more candidates per unit of time, but all such activities are increasingly costly on a per-vacancy basis. The second component is increasing and convex in the vacancy rate and captures the fact that expanding productive capacity is costly in *relative* terms: for example, creating 10 new positions involves a more expensive reorganization of production in a firm with 10 employees than in a firm with 1,000 employees.

In Appendix A we derive several results for the static hiring problem of the firm (9) under this cost function and derive the exact expression for $\mathcal{C}^*(n, n')$ used in the dynamic problem

(10). We show that, by combining first-order conditions, we obtain the optimal choice for e:

$$e\left(n,n'\right) \;=\; \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1-1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{-\frac{\gamma_2}{\gamma_1+\gamma_2}}\left(\frac{n'-n}{n}\right)^{\frac{\gamma_2}{\gamma_1+\gamma_2}}, \tag{12}$$

and, hence, the firm-level job-filling rate $f\left(n,n'\right) \equiv q\left(\theta^*\right) e\left(n,n'\right)$, as well as the optimal vacancy rate:

$$\frac{v}{n}=\left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1-1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{-\frac{\gamma_1}{\gamma_1+\gamma_2}}\left(\frac{n'-n}{n}\right)^{\frac{\gamma_1}{\gamma_1+\gamma_2}}. \tag{13}$$

Equation (12) demonstrates that the model implies a log-linear relation between the job-filling rate and employment growth at the firm level, with elasticity $\gamma_2/(\gamma_1+\gamma_2)$. This is the key empirical finding of DFH, who estimate this elasticity to be 0.82. In fact, one could interpret our functional choice for $\mathcal{C}$ in equation (11) as a reverse-engineering strategy in order to obtain, from first principles, the empirical cross-sectional relation between the establishment-level job-filling rate and the establishment-level hiring rate uncovered by DFH. Put differently, microdata sharply discipline the recruiting cost function of the model.[7]

Why does firm optimality imply that the job-filling rate increases with the growth rate with elasticity $\gamma_2/(\gamma_1+\gamma_2)$? Recruiting intensity e and the vacancy rate (v/n) are substitutes in the production of a target employment growth rate $\left(n'-n\right)/n$ (see the last equation in (9)). Thus, a firm that wants to grow faster than another will optimally create more positions and, at the same time, spend more in recruiting effort. However, the stronger the convexity of $\mathcal{C}$ in the vacancy rate (γ_2), relative to its degree of convexity in effort (γ_1), the more an expanding firm finds it optimal to substitute away from vacancies into recruiting intensity to realize its target growth rate. In the special case when $\gamma_2=0$, all the adjustment occurs through vacancies, and recruiting effort is irresponsive to the growth rate and to macroeconomic conditions, as in the canonical model of Pissarides (2000).

Figure 4 plots the cross-sectional relationship between the vacancy rate and employment growth (panel A) and the job-filling rate and employment growth (panel B) in the model and

[7] Appendix A also shows that, once the optimal choice of e is substituted into (11), $\mathcal{C}$ can be stated solely in terms of the vacancy rate and becomes equivalent to one of the hiring cost functions that Kaas and Kircher (2015) use in their empirical analysis.

Figure 4: Cross-sectional relationships between monthly employment growth $(n' - n)/n$ and the vacancy rate v/n and the job-filling rate eq. Data from DFH online supplemental materials.

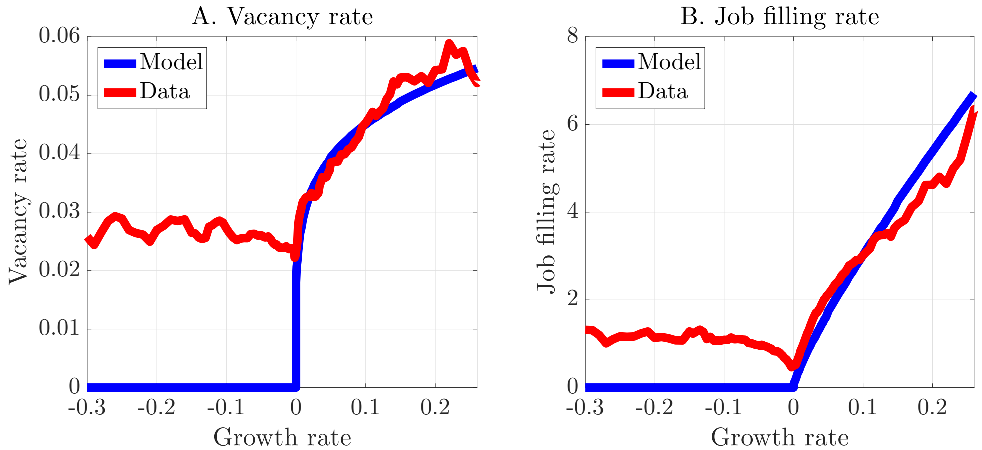

in the DFH data, with the elasticity of the job-filling rate to firm growth $\gamma_2/(\gamma_1 + \gamma_2) = 0.82$.[8] Since the individual hiring function is linear in vacancies, the elasticity of the vacancy rate to firm growth equals $\gamma_1/(\gamma_1 + \gamma_2) = 0.18$.

3.3 Household Problem

The representative household solves

$$\begin{aligned} \mathbb{W}(T, M, D) &= \max_{T', M', C > 0} \; C + \beta \mathbb{W}(T', M', D') \qquad (14) \\ & s.t. \\ C + \bar{Q}T' + PM' &= \omega \bar{L} + (D + P)M + T, \end{aligned}$$

where T are bank deposits, M are shares of the mutual fund composed of all firms in the economy, and D are aggregate dividends per share.[9] The household takes as given the price of bank deposits $\bar{Q}$, the share price P, and the price of the final good, which we normalize to one. From the first-order conditions for deposits and share holdings, we obtain $\bar{Q} = \beta$ and $P = \beta\,(P + D)$ which imply a time-invariant rate of return of $r = \beta^{-1} - 1$ on both deposits and shares. The

[8] In Figure 4, the model implies zero hires for firms with negative growth rates, whereas in the data time aggregation and replacement hires lead to positive vacancy rates and vacancy yields for shrinking firms as well. We return on this point in Section 6.

[9] The initial equity injections into successful start-ups are treated as negative dividends, i.e. they are part of D every period.

household is therefore indifferent over portfolios.

Since firms make take-it-or-leave-it offers to workers (i.e., firms have all the bargaining power) and are competitive, they pay all their workers a wage equal to the individual's flow value of nonemployment ω, which we interpret as output from home production. The total amount of resources available to households for consumption and saving as a result of market and home production is thus simply $\omega\bar{L}$.[10] Because of risk neutrality, the household is indifferent over the timing of consumption.

3.4 Stationary Equilibrium and Aggregation

Let Σ_N, Σ_A, and Σ_Z be the Borel sigma algebras over N and A, and Z. The state space for an incumbent firm is $S = N \times A \times Z$, and we denote with s one of its elements (n, a, z). Let Σ_S be the sigma algebra on the state space, with typical set $\mathcal{S} = \mathcal{N} \times \mathcal{A} \times \mathcal{Z}$, and (S, Σ_S) be the corresponding measurable space. Denote with $\lambda : \Sigma_S \to [0, \infty)$ the stationary measure of incumbent firms at the beginning of the period, following the draw of firm-level productivity, before the exogenous exit shock.

To simplify the exposition of the equilibrium, it is convenient to use $s \equiv (n, a, z)$ and $s_0 \equiv (n_0, a_0 - \chi_0, z)$ as the argument for incumbents' and entrants' decision rules.

A stationary recursive competitive equilibrium is a collection of firms' decision rules $\{i(z), x(s), n'(s), e(s), v(s), a'(s), d(s), k(s)\}$, value functions $\{\mathbb{V}, \mathbb{V}^i, \mathbb{V}^f, \mathbb{V}^h\}$, a measure of entrants λ_e, share price P and aggregate dividends D, wage ω, a distribution of firms λ, and a value for effective labor market tightness θ^* such that: (i) the decision rules solve the firm's problems (4)-(10), $\{\mathbb{V}, \mathbb{V}^i, \mathbb{V}^f, \mathbb{V}^h\}$ are the associated value functions, and λ_e is the mass of entrants implied by (5); (ii) the market for shares clears at $M = 1$ with share price

$$P = \int_S \mathbb{V}(s)\, d\lambda + \lambda_0 \int_Z i(z)\, \mathbb{V}^i(s_0)\, d\Gamma_0$$

[10] If we let the wage be W, then total resources from market and home production equal $W(\bar{L} - U) + \omega U$. The term $\omega\bar{L}$ in the household budget constraint follows from the fact that $W = \omega$. This also explains why unemployment U is not a state variable in the household's problem (14).

and aggregate dividends

$$D=\zeta\int_{S}ad\lambda+(1-\zeta)\int_{S}\{[1-x(s)]\,d\,(s)+x\,(s)\,a\}\,d\lambda-\lambda_0\int_{Z}i\,(z)\,a_0d\Gamma_0;$$

(iii) the stationary distribution λ is the fixed point of the recursion:

$$\begin{aligned}\lambda(\mathcal{N}\times\mathcal{A}\times\mathcal{Z}) \;=\; &(1-\zeta)\int_{S}[1-x(s)]\,\mathbf{1}_{\{n'(s)\in\mathcal{N}\}}\mathbf{1}_{\{a'(s)\in\mathcal{A}\}}\Gamma(\mathcal{Z},z)d\lambda \\ &+\lambda_0\int_{Z}i\,(z)\,\mathbf{1}_{\{n'(s_0)\in\mathcal{N}\}}\mathbf{1}_{\{a'(s_0)\in\mathcal{A}\}}\Gamma(\mathcal{Z},z)d\Gamma_0,\end{aligned} \tag{15}$$

where the first term refers to existing incumbents and the second to new entrants; (iv) effective market tightness θ^* is determined by the balanced flow condition

$$\bar{L}-N(\theta^*)=\frac{F\,(\theta^*)-\lambda_e\,(\theta^*)\,n_0}{p\,(\theta^*)}, \tag{16}$$

where $p\,(\theta^*)$ is the aggregate job-finding rate, $N(\theta^*)$ is aggregate employment

$$N\,(\theta^*)=(1-\zeta)\int_{S}[1-x(s)]n'(s)d\lambda+\lambda_0\int_{Z}i(z)n'(s_0)d\Gamma_0, \tag{17}$$

and $F(\theta^*)$ are aggregate separations

$$F\,(\theta^*)=\zeta\int_{S}nd\lambda+(1-\zeta)\int_{S}x\,(s)\,nd\lambda+(1-\zeta)\int_{S}[1-x\,(s)]\,\big(n-n'(s)\big)^{-}d\lambda, \tag{18}$$

which include all employment losses from firms exiting exogenously and endogenously, plus all the workers fired by shrinking firms, which we have denoted by $(n-n'\,(s))^{-}$.[11] In equations (16)-(18), the dependence of λ_e, N, and F on θ^* comes through the decision rules and the stationary distribution, even though, for notational ease, we have omitted θ^* as their explicit argument.

The left-hand side of (16) is the definition of unemployment—labor force minus employment—whereas the right-hand side is the steady state Beveridge curve, i.e., the law

[11] Entrant firms never fire, as they enter with the lowest value on the support for N, n_0 normalized to 1.

of motion for unemployment

$$U' = U - p(\theta^*)\,U + F(\theta^*) - \lambda_e(\theta^*) n_0 \tag{19}$$

evaluated in steady state. As in Elsby and Michaels (2013), the two sides of (16) are independent equations determining the same variable—unemployment—and combined they yield equilibrium market tightness θ^*.[12] Note that equations (16) and (19) account for the fact that every new firm enters with n_0 workers hired "outside" the frictional labor market (e.g., the firm founders).

Clearly, once θ^* and λ are determined, so is U from either side of (16) and, therefore, V^*. Finally, we note that measured aggregate matching efficiency, in equilibrium, is $\Phi = (V^*/V)^{\alpha}$, where measured and effective vacancies are respectively

$$\begin{aligned} V &= (1-\zeta)\int_S [1-x(s)]\,v(s)\,d\lambda \qquad + \lambda_0 \int_Z i(z)\,v(s_0)\,d\Gamma_0, \\ V^* &= (1-\zeta)\int_S [1-x(s)]\,e(s)\,v(s)\,d\lambda + \lambda_0 \int_Z i(z)\,e(s_0)\,v(s_0)\,d\Gamma_0. \end{aligned}$$

Appendix C provides details on the computation of the decision rules and the stationary equilibrium.

4 Parameterization

We begin with the subset of parameters calibrated externally, then consider those estimated within the model. The main problem we face in parameterizing the model is that the theory does not distinguish between firms and establishments. Ideally we would only use data on firms, since financial constraints apply at the firm level. However JOLTS data is only available by establishment, as are other data sources we use in calibration. We are therefore forced to compromise: we use firm data whenever we have a choice—for example, when we use the

[12] Our computation showed that, typically, $N(\theta^*)$ is decreasing in its argument and the right-hand side of (16) is always positive and decreasing. Thus, the crossing point of the left- and right-hand sides is unique, when it exists. However, an equilibrium may not exist. For example, for very low hiring costs, $N(\theta^*)$ may be greater than $\bar{L}$. Conversely, for large enough operating or hiring costs, no firms will enter the economy. In this case, there is no equilibrium with market production (albeit there is always some home production in the economy).

Table 1: Externally set parameter values

Parameter		Value	Target	Value
Discount factor (monthly)	β	0.9967	Annual risk-free rate	0.04
Mass of potential entrants	λ_0	0.02	Measure of incumbents	1
Size of labor force	$\bar{L}$	24.6	Average firm size (BDS)	23
Elasticity of matching function wrt V_t	α	0.5	JOLTS	

Business Dynamics Statistics (BDS) data—and establishment data when we are limited. Data moments are averages over 2001-2007 unless otherwise specified.

4.1 Externally Calibrated

The model period is one month. We set β to replicate an annualized risk-free rate of 4 percent. Since the measure of potential entrants λ_0 scales λ—see equation (15)—we choose λ_0 to normalize the total measure of incumbent firms to one. We then fix the size of the labor force $\bar{L}$ so that, given a measure one of firms, and the model implied steady-state unemployment rate of 7 percent, the average firm size will be 23 (BDS).[13] In line with empirical studies, we set α, the elasticity of aggregate hires to aggregate vacancies in the matching function, to 0.5. Table 1 summarizes these parameter values.

4.2 Internally Calibrated

Table 2 lists the remaining 19 parameters of the model that are set by minimizing the distance between an equal number of empirical moments and their equilibrium counterparts in the model.[14] It also lists the targeted moments, their empirical values, and their simulated values from the model. Even though every targeted moment is determined simultaneously by all parameters, in what follows we discuss each of them in relation to the parameter for which,

[13]The unemployment rate is $u = \bar{L}/N(\theta^*) - 1$, and with a unit mass of firms the average firm size is simply $N(\theta^*)$. Hence for an unemployment rate of $u = 0.07$, $\bar{L}$ determines average firm size.

