



HAL
open science

Wage Posting and Business Cycles: a Quantitative Exploration

Giuseppe Moscarini, Fabien Postel-Vinay

► **To cite this version:**

Giuseppe Moscarini, Fabien Postel-Vinay. Wage Posting and Business Cycles: a Quantitative Exploration. 2014. hal-03460143

HAL Id: hal-03460143

<https://sciencespo.hal.science/hal-03460143>

Preprint submitted on 1 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Wage Posting and Business Cycles: a Quantitative Exploration

Giuseppe Moscarini*
Yale University
and
NBER

Fabien Postel-Vinay†
University College London
and
Sciences Po

December 2014

Abstract

We provide a quantitative exploration of business cycles in a frictional labor market under contract-posting. The steady-state random search and wage-posting model of Burdett and Mortensen (1998) has become the canonical structural framework for empirical analysis of worker turnover and equilibrium wage dispersion. In this paper, we provide an efficient algorithm to simulate a dynamic stochastic equilibrium version of this model, the Stochastic Burdett-Mortensen (SBM) model, and evaluate its performance against empirical evidence on fluctuations in unemployment, vacancies and wages.

Keywords: Equilibrium Job Search, Dynamic Contracts, Stochastic Dynamics, Business Cycles.

JEL codes: J64, J31, E32.

*Address: Department of Economics, Yale University, PO Box 208268, New Haven CT 06520-8268. Tel. +1-203-432-3596. E-mail giuseppe.moscarini@yale.edu. Web <http://www.econ.yale.edu/faculty1/moscarini.htm>. Moscarini thanks the NSF for support to this research under grant SES 1123021.

†Address: Department of Economics, University College London, Drayton House, 30 Gordon Street, London WC1H 0AX, UK. Tel: +44 117 928 8431. E-mail fabien.postelvinay@gmail.com. Web <https://sites.google.com/site/fabienpostelvinay/>. Postel-Vinay is also affiliated with the Centre for Macroeconomics (CfM, London), CEPR (London) and IZA (Bonn).

1 Introduction

Why are similar workers paid differently? In his 2003 book of this title, Dale Mortensen takes stock of a few decades of investigation of this question, that he had jumpstarted and then developed. His answer is simple: imperfect competition in labor markets. Information and other frictions, which make the outcome of job search time-consuming and random, endows firms with monopsony power that they exploit, in the spirit of Coase (1972), by committing to wage offers. This force depresses all wages towards the opportunity cost of work. But workers cannot commit to their current wage offer, and, while employed, search for better outside offers. In this environment, firms choose a wage policy that balances unit labor costs with hiring and retention. As a result, in equilibrium wages *must* differ among identical firms and workers. In the presence of heterogeneity in productivity and in demand conditions among workers and firms, equilibrium wages still fall short of marginal products, and contain a non-fundamental component of “frictional” inequality.

Burdett and Mortensen (1998, thereafter BM) formalized this powerful insight. Their working paper, first circulated in the 1980s, spurred the vast theoretical and empirical literature culminating in Dale Mortensen’s 2003 book. The BM “wage posting” model quickly emerged as the canonical framework for the analysis of wage inequality, labor turnover, and unemployment. Each of the three exists in conjunction with the other two. Naturally, the scope of this line of research eventually transcended wage inequality alone.

In a series of articles (Moscarini and Postel-Vinay 2009, 2011, 2012, 2013, 2014, thereafter MPV09-MPV14, resp.) we explored both theoretically and empirically the business cycle implications of the wage posting paradigm. Progress in this direction had been stunted by technical difficulties in finding equilibrium in an economy where the law of one price fails. The other canonical model of the labor market now known as “DMP” (Diamond, 1982; Pissarides, 1985; Mortensen and Pissarides, 1994) bypassed this hurdle by assuming that trading partners bargain over their match surplus, hence wages are not allocative. This search-and-matching model still encodes the leading theory of equilibrium unemployment, but runs into difficulties when applied to business cycles. As Shimer (2005) demonstrated, this model cannot reconcile the large cyclical swings in job finding and unemployment rates with the tiny ones in Average Labor Productivity (ALP), or in other plausible sources of aggregate shocks, that we observe in the US economy. The perfectly competitive labor market model failed this test because it requires an implausibly elastic aggregate labor supply; the same issue came back to haunt the search-cum-bargaining model (Hagedorn and Manovskii, 2008). The attention then turned to other sources of wage rigidity. We add to this range of new hypotheses. The simple Coasian assumption of commitment to

wage offers to exploit market power, here conferred by frictions and tempered by on-the-job search, is a natural source of wage rigidity, in an environment that can also explain wage inequality and reallocation and that is very well understood in steady state since BM.