[14]Specifically, the vector of parameters Ψ is chosen to minimize the minimum-distance-estimator criterion function

$$f(\Psi) = \left(\mathbf{m}_{data} - \mathbf{m}_{model}(\Psi)\right)' \mathbf{W} \left(\mathbf{m}_{data} - \mathbf{m}_{model}(\Psi)\right)$$

where $\mathbf{m}_{data}$ and $\mathbf{m}_{model}(\Psi)$ are the vectors of moments in the data and model, and $\mathbf{W} = \text{diag}\left(1/\mathbf{m}^2_{data}\right)$ is a diagonal weighting matrix.

Table 2: Parameter values estimated internally

Parameter		Value	Target	Model	Data
Flow of home production	ω	1.000	Monthly separation rate	0.033	0.030
Scaling of match. funct.	$\bar{\Phi}$	0.208	Monthly job-finding rate	0.411	0.400
Prod. weight on labor	ν	0.804	Labor share	0.627	0.640
Midpoint DRS in prod.	σ_M	0.800	Employment share $n < 50$	0.294	0.306
High-Low spread in DRS	$\Delta\sigma$	0.094	Employment share $n \geq 500$	0.430	0.470
Mass - Low DRS	μ_L	0.826	Firm share $n < 50$	0.955	0.956
Mass - High DRS	μ_H	0.032	Firm share $n \geq 500$	0.004	0.004
Std. dev. of z shocks	ϑ_z	0.052	Std. dev. ann. emp. growth	0.440	0.420
Persistence of z shocks	ρ_z	0.992	Mean Q4 emp. / Mean Q1 emp.	75.16	76.00
Mean $z_0 \sim Exp(\bar{z}_0^{-1})$	$\bar{z}_0$	0.390	$\Delta \log z$: Young vs. Mature	-0.246	-0.353
Cost elasticity wrt e	γ_1	1.114	Elasticity of vac. yield wrt g	0.814	0.820
Cost elasticity wrt v	γ_2	4.599	Ratio vac. yield: $n < 50/n \geq 50$	1.136	1.440
Cost shifter wrt e	κ_1	0.101	Hiring cost (100+) / wage	0.935	0.927
Cost shifter wrt v	κ_2	5.000	Vacancy share $n < 50$	0.350	0.370
Exogenous exit probability	ζ	0.006	Five year survival rate	0.497	0.500
Entry cost	χ_0	9.354	Annual entry rate	0.099	0.102
Operating cost	χ	0.035	Share of job destruction by exit	0.210	0.340
Initial wealth	a_0	10.000	Start-up Debt to Output	1.361	1.280
Collateral constraint	φ	10.210	Aggregate Debt to Assets	0.280	0.350

intuitively, that moment yields the most identification power.

We set the flow of home production of the unemployed ω to replicate a monthly separation rate of 0.03. We choose the shift parameter of the matching function (a normalization of the value of Φ in steady state) in order to replicate a monthly job-finding rate of 0.40. Together, these two moments yield a steady state unemployment rate of 0.07.

We assume a revenue function $y(z, n', k) = z\left[(n')^{\nu} k^{1-\nu}\right]^{\sigma}$. We need not take a stand on whether z represents demand or productivity shocks, or whether $\sigma < 1$ is due to DRS in production or downward-sloping demand.[15] For simplicity, we will refer to the revenue function as if it were a production function: σ represents the span of control and z is total factor productivity.

[15]Given our class of frictions, the revenue function is sufficient. This would not be the case in alternative environments that endogenize components of revenue productivity, for example models with R&D—which affects productivity—or models with customer accumulation—which affects demand.

We introduce a small degree of permanent heterogeneity in the scale parameter σ. Specifically, we consider a three-point distribution with support $\{\sigma_L, \sigma_M, \sigma_H\}$—symmetric about σ_M—leaving four unknown parameters: (i) the value of σ_M; (ii) the spread $\Delta\sigma \equiv (\sigma_H - \sigma_L)$; and (iii)-(iv) the fractions of low and high DRS firms μ_L, μ_H. This heterogeneity allows us to match the skewed firm size distribution, with the parameters chosen to match the shares of total employment and total firms due to firms of size 0-49 and 500+ (BDS). Permanent heterogeneity in productivity could also be used to match these facts, but heterogeneity in σ also generates small old firms alongside young large firms, thus decoupling age and size, which tend to be too strongly correlated in standard firm dynamics models with mean reverting productivity.[16] In other words, heterogeneity in σ captures the appealing idea that there exist some very productive businesses that are small simply because the optimal scale of production for many goods or services is small. This idea will turn out to be important for interpreting the response of firms to a macroeconomic shock.

Firm productivity z follows an AR(1) process in logs: $\log z' = \log Z + \rho_z \log z + \varepsilon$, with $\varepsilon \sim \mathcal{N}(-\vartheta_z^2/2, \vartheta_z)$. We calibrate ρ_z and ϑ_z to match two measures of employment dispersion, one in growth and one in levels: the standard deviation of annual employment growth for continuing establishments in the US Census Bureau's Longitudinal Business Database (Elsby and Michaels, 2013) and the ratio of the mean size of the fourth to first quartile of the firm distribution (Haltiwanger, 2011).[17]

The initial productivity distribution for entrants Γ_0 is exponential. The mean $\bar{z}_0$ is chosen to match the revenue productivity gap between entrants and incumbents, specifically the differential between plants younger than age 1 and older than age 10 (Foster, Haltiwanger, and Syverson, 2016).

We now turn to hiring costs. The cost function (11) has four parameters: the two elasticities (γ_1, γ_2) and the two cost shifters (κ_1, κ_2). From our discussion of equations (11) and (12), recall that the cross-sectional elasticity of the job-filling rate to employment growth, estimated to be 0.82 by DFH, is a function of the ratio of these two elasticities.[18] The second moment used to

[16] See Elsby and Michaels (2013) and Kaas and Kircher (2015) for examples of the use of heterogeneity in permanent productivity.

[17] In the numerical solution and simulation of the model, z remains a continuous state variable.

[18] We cannot map $\gamma_2/(\gamma_1+\gamma_2)$ directly into this value since in DFH, and in the model's simulations for consistency, the growth rate is the Davis-Haltiwanger growth rate normalized in $[-2, 2]$. In practice, as seen in Table 2, the discrepancy between the structural and estimated parameters is very small. Moreover, DFH estimate the

separately identify the two elasticities is the ratio of vacancy yields at small ($n < 50$) and large ($n \geq 50$) establishments (JOLTS). Intuitively, when $\gamma_2 = 0$, recruiting effort is constant across firms and this ratio is one.

We use two targets to pin down the cost shift parameters. The first is the total hiring cost as a fraction of monthly wage per hire, a standard target for the single vacancy cost parameter that usually appears in vacancy posting models. We have a new source for this statistic. The consulting company Bersin and Associates runs a periodic survey of recruitment cost and practices based on over 400 firms—all with more than 100 employees. Once the firms are reweighted by industry and size, the sample is representative of this size segment of the US economy. They compute that, on average, annual spending on all recruiting activities (including internal staff compensation, university recruiting, agencies/third-party recruiters, professional networking sites, job boards, social media, contractors, employment branding services, employee referral bonuses, pay-per-click media, travel to interview candidates, applicant tracking systems, print/media/billboards, and other tools/technologies) divided by the number of hires in 2011 was \$3,479 (see Table 3 in O'Leonard 2011). Given average annual earnings of roughly \$45,000 in 2011, in the model we target a ratio of average recruiting cost to average monthly wage (in firms with more than 100 employees) of 0.928. The second target is the vacancy share of small ($n < 50$) establishments from JOLTS: κ_2 determines the size of hiring costs for small firms and, thus, the amount of vacancies they create.

The parameters χ and ζ have large effects on firm exit. The operating cost χ mostly affects the exit rates of young firms; therefore, we target the five-year firm survival rate which is approximately 50 percent (BDS). The parameter ζ contributes to the exit of large and old firms; hence, we target the fraction of total job destruction due to exit of 34 percent (BDS).[19] To pin down the setup cost χ_0, we target the annual firm entry rate of 10 percent (BDS).[20]

The remaining two parameters are the size of the initial equity injection a_0 and the collateral

relationship between the job-filling rate and the *gross hires rate* rather than employment growth. In our model, the gross hires rate and rate of employment growth of hiring firms coincide, although this would not be the case in a model with replacement hires. We discuss this in Section 6.

[19]Unlike other moments used here from the BDS, job destruction by exit is only available by establishment exit, not firm exit.

[20]When computing moments designed to be comparable to their counterparts in the BDS, we carefully time-aggregate the model to an annual frequency. For example, the entry rate in the BDS is measured as the number of age zero firms in a given year divided by the total number of firms. Computing this statistic in the model requires aggregating monthly entry and exit over 12 months. See Appendix C for details.

Table 3: Nontargeted moments

Moment	**Model**	**Data**	**Source**
Aggregate dividend / profits	0.411	0.400	NIPA
Employment share: *growth* $\in (-2.00, -0.20)$	0.070	0.076	Davis et al. (2010)
Employment share: *growth* $\in (-0.20, 0.20]$	0.828	0.848	Davis et al. (2010)
Employment share: *growth* $\in (0.20, 2.00)$	0.102	0.076	Davis et al. (2010)
Employment share: $Age \leq 1$	0.013	0.028	BDS
Employment share: $Age \in (1, 10)$	0.309	0.212	BDS
Employment share: $Age \geq 10$	0.678	0.760	BDS

Notes: *(i)* NIPA data from Tables 2.1 and 5.1 for 2006, computed as *Dividends/(Corporate profits + (1/3)×Proprietor's Income)*. *(ii)* Growth rate distribution statistics are computed from Table 2 of Davis, Faberman, Haltiwanger, and Rucker (2010), which summarizes the distribution of employment by quarterly establishment growth rates in the Business Employment Dynamics (BED) data from 2001-2006. Growth rates are computed as $g_{it} = (n_{it+1} - n_{it}) / (0.5n_{it} + 0.5n_{it+1})$. We have removed entry and exit since we already match the entry rate in the calibration so consider the distribution only over continuing establishments. Statistics from the model are also quarterly. *(iii)* BDS data are the average distribution of employment across firms from 2001-2007.

parameter φ. To inform their calibration, we target the debt-output ratio of start-up firms computed from the Kauffman Survey (Robb and Robinson, 2014) and the aggregate debt to total assets ratio from the flow of funds accounts.[21]

4.3 Cross-Sectional Implications

We now explore the main cross-sectional implications of the calibrated model, at its steady state equilibrium.

Table 3 reports some empirical moments not targeted in the calibration and their model-generated counterparts. The fact that the ratio of dividend payments to profits in the model is close to its empirical value confirms that the collateral constraint is neither too tight nor too loose. The model can also replicate well the distribution of employment by establishment growth rate and firm age, neither of which was explicitly targeted.

Figure 5 shows that the model is also able to satisfactorily replicate the observed distribution

[21] Robb and Robinson (2014) report \$68,000 of average debt (credit cards, personal and business bank loans, and credit lines) and \$53,000 of average revenue for the 2004 cohort of start-ups in their first year; see their Table 5. From the flow of funds 2005, we computed total debt as the sum of securities and loans and total assets as the sum of all nonfinancial assets plus financial assets net of trade receivables, FDIs, and miscellaneous liabilities (Tables L.103 and L.104, Liabilities of Nonfinancial Corporate and Noncorporate Business), divided by the sum of corporate and noncorporate net worth (Tables B.103 and B.104, Balance Sheet of Nonfinancial Corporate and Noncorporate Business).

Figure 5: Hire and vacancy shares by size class. Model (black) and JOLTS data, 2002-2007 (grey).

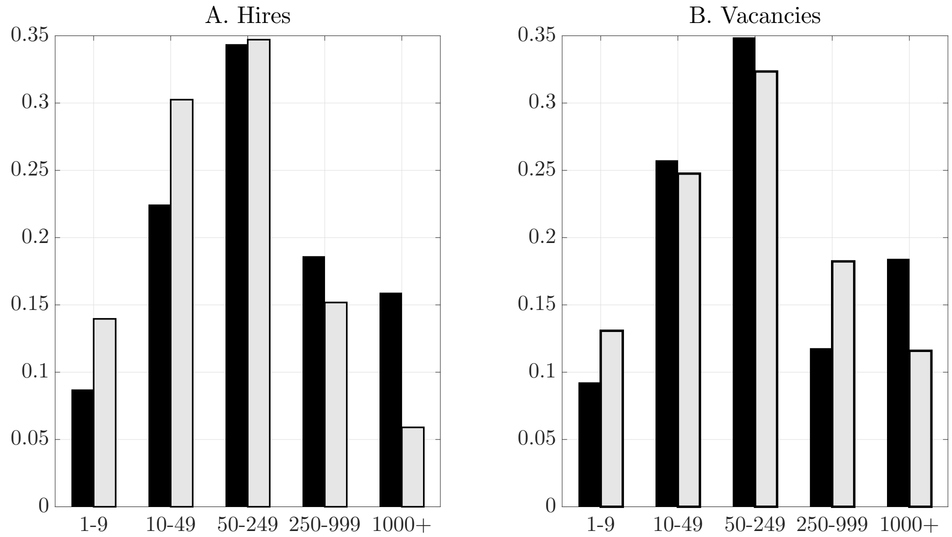

of hires and vacancies by size class (JOLTS).

In Figure 6 we plot the average firm size, job creation and destruction rates, fraction of constrained firms, and leverage (debt/saving over net worth, b/a) for firms from birth through to maturity. Panel A shows that σ_H firms, those with closer to constant returns in production, account for the upper tail in the size and growth rate distributions. On average, though, firm size grows by much less over the life cycle, since these "gazelles"—as they are often referred to in the literature—are only a small fraction of the total ($\mu_H = 0.032$). This lines up well with the data: average firm size grows by a factor of 3.0 between ages 1-5 and 20-25 in the model and 3.1 in the data (BDS). Convex recruiting costs and collateral constraints slow down growth: most firms reach their optimal size around age 10, whereas σ_H firms keep growing for much longer.

Panel B plots job creation and destruction rates by age and is a stark representation of the "up-and-out" dynamics of young firms documented in the literature (Haltiwanger, 2012). Panel C depicts the fraction of constrained firms (defined as those with $k = \varphi a$ and $d = 0$) over the life cycle. In the model, financial constraints bind only for the first few years of a firm's life, when net worth is insufficient to fund the optimal level of capital. Panel D illustrates that leverage declines with age, and after age 10 the median firm is saving (i.e., $b < 0$). Much like in the classical household "income fluctuation problem," in our model firms have a precautionary saving motive because of the simultaneous presence of three elements: (i) a concave payoff function because of DRS, (ii) stochastic productivity, and (iii) the collateral constraint.

Figure 6: Average life cycle of firms in the model

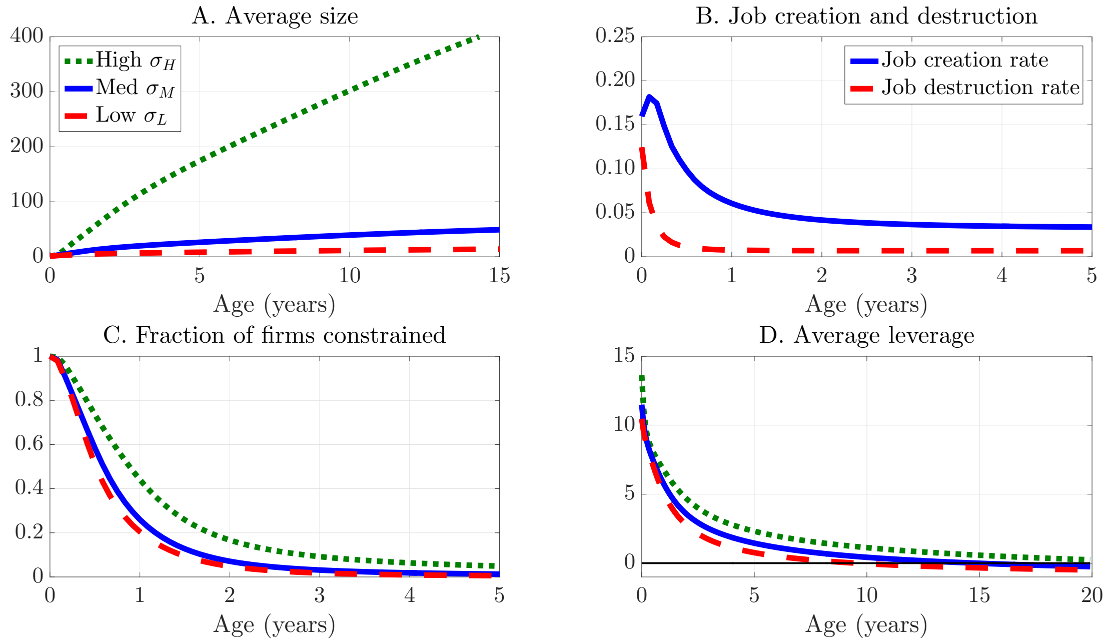

Panel A of Figure 7 shows that recruiting intensity and the vacancy rate are sharply decreasing with age. These features arise because our cost function implies that both optimal hiring effort and the vacancy rate are increasing in the growth rate, and young firms are those with the highest desired growth rates. Moreover, the stronger convexity of $\mathcal{C}$ in the vacancy rate (γ_2), relative to its degree of convexity in effort (γ_1) implies that a rapidly expanding firm prefers to increase its recruiting intensity relatively more than vacancies to realize its target growth rate. Thus, young firms find it optimal to recruit very aggressively for the new positions that they open. As firms age, growth rates fall and this force weakens.

Panel B plots the fraction of total recruiting effort, vacancies, and hiring firms by age. It shows that, relative to the steady state age distribution of hiring firms, the effort distribution is skewed toward young firms, whereas the vacancy distribution is skewed towards older firms. In the model, the age distribution of vacancies is almost uniform: young firms grow faster than old ones and, thus, post more vacancies per worker; however, they are smaller and, thus, they post fewer vacancies for a given growth rate. These two forces counteract each other and the ensuing vacancy distribution over ages is nearly flat. Figure 7 highlights that the JOLTS notion of a vacancy as "open position ready to be filled" is a good metric of hiring effort for old firms, for whom recruiting intensity is nearly constant, whereas it is quite imperfect for young firms

Figure 7: Vacancy and Effort Distributions by Age

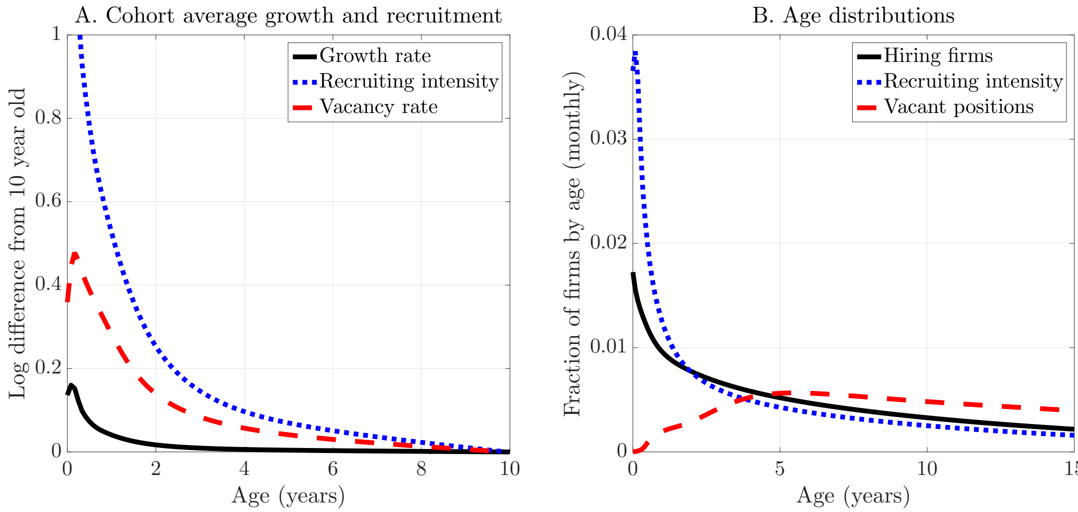

Table 4: Calibration of Aggregate Shocks

Shock		X_0	X_1	$\%\Delta X$	ρ_X	Drop in GDP at impact	$\%\Delta$ GDP after 3 years (half-life)
Productivity	Z_t	1.00	0.96	-4%	0.9764	-11.11%	-5.3%
Financial	φ_t	10.21	2.59	-74.6%	0.9903	-10.82%	-4.9%
Data (Fernald, 2014)						-9.7%	-5.1%

aged 0-5, whose average recruiting intensity, as well as its variance, are much higher than those of mature firms.[22]

5 Aggregate Recruiting Intensity and Macroeconomic Shocks

Our main experiment consists of studying the *perfect foresight* transitional dynamics of the model in response to a onetime, unexpected shock either to aggregate productivity Z or to the financial constraint parameter φ. The economy starts in steady state and the path of the shock reverts back to its initial value, so the economy also returns to its initial steady state.[23]

[22]Unfortunately, JOLTS does not report the age of the establishment, so there are no US data on vacancies and recruiting intensity by age that we can directly compare to our model. Kettemann, Mueller, and Zweimuller (2016) find that, in Austrian data, after controlling for firm fixed effects, job-filling rates are decreasing with firm age.

[23]Appendix C provides details on the solution of the model along these transitional dynamics.

5.1 Calibration of Aggregate Shocks

Let X indicate either the productivity shock or the financial shock, depending on the experiment. The path for $\{X_t\}_{t=0}^{T}$ is such that $X_0 = X_T = \bar{X}$, and $(X_t - \bar{X}) = \rho_X (X_{t-1} - \bar{X})$ for $t \in \{1, \ldots, T-1\}$, where $\bar{X}$ is the value taken in steady state. We must provide values for X_1 and ρ_X. These two values are calibrated to replicate two features of the path for aggregate output described by Fernald (2014): the peak-trough drop and its half-life.[24] First, at the trough, GDP was around 10 percent below trend. Second, GDP returned to around 5 percent below trend three years after the trough. Figure B1 in the Appendix shows the paths for output in the two experiments, which by construction are almost identical. Table 4 provides the details of this calibration exercise.[25]

5.2 Aggregate Dynamics

Figure 8 plots the dynamics of some key aggregate variables. We focus on three of the features of the data that arise in the model in response to the financial shock, but do not in response to the productivity shock.

First, the debt-output ratio drops by a magnitude that is comparable to the data and recovers with a similar persistence.[26] Second, aggregate labor productivity endogenously *increases* by 1.5 percent, close to the 2 percent increase over 2008-2010 measured by McGrattan and Prescott (2012). Tighter financial frictions prevent the expansion of firms. With DRS and firms constrained further away from their frictionless optimal size, labor productivity increases. This is especially true for fast-growing high σ firms, which have a large optimal scale of production and are most affected by the financial constraint. Third, entry declines by 24 percent, which, again, approximates its empirical counterpart of 22 percent.[27] Specifically, young-firm values

[24] See Figure 5 in Fernald (2014). Output is filtered using a biweight kernel with a bandwidth of 48 quarters.

[25] The monthly frequency of the model and slow transition of the distribution of firms back to steady state require solving the transition dynamics over more than 1,200 periods, which is computationally expensive. We therefore economize by setting a grid of evenly spaced values for X_1 and ρ_X for each shock and choose those values that minimize the distance between our two data points and the model.

[26] In the United States since 2008, the debt-output ratio drops by nearly 10 percent and six years later is still 5 percent below its pre-recession level. See footnote 21 for the construction of aggregate debt.

[27] Entry in the data is measured as the number of firms reporting an age of zero divided by the total number of firms in the US Census Bureau's Longitudinal Business Database. Since the survey is annual, the measure excludes firms that enter and exit within a year.

Figure 8: Dynamics of some macroeconomic variables

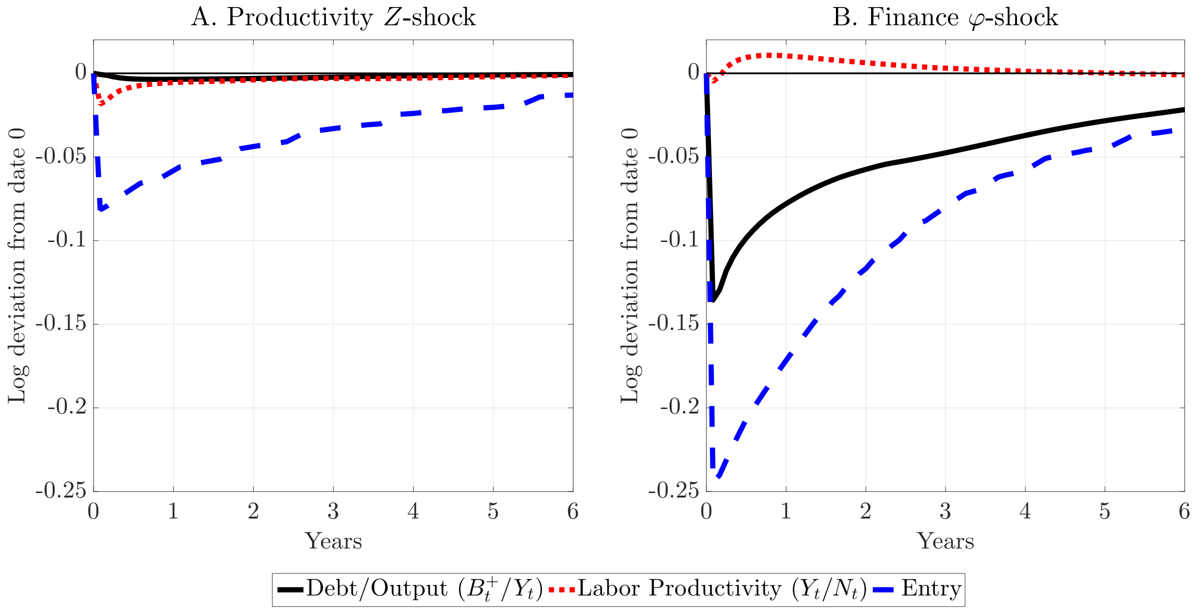

decline sharply, since a large fraction of them are constrained (recall Figure 6), leading to a decline in start-ups. Overall, we conclude that the differential responses of these three variables clearly identify a financial shock in the 2008 recession.

Figure 9 displays the dynamics of the labor market. In both experiments, the response is close to its empirical counterpart shown in Figure 1.[28] The financial shock induces larger and more persistent responses in vacancies, unemployment, and the job-finding rate. Under both scenarios, the decline in aggregate recruiting intensity is sizable, but its magnitude and persistence are, again, larger under the financial shock: Φ_t falls by 25 percent at impact (20 percent under the productivity shock) and five years later it is still 10 percent below its initial value (5 percent under the productivity shock). We conclude that, in the model, the financial shock—the more promising candidate to rationalize the Great Recession based on our discussion of Figure 8—can explain around half of the observed decline in aggregate match efficiency (recall the empirical path in Figure 1).

At first sight, it may be surprising that the response of aggregate recruiting intensity is not too dissimilar across the two shocks, although as we have shown, the entry rate of new firms—which accounts for a disproportionate share of job creation—differs remarkably under the two

[28]In the data, labor market variables move more slowly, but recall that we specified shocks that declined sharply on impact.

Figure 9: Dynamics of labor market variables

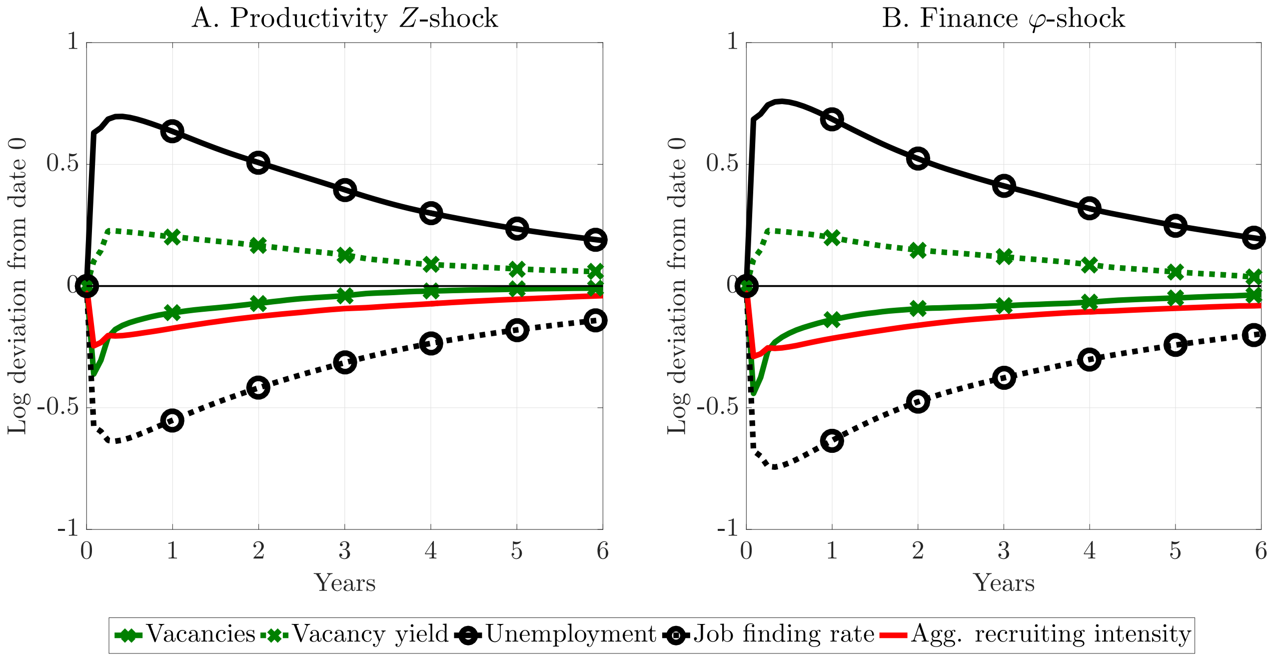

experiments. In what follows, we explain this apparent puzzle.