In our past theoretical work, we outlined the scope and limitations of wage posting models with random search in the presence of aggregate shocks to labor productivity. Our main empirical focus was on the cyclical reallocation of employment among heterogeneous firms (MPV12). Our contribution here is to evaluate the quantitative performance of our MPV13 business cycle wage-posting model against empirical evidence on unemployment and wage fluctuations, as well as on wage inequality and the pace of reallocation as is standard in this approach. To that end, we propose a tractable, stochastic equilibrium version of the BM model — we will refer to it as the “Stochastic BM” (SBM) model — and an operational algorithm to simulate it.

Unlike in bargaining models, in wage- (or, more generally, contract-) posting models such as (S)BM, equilibrium wages are allocative, so they can be unambiguously related to empirical evidence on earnings inequality, dynamics, and cyclicalities. Different varieties of wage-posting models with on-the-job search successfully introduced aggregate shocks, by making one key change to the environment. Menzies and Shi (2011) assume perfect information about posted wages, in the tradition of directed or competitive search, and study business cycle movements in labor market quantities, but not in wages. Closer to our exercise, Robin (2011) maintains random search, but relaxes the full commitment assumption, to introduce Postel-Vinay and Robin (2002)’s sequential auctions. Firms can respond to outside offers to their employees, which critically changes the ex ante incentives of forward-looking firms and workers to post and to accept wage offers. This assumption also greatly simplifies the analysis of business cycles. Robin addresses empirical evidence on both labor market flows and wages. Relative to these two alternatives, computation of equilibrium wages in our SBM model is less straightforward. Yet, as we show, it is still feasible and reasonably fast using an algorithm that we offer here. We build on the simple structure of equilibrium, which is Rank-Preserving (MPV09): more productive and larger firms always offer higher values to all workers, independently of the history of aggregate shocks. Thus, workers always move in the same direction between jobs, equilibrium turnover is trivial to simulate, and generates empirically accurate predictions about job ladder movements over business cycles (MPV14). Our harder task now is to compute the equilibrium contracts, or state-contingent wages, that implement this equilibrium allocation.

The broader goal of this project is to provide a unified explanation, based on a stochastic job ladder, for both unemployment fluctuations and wages, individual dynamics and cyclicalities of earnings and turnover. This is a very ambitious goal. In this first quantitative step, to

give our model a reasonable chance, we introduce a seemingly minor but important change relative to MPV13. Following Pissarides (2009)’s suggestion, we model adjustment frictions on the firm side as a cost that depends on the volume of hires, and not of vacancies or job adverts, as is customary. That is, the firm pays for the output of its recruiting activities, not for the inputs into it. As such, hiring costs are best thought of as training costs. Search for trading partners is still mediated by a matching function. Modeling recruitment costs as hiring costs tames congestion effects, which facilitate hiring in recessions, when unemployment is abundant, thus mute the negative impact of aggregate shocks on job creation (see Christiano et al., 2013). This change in the model requires a new argument to prove that equilibrium remains unique and RP. We find that this property, essential for tractability, requires a restriction on the convexity of the hiring cost function.¹

We gauge our SBM model’s quantitative performance along three main dimensions. First, we assess the model’s ability to amplify TFP shocks, based on its predictions about the volatility and covariances of unemployment, the job finding rate and the vacancy-unemployment (V/U) ratio, in the face of TFP shocks of a plausible magnitude. Second, we examine wage flexibility in the model, as measured by the wage-unemployment semi-elasticity (on which we incidentally provide new evidence from the Survey of Income and Program Participation — SIPP), and the time-series volatility of mean log wages. Third, we re-examine the BM model’s predictions about cross-sectional dispersion in wages, profits, and employer sizes — all issues that have been the focus of many analyses of the steady-state BM model. In so doing, we highlight new connections between the model’s cross-sectional and business-cycle predictions.

Our summary assessment of the successes and failures of this SBM model is as follows. A clear success of the SBM model is its ability to generate plausible amplification of TFP shocks. For example, the predicted volatility ratio of the job finding rate relative to ALP, which has been the focus of much attention in the literature, can be as high 15, well in line with the data. Yet, wages in the SBM model are much more flexible than they appear to be in the data: the model-based wage-unemployment semi-elasticity is up to ten times larger than the one we estimate from SIPP data. How, then, does the model generate plausible amplification? First, even though wages are very flexible, *worker values*, which are what agents care about and what, ultimately, matters for the allocation, are much less so. Second, and most importantly, the model generates amplification by concentrating hiring amongst firms at the bottom of the productivity ladder. Those low-productivity firms generate a

¹Coles and Mortensen (2012) adopt a hiring cost function in the same spirit as ours and impose even more stringent conditions to extend the scope of our early work on transitional dynamics (MPV09). They establish existence and characterization of one RP equilibrium also in the presence of idiosyncratic, firm-level TFP shocks.

low surplus, so that even with flexible wages, small fluctuations in productivity cause large (proportional) variations in profitability (and consequently on hiring) for those firms (as per the intuition put forward in a different context by Hagedorn and Manovskii, 2008).