5.3 The Transmission Mechanism

To understand how macroeconomic shocks transmit to aggregate recruiting intensity, we return to our expression for Φ_t, using λ_t^h to denote the distribution of hiring firms:

$$\Phi_t = \left(\frac{V_t^*}{V_t}\right)^{\alpha} = \left[\int e_{it}\left(\frac{v_{it}}{V_t}\right) d\lambda_t^h\right]^{\alpha}. \tag{20}$$

Substituting the policy function for recruitment effort (12) into the above equation and taking log differences, we obtain:

$$\Delta \log \Phi_t = \underbrace{-\alpha\frac{\gamma_2}{\gamma_1+\gamma_2}\Delta \log q(\theta_t^*)}_{\text{Slackness effect}} + \underbrace{\alpha\Delta \log\left[\int g_{it}^{\frac{\gamma_2}{\gamma_1+\gamma_2}}\left(\frac{v_{it}}{V_t}\right) d\lambda_t^h\right]}_{\text{Composition effect}}. \tag{21}$$

We call the two elements of this equation the *slackness* and *composition* effect, respectively.

The Slackness Effect. The slackness effect is the change in aggregate recruiting intensity Φ_t accounted for by firms' changing effort in response to movements in labor market slackness

Figure 10: Decomposition of Aggregate Recruiting Intensity

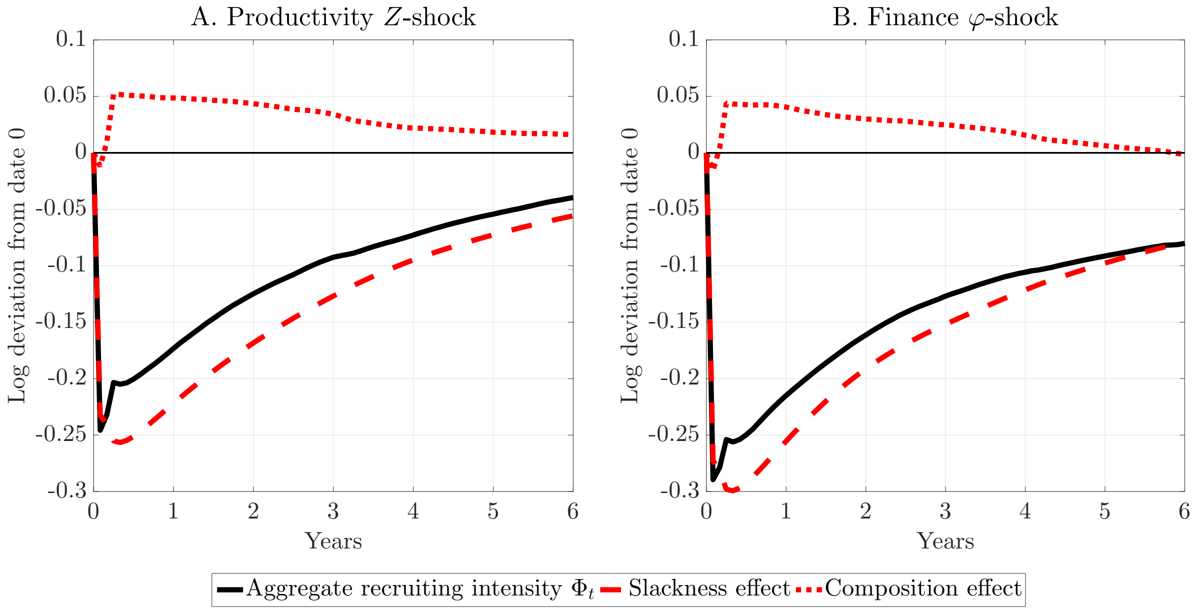

$q(\theta_t^*)$, holding constant growth rates g_{it}, vacancies v_{it}, and the distribution of hiring firms λ_t^h.

In a recession, equilibrium labor market slackness increases, as a spike in job separations increases the measure of unemployed workers, while the reduction in expected profitability reduces vacancy creation. This surge in slackness raises the probability $q(\theta_t^*)$ that any vacancy matches with a job seeker. Therefore, given the hiring technology $g_{it} = q(\theta_t^*)e_{it}v_{it}/n_{it}$, a hiring firm with a target growth rate g_{it} reoptimizes its combination of recruiting inputs e_{it} and v_{it} and decreases both: a slack labor market makes it easier for employers to hire, so employers spend less to attract workers. Since recruiting effort is more sensitive than vacancies to $q(\theta_t^*)$—recall the decision rules (12) and (13)—the slackness effect is always stronger on the effort margin and, in the aggregate, V_t^* declines more than V_t or, equivalently, Φ_t falls in recessions.

The Composition Effect. We define the composition effect residually, thereby including the impact on aggregate recruiting intensity of changes in the distribution of growth rates g_{it} and vacancy policies v_{it} among all hiring firms.

Figure 10 shows how these two components of aggregate recruiting intensity respond to the shocks. These figures reveal that the slackness effect (dashed line) is quantitatively the largest, accounting for almost all of the decline in aggregate recruiting intensity (solid line).

The large magnitude of the slackness effect was, perhaps, expected. Market tightness

Figure 11: Unpacking the composition effect

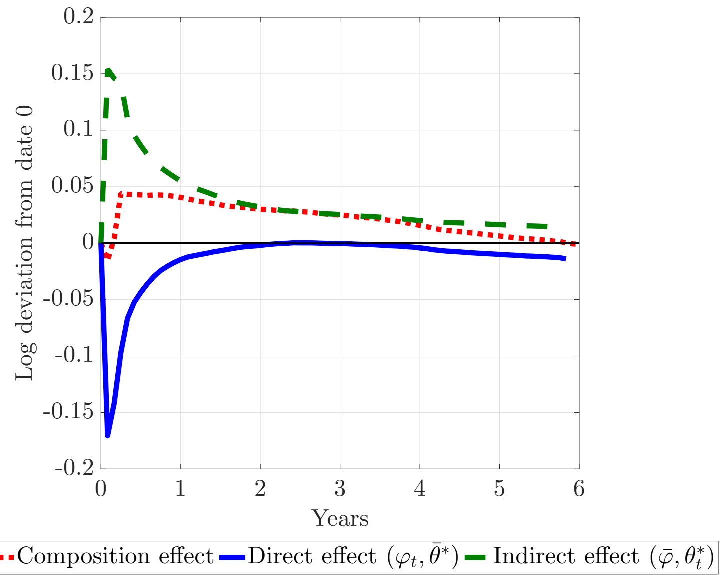

plunges and the elasticity of firm-level recruiting intensity with respect to q is high, nearly one.[29] What is more surprising is that the composition effect is so small and, in particular, after a drop at impact, it becomes positive, i.e., it induces a small countercyclical movement in Φ_t. We now explain this result.

5.3.1 Inspecting the Composition Effect

A useful approach is to split the composition effect into two further elements, which we plot in Figure 11.[30] The first is a *direct composition effect*: the response to the shock in a partial-equilibrium economy, keeping θ_t^* at its steady-state level, denoted $\bar{\theta}^*$. The second is the *indirect composition effect*: the response in an economy under the equilibrium path for θ_t^* induced by the shock while keeping φ_t at its steady state value $\bar{\varphi}$.

[29] We chose to express the slackness effect as a function of θ_t^* because this is a sufficient statistic for aggregate labor market conditions in the firm's hiring problem. One can also obtain an expression for the slackness effect that is a function of the more common measure of tightness θ. Substituting the relationship $q(\theta_t^*) = q(\theta_t)\Phi_t^{-\frac{1-\alpha}{\alpha}}$ in (21) and collecting the terms in Φ_t yields the alternative representation of the slackness effect $\frac{-\alpha[\gamma_2/(\gamma_1+\gamma_2)]}{1-(1-\alpha)[\gamma_2/(\gamma_1+\gamma_2)]}\Delta \log q(\theta_t)$. The denominator is less than one and captures a "multiplier": when Φ is low in the aggregate, firms exert less effort e. This alternative decomposition gives very similar results: if anything, the slackness effect is somewhat stronger.

[30] We illustrate this decomposition only for the tightening of the collateral constraint. Results for the productivity shock are almost identical.

Figure 12: Growth rate distribution and recruitment policies before and after the shock

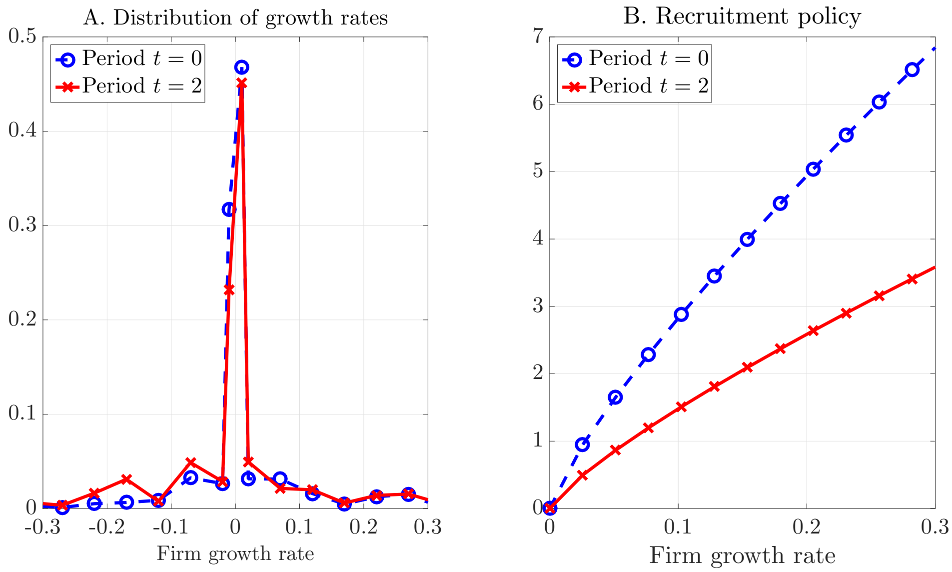

Notes: (i) In period 0 the economy is in steady state, in period 1 the path for the shock is realized, and period 2 then follows. We choose period 2 rather than period 1 since due to the timing of the model, workers that are fired in period 1 do not enter the labor market as unemployed workers until period 2. (ii) The growth rate distribution is computed over bins of width 0.02.

The direct effect reduces aggregate recruiting intensity on impact, since the drop in the collateral parameter lowers firm growth rates and reallocates hiring away from young, fast-growing firms that account for the bulk of recruiting intensity in the economy. Note that the direct component reverts rapidly towards zero. This is due to the fact that decline in φ_t induces positive selection among the hiring firms. The fraction of firms hiring drops from 55 percent in steady state to 22 percent following the shock. So these firms—both incumbents and entrants—have on average higher productivity and thus grow slightly more—a force that pushes aggregate recruiting intensity back up.

The indirect effect increases aggregate recruiting intensity on impact. As $q(\theta_t^*)$ rises, growing firms can meet job seekers more easily, intertemporally substituting their planned hiring. As they grow more quickly, they exert more recruiting effort, pushing up aggregate recruiting intensity. Selection of hiring firms on productivity tempers this effect as well: the increase in $q(\theta^*)$ *reduces* the average productivity of hiring firms, since some firms that did not hire in steady state do hire under higher $q(\theta^*)$, thereby dampening aggregate recruiting intensity.

Overall, the direct and indirect components show large movements, but these movements offset each other and the composition effect remains small throughout the transition.

Figure 12 provides another way to appreciate why the slackness effect is bound to dominate the composition effect. Panel A describes the behavior of the (unweighted) distribution of firm growth rates. Relative to steady state ($t = 0$), in the period following the shock ($t = 2$), firing firms contract faster and hiring firms expand slightly faster (thus, the dispersion of growth rates increases, as we discuss in some detail below).[31] Panel B shows how the slackness effect contributes to lower recruiting intensity at any given hiring rate (recall eq. 12). These two panels show that the choice of hiring firms to change their effort as market tightness varies over time dominates the compositional changes across growth rates in the pool of hiring firms.

The analysis in this section highlights the role of general equilibrium forces in the dynamics of aggregate recruiting effort of firms. A casual look at the microeconomic relationship between job-filling and hiring rates may induce one to conclude that economy-wide recruiting intensity declines after a negative macro shock because the shock curtails the speed at which hiring firms expand. Such a force is present and reflects the direct composition effect. But this logic ignores the adjustment of equilibrium market tightness that sets in motion the slackness effect and the indirect composition effect, the other essential—and quantitatively dominant—pieces of the transmission mechanism.

5.3.2 Relationship with Kaas and Kircher (2015)

In Kaas and Kircher's model of competitive search, aggregate recruiting intensity can be expressed as an average of meeting rates in each market, where each meeting rate is a concave function of market tightness. In terms of our notation, $\Phi_t^{KK} = \int q(\theta_{mt})(v_m/V)dm$, where m indexes markets. The authors find that, during productivity-driven recessions, the dispersion of tightness across markets increases, leading to a decline in Φ_t^{KK}. They ascribe the procyclicality of aggregate recruiting intensity chiefly to this mechanism.[32]

A version of this mechanism is present in our model as well. An increase in the standard deviation of growth rates will have a negative effect on Φ_t since the second term in (21) is concave in g_{it} (i.e., $\gamma_2/(\gamma_1 + \gamma_2) < 1$). Note that this source of fluctuations in Φ_t will enter exclusively into the composition effect.

[31] Figure B2 in Appendix B plots the employment-weighted kernel density function of the distribution of firm-level growth rates in the model. This reproduces well its data counterpart, Figure 5 in Davis, Faberman, and Haltiwanger (2012a).

[32] This mechanism is explained on pages 3053-3054 of their article.

Is this mechanism quantitatively important in our model? Our calibrated model is well suited to answer this question since (i) we match the empirical standard deviation of growth rates (recall Table 2), and (ii) the financial shock generates an empirically reasonable increase in dispersion: a 45 percent increase in the standard deviation of growth rates, compared to a 39 percent increase in the data (Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry, 2012).[33] To gauge the importance of this mechanism, we compare our measure to one in which this curvature effect is absent, computing

$$\alpha \Delta \log \left[\int g_{it}^{\frac{\gamma_2}{\gamma_1+\gamma_2}} \left(\frac{v_{it}}{V_t} \right) d\lambda_t^h \right] - \alpha \Delta \log \left[\int g_{it} \left(\frac{v_{it}}{V_t} \right) d\lambda_t^h \right]^{\frac{\gamma_2}{\gamma_1+\gamma_2}} .$$

The first term is the composition effect, and the second is its counterpart where we raise the integral (not the integrand) to the exponent.[34] Following a financial shock, we find that, between $t = 0$ and $t = 2$, this measure equals -4 percent. Therefore, this mechanism explains only around 15 percent of the decline in aggregate recruiting intensity generated by our model. Its contribution is limited by the fact that empirically $\gamma_2/(\gamma_1 + \gamma_2)$ is close to 1, so the degree of concavity of the integrand in the composition effect is small. We conclude that the key transmission mechanism of our model, the slackness effect, is different from that emphasized by Kaas and Kircher (2015).

5.4 Cross-Sectional Dynamics

In addition to explaining the aggregate dynamics of the vacancy yield through recruiting intensity, our model also accounts for the cross-sectional dynamics of vacancy yields by size, as documented by Moscarini and Postel-Vinay (2016). We argue that the heterogeneity in the extent to which the financial friction binds across firms is key to understanding the latter.

We begin by splitting firms in the model into financially constrained firms and uncon-

[33] This is a notable feature of our model in response to a financial shock which provides another over-identifying test. Even without a shock to the dispersion of firm-level productivity growth we attain a significant increase in the dispersion of employment growth: as explained above, the shock adversely affects some firms, reducing their growth rates, whereas some other firms respond to the surge in labor market slackness by growing faster.

[34] We are in effect computing $\mathbb{E}\left[f(X)\right] - f\left(\mathbb{E}\left[X\right]\right)$, where $f(X) = X^{\frac{\gamma_2}{\gamma_1+\gamma_2}}$ and the random variable X is the hiring firm growth rate, which is distributed with a density $h(g_i) = \frac{v_{it}}{V_t} \lambda_{it}^h$.

Figure 13: Vacancy Yield by Size in the Model and the Data (JOLTS)

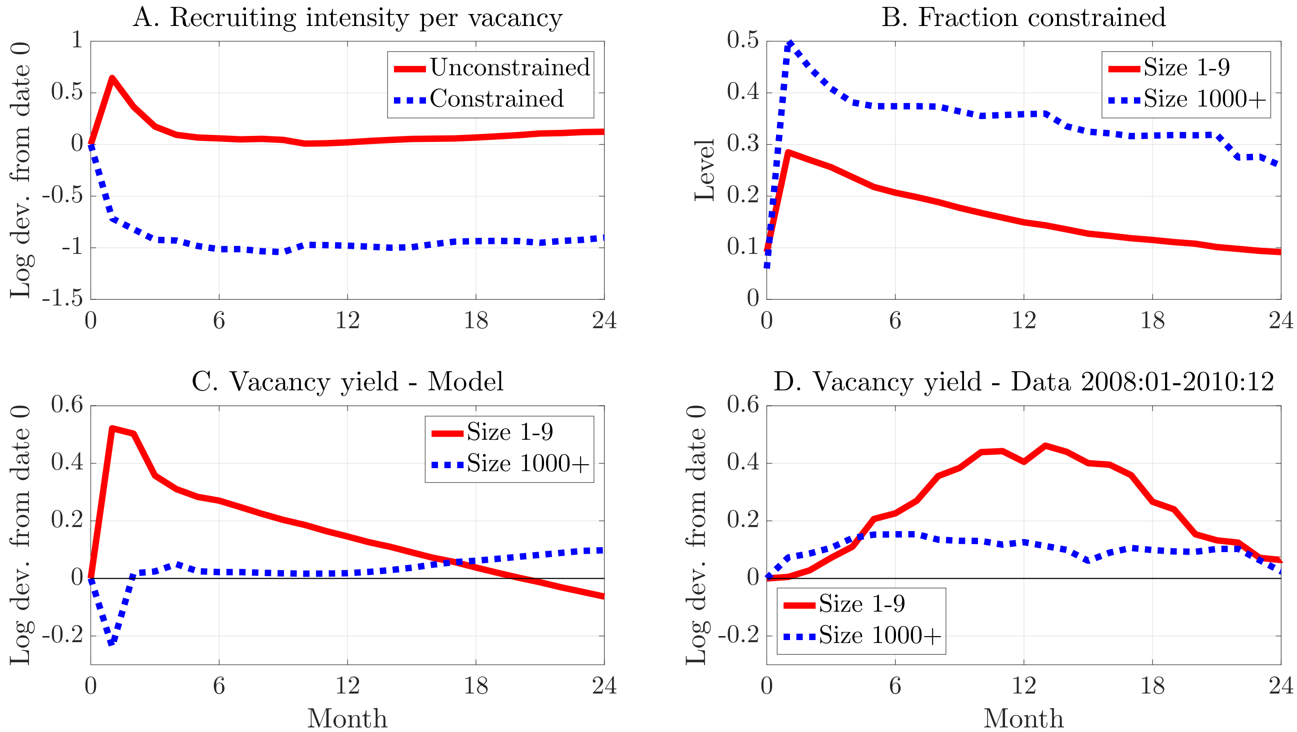

strained firms.[35] Panel A of Figure 13 shows that recruiting intensity dynamics differ markedly between the two types of firms. Among constrained firms, the financial shock causes a sharp drop in the growth rate and, therefore, in the recruiting intensity of those hiring. Unconstrained hiring firms, instead, increase their hiring in response to the surge in labor market slackness, choosing higher recruiting intensity. The constrained firms drive the direct component of the composition effect, whereas unconstrained firms drive the indirect component.

Turning to size, which is observable in JOLTS, panel B shows that, following the macro shock, the fraction of constrained firms rises significantly across all sizes but does so in particular among large firms. In the model, these are the young, fast-growing firms with high span of control parameter (σ_H). Panel C illustrates that the vacancy yield of these large firms is flat: reduced recruitment effort due to the financial shock offsets the effect of a slacker labor market, which would usually lead to higher vacancy yields in a recession. Meanwhile, the vacancy yield of small firms increases, as they receive the full effect of a slacker labor market.

Panel D shows that this narrative implied by our model is borne out in the data and pro-

[35] Financially constrained firms in the model are firms for which both the collateral and dividend constraints bind.

Figure 14: Size of the composition effect under different values of α

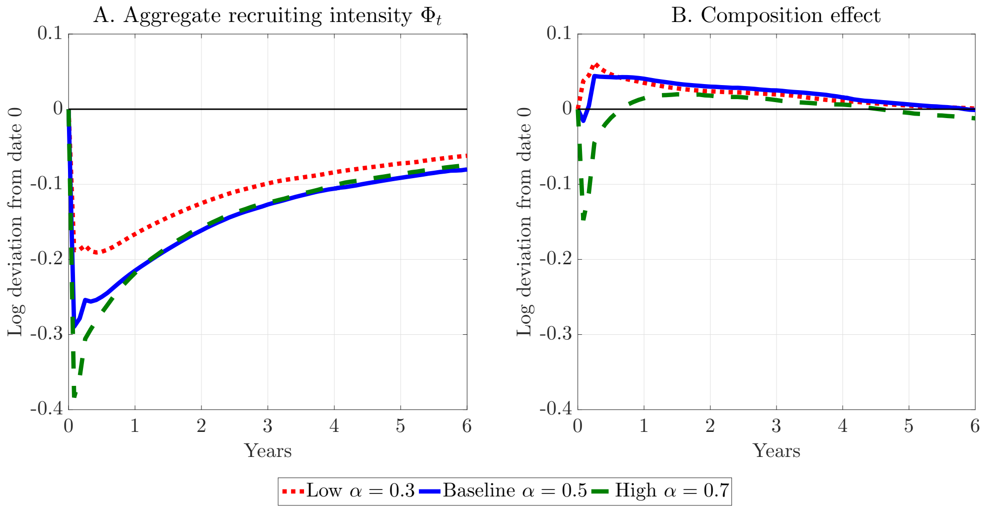

vides an overidentifying test of the model. During the Great Recession, the vacancy yields of small firms increased substantially, whereas while those of large firms remained flat. It also confirms that our characterization of the labor market—comprising fast-growing high-scale firms responding directly to a macroeconomic shock, and small low-scale firms responding indirectly to market tightness—is a useful lens for thinking about hiring dynamics in the US labor market.

6 Robustness

In this section, we analyze the robustness of our main finding regarding how shocks are transmitted to aggregate recruiting intensity: a large slackness effect and a composition effect strongly tempered by its indirect component.

We start by examining different model calibrations that could yield a larger composition effect. Next, we study how permanent heterogeneity in vacancy filling rates across industries—for example due to different recruiting methods—affects the size of the composition effect. Then we discuss how the inclusion of quits, replacement hires, and on-the-job search could affect our conclusions. Finally, we reflect on whether the strong offsetting force that counteracts the composition effect is only germane to financial frictions or would, possibly, survive under other mechanisms that may underlie the observed rich firm dynamics by age and size.

6.1 Alternative Calibration

As is clear from equation (21), the magnitude of the composition effect is especially sensitive to the value of α, the elasticity of hires with respect to vacancies. Figure 14 plots the response of aggregate recruiting intensity (panel A) and the composition effect (panel B) for three values of α in the neighborhood of existing estimates. In the range below 0.5, our baseline value, the total composition effect is small at impact and turns positive quickly as its indirect component takes over. However, for $\alpha = 0.7$, the composition effect becomes sizable at impact and remains negative for almost a year after the shock.

To understand this result, note that the strength of the indirect component of the composition effect (which facilitates firm growth during a recession as labor market tightness falls) is determined by how much the meeting rate $\log(q_t) = -(1-\alpha)\log\theta_t^*$ rises in a downturn. Hence, a large value of α mutes the countercyclical indirect component, inducing larger procyclical movements in the composition effect. A stronger composition effect also explains the deeper drop in Φ_t, as shown in panel A of Figure 14.[36]

6.2 Sectoral Heterogeneity in Recruiting Technology

Ours is a one-sector model of the aggregate economy in which all firms face the same recruiting technology. DFH document that different sectors of the economy display consistently different vacancy yields. To the extent that such discrepancies in vacancy yields stem from systematic differentials in growth rates across sectors, then our model will capture these since we generate a realistic distribution of firm growth rates (Table 3).[37] If, however, they are due to permanent characteristics of the recruiting technology across sectors, then a macro shock that changes the sectoral shares of hiring firms will affect aggregate match efficiency and should appear in the composition effect.[38]

In the context of the Great Recession, this point is especially relevant. The Construction sector is an outlier in terms of its frictional characteristics (its vacancy yield is about 2.5 times as

[36]Note that the slackness effect is not as sensitive to α because, as seen in equation (21), $\log q$, which contains the term $1-\alpha$ is also multiplied by α.

[37]Indeed, DFH Figure B.5 shows that the cross-sector variation in average hiring rates is strongly correlated with the cross-sector variation in vacancy yields.

[38]We thank Steve Davis for suggestions that led to the inclusion of this section.

large as in the economy as a whole), and it was hit particularly hard in the recession. One would therefore expect Construction to play a significant role in the national movement of aggregate recruiting intensity, despite its small share of employment (Davis, Faberman, and Haltiwanger, 2012b).

A fully specified multisector model is beyond the scope of this paper, but we can nevertheless estimate the size of this *sectoral composition effect* using the structure of our model and industry-level data on vacancy yields and vacancy shares from JOLTS.[39] Suppose that the firm-level hiring technology in each sector $s = 1, \ldots, S$ is subject to a sector-specific recruitment efficiency shifter ϕ_s,

$$h_{ist} = \phi_s q\left(\theta_t^*\right) e_{ist} v_{ist}, \tag{22}$$

leading to a modified expression for aggregate recruiting intensity,

$$\Phi_t = \left[\int_i \phi_s e_{ist} \frac{v_{ist}}{V_t} di\right]^\alpha, \tag{23}$$

and the optimal choice of firm-level recruiting intensity,

$$e_{ist} = Constant \times \phi_s^{-\frac{\gamma_2}{\gamma_1+\gamma_2}} q\left(\theta_t^*\right)^{-\frac{\gamma_2}{\gamma_1+\gamma_2}} g_{ist}^{\frac{\gamma_2}{\gamma_1+\gamma_2}}. \tag{24}$$

Firm-level recruiting intensity depends negatively on sector-specific efficiency since firms belonging to sectors with a high recruiting efficiency can use less effort to realize any desired growth rate.

To decompose aggregate recruiting intensity, we can again substitute the optimal policy (24) into (23) to arrive at

$$\Phi_t = Constant \times q\left(\theta_t^*\right)^{-\frac{\alpha\gamma_2}{\gamma_1+\gamma_2}} \times \left[\sum_{s=1}^{S} \phi_s^{\frac{\gamma_1}{\gamma_1+\gamma_2}} \left(\frac{v_{st}}{V_t}\right) \int_{\{i\in s\}} \left[g_{ist}^{\frac{\gamma_2}{\gamma_1+\gamma_2}} \frac{v_{ist}}{v_{st}}\right] di\right]^\alpha.$$

The effect we are trying to determine comes from the interaction of permanent differences in

[39] In what follows, we maintain the assumption that all firms hire in the same labor market. Accordingly, one could read our exercise as the counterpart to one conducted on the worker side, in which different groups of job seekers enter the same labor market but are weighted by some fixed level of search efficiency. For example, see Hall and Schulhofer-Wohl (2015) and Hornstein and Kudlyak (2016).