This amplification mechanism highlights a trade-off between aggregate volatility of unemployment and cross-sectional dispersion in employer size. We achieve concentration of hiring at low-productivity firms by assuming a very highly convex recruitment (or training) cost function. This strong degree of diminishing returns to hiring tames the incentives of highly-productive firms to hire much more than low-productivity firms. Dispersion in gross hiring flows between high- and low- productivity firms is therefore very small and, as a consequence, so is dispersion in equilibrium size between firms. The largest firm in the simulated economy is only about four times the size of the smallest one, a number which is obviously nowhere near what is observed in the data. This tension between size dispersion and amplification is arguably the key link between the model's cross-sectional and dynamic predictions.

Next, the SBM model tends to understate cross-sectional wage dispersion. This is not a new finding: as pointed out by Hornstein, Krusell and Violante (2011), if the residual wage dispersion observed in US data was truly the result of search frictions, a jobless worker searching for jobs should wait to sample a high paying job. But average unemployment duration in the U.S., as measured in the Current Population Survey, is short.² On-the-job search helps to resolve this tension, at least qualitatively, because accepting quickly a low offer does not 'burn' all of the option value of looking for a better wage. But estimation of the steady-state BM model found that the opportunity cost of work must be extremely low to reconcile empirical observations on the pace of transition and on wage inequality. Like its steady-state relative, the SBM model is affected by those issues, which are largely orthogonal to time-series fluctuations.

Finally, as we highlighted in previous work (MPV13), the model offers a natural explanation for our observation (MPV12) that net job creation is more negatively correlated with aggregate unemployment at large than at small firms. In the SBM model this still holds, although, as mentioned, exactly the opposite is true of gross job creation, which is the source of aggregate volatility in job finding rates. Intuitively, in an aggregate expansion, small employers hire proportionally more new workers than larger competitors, but lose even more employees to their poaching, so their employment grows more slowly on net. The opposite

²Using data from the Survey of Income and Program participation (SIPP), Fujita and Moscarini (2013) show that unemployment duration in the US is shortest for the large share of workers who end up being recalled by their former employers, and is in fact much longer than commonly thought, even in good times, for those unemployed workers who are hired by new employers, which is presumably what search models are really about. The CPS does not contain information that allow to make this distinction.

occurs in recessions, the Great Recession being a perfect example (MPV14).³ The disconnect in the model between the cyclical nature of gross and net job creation among heterogeneous firms is critical. Bilal et al. (2012) show, in a DMP-style model, that a similar tension between aggregate volatility and cross-sectional inequality on the worker side cannot be resolved. They observe that the importance of small surplus emphasized by Hagedorn and Manovskii (2008) applies to the marginal, and not to the average, labor supply. Hence, they ask whether worker ex ante skill heterogeneity can amplify the effect of aggregate shocks on the average job finding rate, by concentrating job finding volatility among low-skill individuals, whose surplus from work over leisure is small. They show that this is indeed the case only if the calibrated model also delivers job finding rates for high wage workers that are far too sticky relative to the data. Just like large, productive firms exhibit surprisingly cyclical net job creation rates, so high-wage, productive workers face surprisingly cyclical exit rates from unemployment.

The rest of the paper is organized as follows. Section 2 introduces the SBM model and characterizes (conditions) for equilibrium to be RP, hence computable. Section 3 studies steady state equilibrium and its comparative statics properties when changing aggregate productivity. This is a warm-up towards our quantitative analysis, which is based on empirical moments presented in Section 4, and is illustrated in Section 5, calibration and simulation results. After some concluding remarks, the Appendix contains the proofs of the formal arguments and presents the structure of the computation algorithm.

2 A wage-posting model of the business cycle

2.1 The environment

Our Stochastic BM (SBM) model is a variant of MPV13's business cycle model of a frictional labor market with random search and wage posting. In this version, we introduce firm-level hiring costs that depend on the output of the hiring process, namely on the flow of hires, rather than on the input, vacancies or job ads. So hiring costs are best thought of as a combination of search and training costs. As emphasized by Pissarides (2009), hiring costs are independent of competition in the market, hence, unlike vacancy costs, they do not fall in recessions, when labor demand is lower, and do not mute the effect of aggregate shocks.⁴

³Kahn and McEntarfer (2014) use matched employer-employee US data to order employers on the job ladder by the average wage they pay, and find direct evidence for this cyclical poaching pattern.