Figure 15: Sectoral composition effect

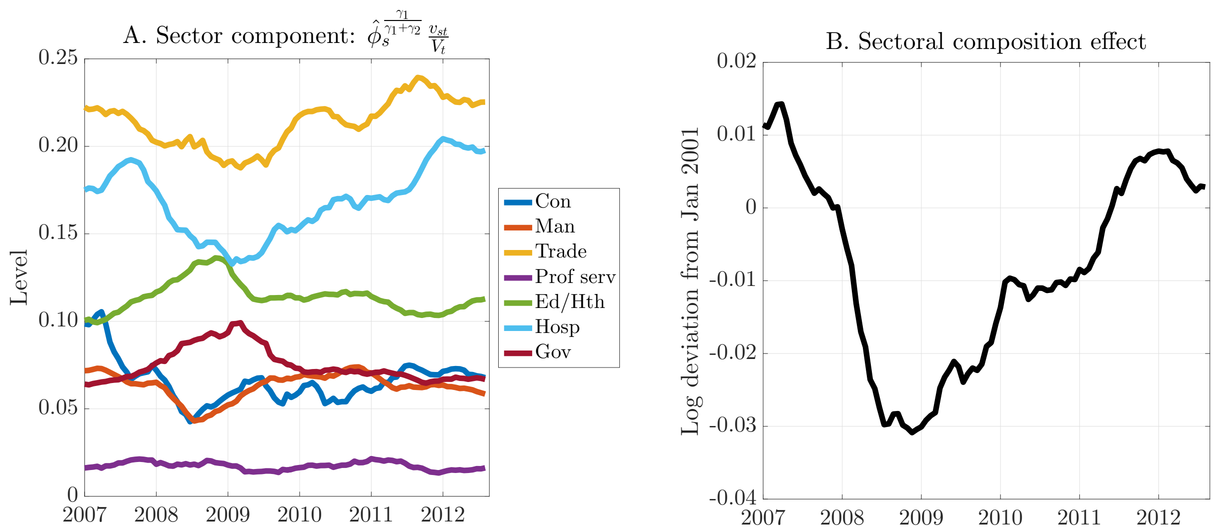

Notes: (i) Panel A plots the vacancy yield in the seven largest industries. In Panel B, we have normalized the log of the sectoral composition effect to zero in January 2008.

match efficiencies across sectors ϕ_s and the sectoral composition of hiring firms given by the vacancy share v_{st}/V_t. Therefore, we assume that the distribution of growth rates and vacancies is identical within each sector and, thus, the integral term is constant across sectors. Under this assumption, we obtain a counterpart to our previous decomposition of aggregate recruiting intensity, with an additional term characterizing the sectoral composition effect:

$$\Delta \log \Phi_t = -\alpha \frac{\gamma_2}{\gamma_1 + \gamma_2} \Delta \log q\left(\theta_t^*\right) + \alpha \Delta \log \left[\int_i g_{it}^{\frac{\gamma_2}{\gamma_1+\gamma_2}} \frac{v_{it}}{V_t} di\right] + \underbrace{\alpha \Delta \log \left[\sum_{s=1}^{S} \phi_s^{\frac{\gamma_1}{\gamma_1+\gamma_2}} \frac{v_{st}}{V_t}\right]}_{\text{Sectoral composition effect}}. \quad (25)$$

Note that the exponent on ϕ_s is less than one. A sector with a higher ϕ_s will be exogenously more productive in creating matches, increasing recruiting intensity with an elasticity of one with respect to its vacancy share; however, the firms in that sector will endogenously decrease effort with an elasticity of $\frac{\gamma_2}{\gamma_1+\gamma_2}$, leaving the net elasticity of $\frac{\gamma_1}{\gamma_1+\gamma_2}$.

Computing the last term in (25) requires data on vacancy shares by sector, readily available from JOLTS and data on sectoral match efficiency. Under our assumptions, (22) and (23) imply

$$\phi_s^{\frac{\gamma_1}{\gamma_1+\gamma_2}} = \frac{H_{st}/V_{st}}{H_{kt}/V_{kt}},$$

where match efficiency of the sector k is normalized to one without loss of generality.[40]

Using data on all eleven two-digit industries from JOLTS, we plot the sectoral component $\phi_s^{\frac{\gamma_1}{\gamma_1+\gamma_2}} \frac{v_{st}}{V_t}$ for the largest seven sectors in panel A of Figure 15, and the total sectoral composition effect in panel B. We find that this component generates an additional 4 percent drop in aggregate recruiting intensity during the Great Recession—mostly due to the decline in the vacancy shares of Construction, Manufacturing, and Hospitality and Leisure. Even though adding this mechanism shifts the decomposition more toward the composition effect, we tentatively conclude that it does not modify our conclusion that the slackness channel is dominant. Obviously, the economy has a lot more structural heterogeneity than that implied by our coarse partition into seven industries. Incorporating additional relevant sources of heterogeneity remains an open area for future research.

6.3 Replacement Hiring and On-the-Job Search

In our baseline model, we have abstracted from replacement hiring associated with quits and search on the job, two related and prominent features of labor markets. We now assess to what extent these omissions could affect our conclusions.

6.3.1 A Larger Composition Effect with Quits and Replacement Hiring?

We have solved the model under a range of values for an exogenous quit rate between one and three percent per month, and found our results to be quantitatively very robust. The reason is that in our model, as in the data, the bulk of hires are made by firms with positive net hiring rates. However, this model is stylized in that it assumes a quit rate that is invariant in the cross section and over time. These assumptions might lead us to understate the composition effect.

[40] To estimate ϕ_s, we use ratios of average sectoral vacancy yields from 2005 to 2006. We take Professional Business Services as the normalizing sector, since its average vacancy yield of 1.30 is the sectoral median. We use data for all nine sectors available in JOLTS.

In the cross section, the data show that among shrinking firms quits are especially high, but at the same time some of these firms display positive gross hiring rates.[41] An adverse aggregate shock pushes many firms into this position of negative net growth and replacement hire. In our model, these firms contribute to the determination of aggregate recruiting intensity before the shock (because they hire) but not after the shock (because downsizing firms do not hire). If the model allowed for replacement hires, some of these firms would instead contribute to aggregate recruiting intensity even after the shock, and do so with a lower gross hiring rate and, therefore, a lower recruiting intensity.[42] In our analysis, this shift would be captured by a more negative composition effect.

In the aggregate, the data show that quits are strongly procyclical, falling sharply during a recession. The average gross hiring rate therefore also declines, because of a reduction in replacement hires, and as a consequence, recruiting intensity decreases across all hiring firms. In our analysis, this result would also be captured by a more negative composition effect.

6.3.2 A Smaller Slackness Effect with On-the-Job Search?

When a large portion of job seekers are employed, the response of market tightness to spikes in layoffs to unemployment—such as those following financial and productivity shocks—would be smaller. This mechanism has the potential to weaken the slackness effect.

In an economy where firms engage in take-it-or-leave-it offers to risk-neutral workers, modeling search on the job is relatively simple once it is assumed that (i) firms commit to not responding to the poaching competitor when an employed worker receives an outside offer and (ii) the worker—who is indifferent between staying and going—quits. In addition, we make the following minimal amendments to the model: (i) all employed workers search with a relative search intensity of s determining the effective units of search of an employed worker relative to an unemployed worker (whose intensity is normalized to 1); and (ii) the matching function is modified to take the total measure of effective search units $S_t = U_t + sN_t$ as an input, where $N_t = \bar{L} - U_t$ is the measure of employed workers. The firm-level hiring technology remains

[41] For example, a negative productivity shock leads a firm to fire some of its worst workers. Meanwhile, some of its best workers quit to find a more productive match, leaving the firm to replace some of these quits with new hires.

[42] More specifically, they would be part of the integral in the second term of (21).

$h_{it} = q_t e_{it} v_{it}$, but the law of motion for firm employment is now

$$n_{it+1} = n_{it} + h_{it} - f_{it} - sp\left(\theta_t^*\right) n_{it} \, ,$$

where $p\left(\theta_t^*\right)$ is the job-finding rate of the unemployed. By constant returns to scale in the matching functions, $sp\left(\theta_t^*\right)$ is the job-finding rate of employed workers. As a result, the law of motion for unemployment becomes

$$U_{t+1} = U_t + F_t - \left[\frac{U_t}{U_t + sN_t}\right] H_t,$$

where $U_t / \left(U_t + sN_t\right)$ is the fraction of total hires that come from unemployment.

In choosing a value for on the job search intensity, note that s is equal to the ratio of employment-employment (EE) to unemployment-employment (UE) transition rates. Following Fujita and Moscarini (Forthcoming) and, thus, excluding recalls and workers on temporary layoffs from UE, we obtain $s = 0.09$ for the prerecession period.

What is the impact of on-the-job search on the slackness effect? Consider an increase in the firing rate due to a negative macro shock. In the baseline model, the monthly firing rate is $F_t / N_t = 0.03$. Suppose that this ratio were to spike in a recession, doubling. Without on-the-job search, the mass of effective search units increases nearly one for one, by 0.03. In the model with on-the-job search, $S_t = (1-s)U_t + s\bar{L}$, so although the number of unemployed workers rises by 0.03, the measure of total job seekers increases by $(1-s) \times 0.03 = 0.027$. Therefore, labor market tightness falls by less and the slackness effect is somewhat weakened, as expected, but this correction is quantitatively small. The reason is that, although the *stock* of employed workers is large, their average search intensity is low relative to that of the unemployed. Moreover, if one were to also allow s to vary over the cycle and match the data, then the relative intensity of the employed would be countercyclical.[43] This force would partially counteract the initial correction, thus making the total effect of on-the-job search on the dynamics of market tightness even smaller.

We acknowledge that ours is only a back-of-the-envelope calculation and that a thorough analysis would require a more satisfactory representation of on-the-job search behavior (i.e.,

[43] Figure B3 in Appendix B documents the cyclicality of the relative search intensity of the employed.

one where workers are not indifferent between staying and moving). Frictional models of the labor market with both a realistic firm size distribution induced by DRS in production and a rich job ladder whereby high-productivity, high-wage firms can poach workers more easily from other firms—and thus the vacancy filling rate is increasing in the firm type because it is further up the ladder—have not yet been fully developed.[44] Whether such class of models has novel forces at work relative to those emphasized here remains to be established.

6.4 Alternative Frictions

One of the main insights of our analysis is that financially unconstrained firms unravel the response of constrained firms and mitigate the composition effect of a macro shock on aggregate recruiting intensity. A natural questions is whether this result specific to models where the key source of firm dynamics over the life cycle is a financial friction? Frictions of a different nature may underlie the observed rich up-and-out dynamics of young firms: how robust is our insight to these generalizations?

Without fully solving alternative models, offering a precise answer to this question is challenging. However, we conjecture that our result is more general than it may appear at first sight. Any successful model of firm dynamics combined with labor market frictions would feature the following minimal set of ingredients: (i) heterogeneity across firms induced by idiosyncratic shocks, and (ii) a friction —over and above hiring costs tied to search/matching— that slows down growth for young firms over the life cycle and that, interacted with shocks, generates up-and-out dynamics. In our setup: (i) is generated by productivity shocks and (ii) by financial market imperfections.

A common property of any such model is that firms escape the friction conditional on surviving long enough. Consider, for example, two popular alternative sources of life-cycle dynamics: learning and customer capital. If the friction is tied to learning about one's own productivity, after enough time in the market, much of the fixed individual productivity effect will be revealed. If the friction is tied to the necessity of attracting a customer base, after enough time in the market, the firm would have built a demand for its product. However, in both cases,

[44] Promising environments are those developed by Lentz and Mortensen (2008) with random search and Schaal (2017) with directed search. Even though they do not study the determinants of and the transmission of shocks to recruiting intensity, as we do, their frameworks lend naturally to addressing such questions as well.

even these older unconstrained firms are still subject to shocks that change their optimal size. Therefore, there will always be a share of firms in the economy that experience positive shocks and are gradually (because of the hiring costs) reaching their new higher employment target. When a negative aggregate shock hits the economy, many of these firms would still want to hire. As in our model, they would then take advantage of the slack labor market by growing even faster and exerting even more recruiting effort.

While this logic suggests that the composition effect would remain relatively small across a wider range of firm dynamics models, ours remains a conjecture that can be verified only by explicitly solving and plausibly parameterizing these models.

7 Conclusions

The existing literature on the cyclical fluctuations of aggregate match efficiency has focused almost exclusively on explanations involving the worker side of the labor market, such as occupational mismatch, shifts in job search intensity of the unemployed over the cycle, and compositional changes among the pool of job seekers. In this paper, we have shifted the focus to the firm side and, building on the microeconomic evidence in Davis, Faberman, and Haltiwanger (2013), developed a macroeconomic model of aggregate recruiting intensity.

The model, parameterized to replicate a range of cross-sectional facts about firm dynamics and hiring behavior, is able to explain about half of the collapse in aggregate match efficiency during the Great Recession through a sharp decline in firms' recruiting intensity. Our analysis of the transmission mechanism points toward the importance of general equilibrium forces: aggregate recruiting intensity declined mainly because the number of available job seekers per vacancy increased (i.e., labor market tightness declined), making it easier for firms to achieve their recruitment targets without having to spend as much on recruitment costs. Changes in the within-sector composition of the pool of hiring firms, for example, due to the fall in new firm entry that is well matched by the model, did not play a large role. The shift in sectoral composition—in particular, the bust in Construction and other sectors with structurally high job-filling rates, instead contributed to the measured deterioration in aggregate recruiting effort.

Besides its contribution to understanding the determinants of movements in match efficiency, and thus the job-finding rate—a key object for labor market analysis—our theory has broader implications for macroeconomics. First, as, for example, Faberman (2014) discusses, making progress in understanding how firms' hiring decisions respond to macroeconomic conditions is important since job creation policies that fail to recognize the determinants of employers' recruitment effort may fall short in achieving their goal. Our model predicts that subsidizing firm hiring (abstracting from the offsetting effects of higher tax rates) will increase the average firm growth rate and induce a rise in recruiting intensity, whereas a subsidy to workers' job search that decreases market tightness will induce a decline in recruiting intensity, through the slackness effect discussed in the paper. Second, a richer model of employer recruiting behavior can lead to better estimates of the marginal cost of labor and, therefore, result in improved measures of the labor wedge and of the relative importance of labor and product market wedges (Bils, Klenow, and Malin, 2014). In this respect, our model suggests that the price of labor faced by firms may be more procyclical than what would appear from naively using wages as a proxy.

References

ACEMOGLU, D., AND W. B. HAWKINS (2014): "Search with Multi-Worker Firms," *Theoretical Economics*, 9(3), 583–628.

BILS, M., P. J. KLENOW, AND B. A. MALIN (2014): "Resurrecting the Role of the Product Market Wedge in Recessions," NBER Working Paper 20555, National Bureau of Economic Research.

BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2012): "Really Uncertain Business Cycles," NBER Working Paper 18245, National Bureau of Economic Research.

BUERA, F. J., R. N. F. JAEF, AND Y. SHIN (2015): "Anatomy of a Credit Crunch: From Capital to Labor Markets," *Review of Economic Dynamics*, 18(1), 101–117.

CHODOROW-REICH, G. (2014): "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis," *Quarterly Journal of Economics*, 129(1), 1–59.

COOPER, R., J. C. HALTIWANGER, AND J. L. WILLIS (2007): "Search Frictions: Matching Aggregate and Establishment Observations," *Journal of Monetary Economics*, 54, 56–78.

DAVIS, S. J., R. J. FABERMAN, AND J. C. HALTIWANGER (2012a): "Labor Market Flows in the Cross Section and Over Time," *Journal of Monetary Economics*, 59(1), 1–18.

——— (2012b): "Recruiting Intensity during and after the Great Recession: National and Industry Evidence," *The American Economic Review, Papers and Proceedings*, 102(3), 584–88.

——— (2013): "The Establishment-Level Behavior of Vacancies and Hiring," *Quarterly Journal of Economics*, 128(2), 581–622.

DAVIS, S. J., R. J. FABERMAN, J. C. HALTIWANGER, AND I. RUCKER (2010): "Adjusted Estimates of Worker Flows and Job Openings in JOLTS," in *Labor in the New Economy*, ed. by K. G. Abraham, J. R. Spletzer, and M. Harper, pp. 187–216. Chicago, IL: University of Chicago Press.

DAVIS, S. J., AND J. C. HALTIWANGER (2015): "Dynamism Diminished: The Role of Credit Conditions and Housing Markets," Mimeo, University of Maryland.

ECKSTEIN, Z., O. SETTY, AND D. WEISS (2014): "Financial Risk and Unemployment," CEPR Discussion Paper 10596, C.E.P.R. Discussion Papers.

ELSBY, M. W. L., AND R. MICHAELS (2013): "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows," *American Economic Journal: Macroeconomics*, 5(1), 1–48.

EVANS, D. S., AND B. JOVANOVIC (1989): "An Estimated Model of Entrepreneurial Choice under Liquidity Constraints," *Journal of Political Economy*, 97(4), 808–27.

FABERMAN, R. J. (2014): "Recruiting Intensity," *IZA World of Labor*, pp. 1–21.

FERNALD, J. G. (2014): "Productivity and Potential Output before, during, and after the Great Recession," in *NBER Macroeconomics Annual*, ed. by J. A. Parker, and M. Woodford, vol. 29, pp. 1–51. Chicago, IL: University of Chicago Press.

FOSTER, L., J. C. HALTIWANGER, AND C. SYVERSON (2016): "The Slow Growth of New Plants: Learning about Demand?," *Economica*, 83(329), 91–129.

FUJITA, S., AND G. MOSCARINI (Forthcoming): "Recall and Unemployment," *American Economic Review*.

HALL, R. E., AND S. SCHULHOFER-WOHL (2015): "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Jobseekers," NBER Working Paper 20939, National Bureau of Economic Research.

HALTIWANGER, J. C. (2011): "Firm Dynamics and Productivity Growth," EIB Paper 5/2011, European Investment Bank, Economics Department.

——— (2012): "Job Creation and Firm Dynamics in the United States," *Innovation Policy and the Economy*, 12(1), 17–38.

HALTIWANGER, J. C., R. S. JARMIN, AND J. MIRANDA (2010): "Who Creates Jobs? Small vs. Large vs. Young," NBER Working Paper 16300, National Bureau of Economic Research.

HOPENHAYN, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–50.

HORNSTEIN, A., AND M. KUDLYAK (2016): "Estimating Matching Efficiency with Variable Search Effort," Working Paper 16-13R, Federal Reserve Bank of Richmond.

JOVANOVIC, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50(3), 649–70.

KAAS, L., AND P. KIRCHER (2015): "Efficient Firm Dynamics in a Frictional Labor Market," *American Economic Review*, 105(10), 3030–60.

KETTEMANN, A., A. I. MUELLER, AND J. ZWEIMULLER (2016): "Wages, Workers and Vacancy Durations: Evidence from Linked Data," Mimeo, Columbia GSB.

LEDUC, S., AND Z. LIU (2017): "The Slow Job Recovery in a Macro Model of Search and Recruiting Intensity," Working Paper Series 2016-9, Federal Reserve Bank of San Francisco.

LENTZ, R., AND D. T. MORTENSEN (2008): "An Empirical Model of Growth Through Product Innovation," *Econometrica*, 76(6), 1317–73.

MCGRATTAN, E. R., AND E. C. PRESCOTT (2012): "The Great Recession and Delayed Economic Recovery: A Labor Productivity Puzzle?," in *Government Policies and the Delayed Economic Recovery*, ed. by L. E. Ohanian, J. B. Taylor, and I. J. Wright, pp. 115–54. Stanford, CA: Hoover Institution.

MEHROTRA, N. R., AND D. SERGEYEV (2013): "Sectoral Shocks, the Beveridge Curve and Monetary Policy," Mimeo, Brown University.

——— (2015): "Financial Shocks and Job Flows," Mimeo, Brown University.

MIRANDA, M. J., AND P. L. FACKLER (2002): *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press.

Moscarini, G., and F. Postel-Vinay (2016): "Did the Job Ladder Fail after the Great Recession?," *Journal of Labor Economics*, 34(S1), S55–S93.

Mukoyama, T., C. Patterson, and A. Şahin (2014): "Job Search Behavior over the Business Cycle," Staff Report 689, Federal Reserve Bank of New York.

O'Leonard, K. (2011): "Talent Acquisition Factbook 2011: Benchmarks and Trends of Spending, Staffing and Key Talent Metrics," Discussion paper, Bersin and Associates.

Petrosky-Nadeau, N., and E. Wasmer (2013): "The Cyclical Volatility of Labor Markets under Frictional Financial Markets," *American Economic Journal: Macroeconomics*, 5(1), 193–221.

Pissarides, C. A. (2000): *Equilibrium Unemployment Theory*. Cambridge, MA: MIT Press.

Robb, A. M., and D. T. Robinson (2014): "The Capital Structure Decisions of New Firms," *Review of Financial Studies*, 27(1), 153–79.

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014): "Mismatch Unemployment," *American Economic Review*, 104(11), 3529–64.

Schaal, E. (2017): "Uncertainty and Unemployment," *Econometrica*, Forthcoming.

Sedláček, P. (2014): "Match Efficiency and Firms' Hiring Standards," *Journal of Monetary Economics*, 62, 123–33.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25–49.

——— (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15(2), 127–48.

Siemer, M. (2014): "Firm Entry and Employment Dynamics in the Great Recession," Finance and Economics Discussion Series 2014-56, Board of Governors of the Federal Reserve System (U.S.).

Wasmer, E., and P. Weil (2004): "The Macroeconomics of Labor and Credit Market Imperfections," *American Economic Review*, 94(4), 944–63.

Young, E. R. (2010): "Solving the Incomplete Markets Model with Aggregate Uncertainty Using the Krusell-Smith Algorithm and Non-stochastic Simulations," *Journal of Economic Dynamics and Control*, 34(1), 36–41.

Appendix For Online Publication

This Appendix is organized as follows. Section A contains the derivations of the hiring cost function that we introduced in Section 4. Section B provides additional figures referenced in the main text. Section C details the algorithms for the computation of the stationary equilibrium, transitional dynamics, and estimation of the model's parameters.

A The hiring cost function

In this section we show that, once we postulate the hiring cost function

$$\mathcal{C}(n,e,v) = \left[\frac{\kappa_1}{\gamma_1}e^{\gamma_1} + \frac{\kappa_2}{\gamma_2+1}\left(\frac{v}{n}\right)^{\gamma_2}\right]v, \tag{A1}$$

then, through the firm's optimization, we obtain a log-linear cross-sectional relationship between the job-filling rate and the employment growth rate that is consistent with the empirical findings in DFH. Next, by substituting the firm's first-order conditions into $(A1)$, we derive a formulation of the cost only in terms of (n,n') that we use in the intertemporal problem (10) in the main text.

As we explained in Section 3.1, the firm solves a static cost minimization problem: given a choice of n', it determines the lowest cost combination of (e,v) that can deliver n'. The hiring firm's cost minimization problem is

$$\begin{aligned} \mathcal{C}(n,n') \quad &= \quad \min_{e,v} \quad \left[\frac{\kappa_1}{\gamma_1}e^{\gamma_1} + \frac{\kappa_2}{\gamma_2+1}\left(\frac{v}{n}\right)^{\gamma_2}\right]v \\ s.t. \quad &: \qquad n'-n \leq q(\theta^*)\,ev \\ & \qquad\qquad v \geq 0 \end{aligned} \tag{A2}$$

Convexity of the cost function $(A1)$ in (e,v) requires $\gamma_1 \geq 1$ and $\gamma_2 \geq 0$. After setting up the Lagrangian, one can easily derive the two first-order conditions with respect to e and v that, combined, yield a relationship between the optimal choice of e and the optimal choice of the vacancy rate v/n:

$$e = \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1-1}\right)\right]^{\frac{1}{\gamma_1}}\left(\frac{v}{n}\right)^{\frac{\gamma_2}{\gamma_1}}. \tag{A3}$$

Note that if $\gamma_2 = 0$, as in Pissarides (2000), recruiting intensity is equal to a constant for all firms and it is independent of aggregate labor market conditions —both counterfactual implications. The following changes in parameters (ceteris paribus) result in a substitution away from vacancies and towards effort: $\uparrow \kappa_2, \downarrow \kappa_1, \uparrow \gamma_2$, and $\downarrow \gamma_1$. The effect of the cost shifter is obvious. A higher curvature on the vacancy rate in the cost function ($\uparrow \gamma_2$) makes the marginal cost of creating vacancies rise faster than the marginal cost of recruiting effort. Since the gain in terms of additional hires from a marginal unit of effort or vacancies is unaffected by γ_2, it is optimal for the firm to use relatively more effort.

Now, substituting the law of motion for employment at the firm level into (A3), we obtain the optimal recruitment effort choice, expressed only as a function of the firm-level variables (n, n'):

$$e\left(n, n'\right) = \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1 - 1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{-\frac{\gamma_2}{\gamma_1+\gamma_2}} \left(\frac{n' - n}{n}\right)^{\frac{\gamma_2}{\gamma_1+\gamma_2}} \tag{A4}$$

which, in turn implies, for the job-filling rate,

$$f\left(n, n'\right) = q\left(\theta^*\right) e\left(n, n'\right) = \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1 - 1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{\frac{\gamma_1}{\gamma_1+\gamma_2}} \left(\frac{n' - n}{n}\right)^{\frac{\gamma_2}{\gamma_1+\gamma_2}}. \tag{A5}$$

This equation demonstrates that the model implies a log-linear relation between the job-filling rate and employment growth at the firm level, with elasticity $\gamma_2/\left(\gamma_1 + \gamma_2\right) < 1$, as in the data.

Finally, substituting (A5) into the firm-level law of motion for employment yields an expression for the vacancy rate

$$\frac{v}{n} = \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1 - 1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{-\frac{\gamma_1}{\gamma_1+\gamma_2}} \left(\frac{n' - n}{n}\right)^{\frac{\gamma_1}{\gamma_1+\gamma_2}}. \tag{A6}$$

Now, note that by substituting the optimal choice for recruitment effort (A3) into (A1), we obtain the following formulation for the cost function:

$$\mathcal{C}\left(n, v\right) = \left[\kappa_2 \frac{\gamma_1 + \gamma_2}{\left(\gamma_1 - 1\right)\left(\gamma_2 + 1\right)}\left(\frac{v}{n}\right)^{\gamma_2}\right] v, \tag{A7}$$

which is one of the specifications invoked by Kaas and Kircher (2015).

Finally, if we use (A6) in (A7), we obtain a version of the cost function only as a function of

(n, n') that we can use directly in the dynamic problem (10):

$$\mathcal{C}^*\left(n, n'\right) = \kappa_2 \left[\frac{\gamma_1 + \gamma_2}{\left(\gamma_1 - 1\right)\left(\gamma_2 + 1\right)}\right] \left\{ \left[\frac{\kappa_2}{\kappa_1}\left(\frac{\gamma_1}{\gamma_1 - 1}\right)\right]^{\frac{1}{\gamma_1+\gamma_2}} q\left(\theta^*\right)^{-\frac{\gamma_1}{\gamma_1+\gamma_2}} \left(\frac{n' - n}{n}\right)^{\frac{\gamma_1}{\gamma_1+\gamma_2}} \right\}^{1+\gamma_2} n.$$

B Additional figures

Figure B1: Dynamics of output under the two shocks

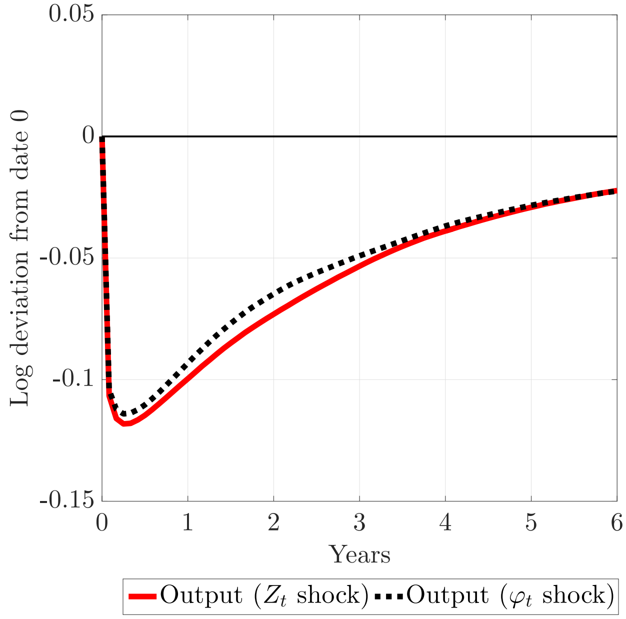

Figure B2: Employment-weighted growth rate distribution in the model

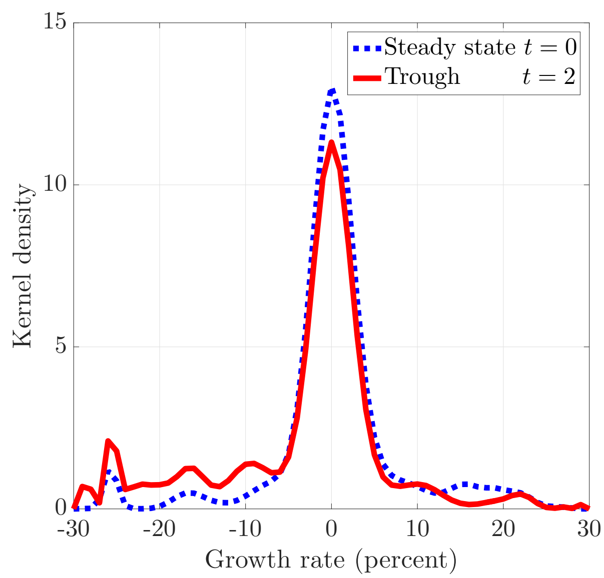

Figure B3: Relative search intensity of employed workers

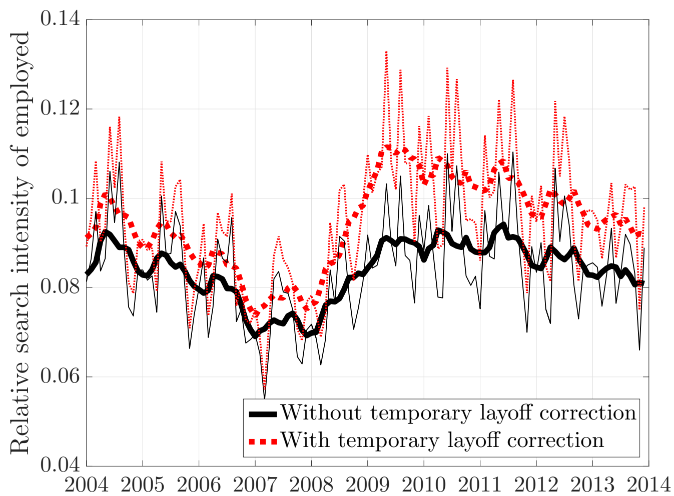

Note: The two series are constructed as the ratio of the EE flow rate to the UE flow rate. In one case, the UE rate is taken directly from the BLS, and in the other case, we adjust the UE rate by subtracting unemployed workers on temporary layoff from the denominator and rehires from temporary layoffs from the numerator.