⁴This assumption is taken up in a growing number of contributions to the topic at hand, including the closely related paper by Coles and Mortensen (2013). In their paper, recruitment costs are specified as $Lc(h/L)$, where L is initial firm size and h is the inflow of new hires, so that h/L is the number of new hires per incumbent worker in the firm.

Time $t = 0, 1, 2 \dots$ is discrete. The labor market is populated by a unit-mass of workers, who can be either employed or unemployed, and by a unit measure of firms.⁵ Workers and firms are risk neutral, infinitely lived, and maximize payoffs discounted with common factor $\beta \in (0, 1)$. Firms operate constant-return technologies with labor as the only input and with productivity scale $\omega_t p$, where $\omega_t \in \Omega$ is an aggregate component, evolving according to a stationary first-order Markov process $Q(d\omega_{t+1} | \omega_t)$, and p is a fixed, firm-specific component, distributed across firms p according to a c.d.f. Γ over some positive interval $[\underline{p}, \bar{p}]$.

The labor market is affected by search frictions in that unemployed workers can only sample job offers sequentially with some probability $\lambda_t \in (0, 1]$ at time t and, while searching, enjoy a value of leisure b_t . Employed workers earn a wage and also sample job offers with probability $s\lambda_t \in (0, 1]$ each period, so that s is the search intensity of employed relative to unemployed job seekers. Workers receive at most one job offer per time period. Each employed worker is separated from his employer and enters unemployment every period with probability $\delta_t \in (0, 1]$. All firms of equal productivity p start out with the same labor force $L_0(p)$. We denote by $L_t(p)$ the size of a firm of type p at time t on the equilibrium path, and $N_t(p) = \int_{\underline{p}}^p L_t(x) d\Gamma(x)$ the cumulated population distribution of employment across firm types. So $N_0(p)$ is the (given) initial measure of employment at firms of productivity *at most* p , $N_t(\bar{p})$ is employment and $u_t = 1 - N_t(\bar{p})$ the unemployment (rate) at time t .

We maintain throughout the assumption that the destruction rate is exogenous and a function of the aggregate productivity state $\delta_t = \delta(\omega_t)$. Similarly for the flow value of non production $b_t = b(\omega_t)$. The job-contact probability λ_t instead is determined in equilibrium by a matching function. Each period, the firm can hire h workers at cost $c(h)$, with $c(\cdot)$ positive, strictly increasing and convex, continuously differentiable. To do so, before workers have a chance to search, a firm can post $a \geq 0$ job adverts (vacancies). Own job adverts determine the firm's sampling weight in workers' job search, while total job adverts determine the rate at which any one advert returns contacts with workers. Let $a_t(p)$ denote the adverts posted on the equilibrium path by a firm of productivity p , size $L_{t-1}(p)$, and define aggregate adverts A_t and aggregate search effort by workers Z_t as

$$\begin{aligned} A_t &= \int_{\underline{p}}^{\bar{p}} a_t(p) d\Gamma(p) \\ Z_t &= 1 - N_{t-1}(\bar{p}) + (1 - \delta_t) s N_{t-1}(\bar{p}). \end{aligned} \tag{1}$$

The latter adds the previously unemployed to the previously employed who are not displaced and draw a chance to search this period t . In each time period, employed and unemployed

⁵A firm can be inactive when its productivity is too low relative to the worker value of leisure. So the unit measure of firms includes all potential producers, active and inactive. That the mass of firms and workers both have measure one is obviously innocuous and only there to simplify the notation.

search simultaneously. Then:

$$\eta_t A_t = \lambda_t Z_t = m(A_t, Z_t) \leq \min \langle A_t, Z_t \rangle \quad (2)$$

where η_t is the chance for any advert to contact a worker, $m(\cdot)$ is a linearly homogeneous matching function, increasing and concave in each argument.

The timing within a period is as follows. Given a current state ω_t of aggregate labor productivity and distribution of employed workers N_t :

1. firms produce and sell output and pay workers in state ω_t ; the flow benefit b_t accrues to unemployed workers;
2. the new state ω_{t+1} of aggregate labor productivity is realized;
3. employed workers can quit to unemployment;
4. filled jobs are destroyed exogenously with chance δ_{t+1} ;
5. firms post job adverts a_{t+1} ;
6. the remaining employed workers receive an outside offer with chance $s\lambda_{t+1}$ and decide whether to accept it or to stay with the current employer; simultaneously, each previously unemployed worker receives an offer with probability λ_{t+1} ;
7. firms hire workers and pay hiring costs that depend on how many workers they hire.