C Computational details

C.1 Value and policy functions

We use collocation methods to solve the firm's value function problem (4)-(7). Let $s = (n, a, z)$ be the firm's idiosyncratic state, abstracting from heterogeneity in σ which is permanent. We solve for an approximant of the *expected* value function $\mathbf{V}^e(n', a', z)$ which gives the firm's expected value conditional on current decisions for net worth and employment:

$$\mathbb{V}^e(n', a', z) = \int_{\mathcal{Z}} \mathbb{V}(n', a', z') d\Gamma(z, z'),$$

where the integrand is the value given in (6).

We set up a grid of collocation nodes $S = N \times A \times Z$ where $N = \{n_1, \ldots, n_{N_n}\}$, with $N_n = N_a = N_z = 10$. We construct Z by first creating equi-spaced nodes from 0.001 to 0.999, which we then invert through the cumulative distribution function of the stationary distribution implied by the AR(1) process for z to obtain Z. This ensures better coverage in the higher probability regions for z. We choose A and N to have a higher density at lower values. The upper bound for employment, $\bar{n}$, is chosen so that the optimal size of the highest productivity firm $n^*(\bar{z})$ is less than $\bar{n}$. We choose the upper bound for assets, $\bar{a}$, so that the maximum optimal capital $k^*(\bar{z})$ can be financed, that is, $k^*(\bar{z}) < \varphi \bar{a}$. Note that N, A, and Z are parameter dependent, and therefore recomputed for each new vector of parameters considered in estimation.

We approximate $\mathbb{V}^e(s)$ on S using a linear spline with $N_s = N_n \times N_a \times N_z$ coefficients. Given a guess for the spline's coefficients, we iterate towards a vector of coefficients that solve the system of N_s Bellman equations, which are linear in the N_s unknown coefficients. Each iteration proceeds as follows. Given the spline coefficients, we use golden search to compute the optimal policies for all states $s \in S$ and the value function $\mathbb{V}(s)$. We then fit another spline to $\mathbb{V}(s)$, which facilitates integration of productivity shocks $\varepsilon \sim \mathcal{N}(0, \vartheta_z)$. To compute $\mathbb{V}^e(s)$ on S, we approximate the integral by

$$\mathbb{V}^e(n, a, z) = \sum_{i=1}^{N_\varepsilon} w_i \mathbb{V}\left(n, a, \exp(\rho_z \log(z) + \varepsilon_i)\right).$$

Here, $N_\varepsilon = 80$, and the values of ε_i are constructed by creating a grid of equi-spaced nodes between 0.001 to 0.999, then using the inverse cumulative distribution function of the shocks (normal) to create a grid in ε. The weights w_i are given by the probability mass of the normal distribution centered on each ε_i. Note that this differs from quadrature schemes in which one is trying to minimize the number of evaluations of the integrand, usually with N_ε around four. Since $\mathbb{V}(s)$ is already given by an approximant at this step, and the integral is only computed once each iteration, this is not a concern and we compute the integral very precisely. We then fit an updated vector of coefficients to $\mathbb{V}^e(s)$ and continue.[45]

C.2 Stationary distribution

To construct the stationary distribution, we use the method of nonstochastic simulation from Young (2010), modified to accommodate a continuously distributed stochastic state. We create a new, fine grid of points S^f on which we approximate the stationary distribution using a histogram, setting $N_n^f = N_n^f = N_z^f = 100$. Given our approximation of the expected continuation value, we solve for the policy functions $n'\left(s^f\right)$ and $a'\left(s^f\right)$ on the new grid and use these to create two transition matrices $\mathbf{Q}_n$ and $\mathbf{Q}_a$, which determine how mass shifts from points $s^f \in S^f$ to points in N^f and A^f, respectively. We construct $\mathbf{Q}_x$ as follows for $x \in \{a, n\}$:

$$\mathbf{Q}_x[i,j] = \left[\mathbf{1}_{x'\left(s_i^f\right)\in\left[X_{j-1}^f, X_j^f\right]} \frac{x'\left(s_i^f\right) - X_j^f}{X_j^f - X_{j-1}^f} + \mathbf{1}_{x'\left(s_i^f\right)\in\left[X_j^f, X_{j+1}^f\right]} \frac{X_{j+1} - x'(s_i)}{X_{j+1}^f - X_j^f}\right],$$

for $i = 1, \ldots, N_s^f$ and $j = 1, \ldots, N_x^f$.[46] This approach ensures that aggregates computed from the stationary distribution will be unbiased. For example if $x'(s) \in (X_j, X_j + 1)$, then masses w_j and w_{j+1} are allocated to X_j and X_{j+1} such that $w_j X_j + w_{j+1} X_{j+1} = x'(s)$. The transition matrix for the process for z is computed by $\mathbf{Q}_z = \sum_{i=1}^{N_\varepsilon} w_i \mathbf{Q}_z^i$, where $\mathbf{Q}_z^i$ is computed as above under $z'(s^f) = \exp(\rho_z \log z + \varepsilon_i)$. The overall incumbent transition matrix $\mathbf{Q}$ is simply the tensor product $\mathbf{Q} = \mathbf{Q}_z \otimes \mathbf{Q}_a \otimes \mathbf{Q}_n$.

To compute the stationary distribution, we still need the distribution of entrants. To allow for entry cutoffs to move smoothly, we compute entrant policies on a dense grid of $N_z^0 = 500$ productivities. This is clearly important for us since it ensures that entry does not jump in the

[45]In practice, instead of this simple iterative approach to solve for the coefficients, we follow a Newton algorithm as in Miranda and Fackler (2002), which is two orders of magnitude faster. The Newton algorithm requires computing the Jacobian of the system of Bellman equations with respect to the coefficient vector. The insight of Miranda and Fackler (2002) is that this is simple to compute given the linearity of the system in the coefficients.

[46]If exit is optimal on grid point s_i^f, then we set row i of $\mathbf{Q}_x$ to zero.

transition dynamics or across parameters in calibration. The grid Z^0 is constructed by taking an equally spaced grid in $[0.01, 0.99]$ and inverting it through the cumulative distribution function of potential entrant productivities (exponential). Let the corresponding vector of weights be given by $\mathbf{P}_0$. Given the approximation of the continuation value $\mathbb{V}^e(s)$, we can solve the potential entrant's policies $n_0'(s_0)$ and $a_0'(s_0)$, conditional on entry. We can then solve the firm's discrete entry decision. Finally, we compute an equivalent transition matrix $\mathbf{Q}_0$ using these policies, where nonentry results in a row of zeros in $\mathbf{Q}_0$.

The discretized stationary distribution $\mathbf{L}$ on S^f is then found by the following approximation to the law of motion (15)

$$\mathbf{L} = (1-\zeta)\mathbf{Q}'\mathbf{L} + \lambda_0 \mathbf{Q}_0'\mathbf{P}_0,$$

which is a contraction on $\mathbf{L}$, solved by iterating on a guess for $\mathbf{L}$. The final stationary distribution is found by choosing λ_0 such that $\sum_{i=1}^{N_s^f} L_i = 1$.

C.3 Computation of moments

We compute an aggregate moment X by integrating λ over firm policies $x(s)$. Using the above approximation, this is simply $X = \mathbf{L}'x(s)$.

For age-based statistics, our moments in the data refer to firm ages in years. We therefore generate an "age zero" measure of firms by allowing for 12 months of entry. We then iterate this distribution forward to compute age statistics such as average debt to output for age 1 firms or the distribution of vacancies by age.

For statistics such as the average annual growth rate conditional on survival, we need to simulate the model. In this case, we draw 100,000 firms on S^f in proportion to $\mathbf{L}$ and simulate these forward solving (rather than interpolating) firm policies each period and evolving productivity with draws from the continuous distribution of innovations ε. To remove the effect of the starting grid, we simulate for 36 months and compute our statistics, comparing firms across months 24 and 36.

C.4 Estimation

The model has a large number of unknown parameters and a criterion function that is potentially nonsmooth. Furthermore, the model does not have an equilibrium for large regions of the parameter space.[47] For these reasons, using a sequential optimizer that takes the information from successive draws from the parameter space and updates its guess is prohibitive. For example, a Nelder-Mead optimizer both needs to be returned values for the objective function at each evaluation and needs to make many evaluations of the function when taking each "step".

Our solution is to use an algorithm that we can very easily parallelize, that efficiently explores the parameter space, and for which we can ignore cases with no equilibrium. We set up a hyper-cube in the parameter space and then initialize a Sobol sequence to explore it. A Sobol sequence is a quasi-random low-discrepancy sequence that maintains a maximum dispersion in each dimension and far outperforms standard random number generators. We then partition the sequence and submit each partition to a separate CPU on a high performance computer (HPC). From each evaluation of the parameter hyper-cube, we save the vector of model moments and regularly splice these together, choosing one that minimizes the criterion function. Starting with wide bounds on the parameters, we run this procedure a number of times, shrinking the hyper-cube each time.

This procedure has a number of benefits. First, we trade in the optimization steps associated with a traditional solver for scale. Instead of using a 10 CPU machine to run a Nelder-Mead algorithm, we can simultaneously solve the model on 300+ CPUs. Second, the output of the exercise gives a strong intuition for the identification of the model. From an optimizer one may retrieve the moments of the model along the path of the parameter vector chosen by the algorithm. In our case, we retrieve thousands of evaluations, knowing that the low-discrepancy property of the Sobol sequence implies that for an interval of any one parameter, the remaining parameters are drawn uniformly. Plotting moments against parameters therefore shows the effect of a parameter on a certain moment, conditional on local draws of all other parameters. Plotting a histogram of the moments returned, as in Figure C1, gives a strong indication as to which moments may be difficult to match for the current bounds of the parameter space.

C.5 Transition dynamics

We solve for transition dynamics as follows. Consider the case of a shock to aggregate productivity Z. We specify a path for $\{Z_t\}_{t=0}^{T}$ with $Z_0 = Z_T = \bar{Z}$. Given a conjectured path for

[47] For example, if the value of home production is very low then unemployment derived from the labor demand condition may be *negative*. Wages are so low that labor demand eclipses the fixed supply $\bar{L}$.

Figure C1: Histograms of selected moments from final calibration exercise

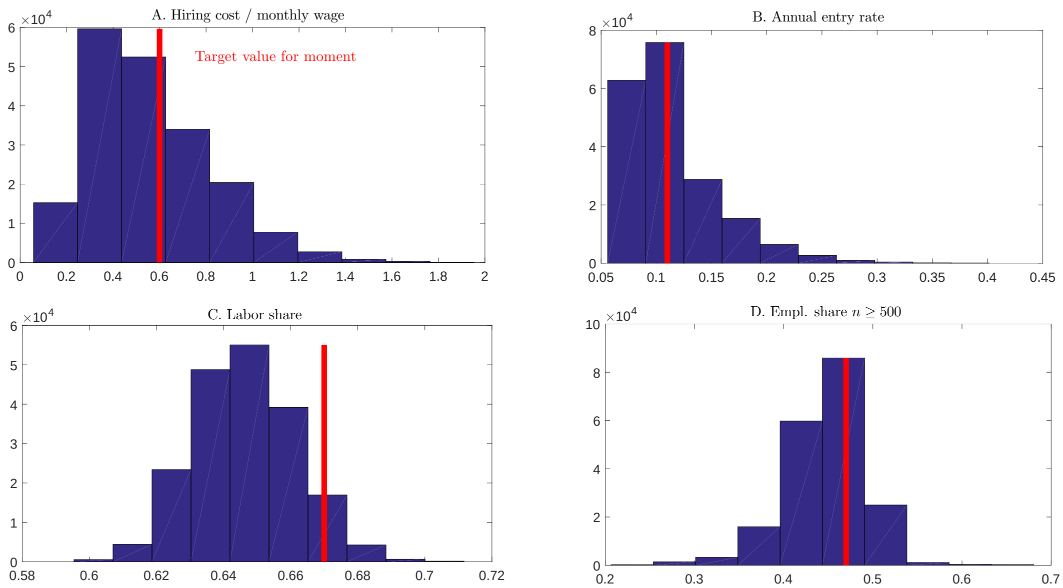

equilibrium market tightness $\{\tilde{\theta}_t^*\}_{t=0}^{T}$ and the assumption that the date T continuation values of the firm are the same as they are in steady state, one can solve *backward* for expected value functions $\mathbf{V}_t^e$ at all dates $T-1, T-2, \ldots, 1$. Setting the aggregate states $U_0 = \bar{U}$ and $\lambda_0 = \bar{\lambda}$, and using the conjectured path θ_t^*, the shocks, and continuation values, one can then solve *forward* for a new market clearing $\theta_t^{*,\prime}$ that equates unemployment from labor demand U_t^{demand} and worker flows U_t^{flows} in every period using the labor demand and evolution of unemployment equations,

$$\begin{aligned} U_{t+1}^{flows} &= U_t - H(\theta_t^*) + F(\theta_t^*) - \lambda_{e,t} n_0 \\ U_{t+1}^{demand} &= \bar{L} - \int n'(s, \theta_t^*, A_t, \mathbf{V}_t^e) d\lambda_t \end{aligned}$$

Once we reach $t = T$, we set $\tilde{\theta}_t^* = \theta_t^*$ and iterate until the proposed path and equilibrium path for market tightness converge.