Finally, in order to avert unnecessary complications, and to simplify the illustration, we assume that the state space Ω is finite, the distribution of firm types, Γ , has continuous and everywhere strictly positive density $\gamma = \Gamma'$ over $[\underline{p}, \bar{p}]$, and the initial measure of employment across firm types, N_0 , is continuously differentiable in p .

2.2 The firm's contract-posting and hiring problem

Each firm of type p chooses and commits to an employment contract, namely a state-contingent wage $w_t(p)$ depending on aggregate productivity ω_t and the employment distribution $N_{t-1}(\cdot)$, to maximize the present discounted value of profits at time 0, given other firms' contract offers. The dependence of the wage on the state $\{\omega_t, N_{t-1}(\cdot)\}$ is marked as a shorthand by the time index of the wage. The firm is further subjected to an *equal treatment constraint*, whereby it must pay the same wage to all its workers. Under commitment, such a wage function implies an equilibrium value $V_t(p)$ for any worker to work for that firm.

We introduce some notation. Let

$$F_t(W) = \frac{1}{A_t} \int_{\underline{p}}^{\bar{p}} \mathbb{I}\{V_t(p) \leq W\} a_t(p) d\Gamma(p). \quad (3)$$

denote the c.d.f. of values posted by all firms and offered to searching workers, with $\bar{F}_t(\cdot) = 1 - F_t(\cdot)$, i.e. F_t is the distribution of values from which job searchers draw from,

$$G_t(W) = \frac{1}{N_t(\bar{p})} \cdot \int_{\underline{p}}^{\bar{p}} \mathbb{I}\{V_t(p) \leq W\} dN_t(p) \quad (4)$$

denote the c.d.f. of values accruing to the currently employed workers, and let

$$U_t = b_t + \beta \mathbf{E}_t \left[(1 - \lambda_{t+1}) U_{t+1} + \lambda_{t+1} \int \max\langle v, U_{t+1} \rangle dF_{t+1}(v) \right] \quad (5)$$

be the value of unemployment. The unemployed worker collects a flow value b_t and, next period, when aggregate productivity becomes ω_{t+1} , she draws with chance λ_{t+1} a job offer from the distribution of offered values F_{t+1} , that she accepts if its value exceeds that of staying unemployed. Each firm now has a sampling weight in F equal to its (normalized) job posting, a/A .

A firm that observes state ω_{t+1} and decides to post a continuation value $W_{t+1} < U_{t+1}$ loses all workers, who quit to unemployment, so $L_{t+1} = 0$. Otherwise, a firm of current size L_t posting *any* value $W_{t+1} \geq U_{t+1}$ in aggregate state $\{\omega_{t+1}, N_t(\cdot)\}$ loses workers to unemployment, with chance δ_{t+1} , and to other firms, if its workers draw offers, with chance $s\lambda_{t+1}$, which are more valuable, with chance $\bar{F}_{t+1}(W_{t+1})$. The firm chooses to hire h_{t+1} workers. Formally:

$$L_{t+1} = L_t (1 - \delta_{t+1}) (1 - s\lambda_{t+1} \bar{F}_{t+1}(W_{t+1})) + h_{t+1} \quad (6)$$

where hires $h_{t+1} = a_{t+1} \eta_{t+1} P_{t+1}(W_{t+1})$ equal the measure of vacancies posted times the contact rate of each vacancy times the probability that the offer is accepted:

$$P_{t+1}(W_{t+1}) = \frac{1 - N_t(\bar{p}) + s(1 - \delta_{t+1}) N_t(\bar{p}) G_{t+1}(W_{t+1})}{Z_{t+1}} \quad (7)$$

In (7), the denominator is the measure of workers who can make contact, and the numerator counts only those who accept the offer, namely all the unemployed and only the fraction of employed who will earn less than W_{t+1} by staying where they are.

Two payoff-relevant state variables of the firm's problem are aggregate productivity ω_t and the distribution of values offered in the market, F_t , which is infinitely-dimensional and determines the retention effects of posting a contract value. These two state variables are

aggregate and exogenous to the firm. For notational simplicity, we again subsume in a time index the dependence on (ω_t, F_t) of firm's and workers' values.

The firm's problem can be formulated recursively (Spear and Srivastava, 1987) by introducing an additional, fictitious state variable, namely the continuation utility \bar{V} that the firm promised at time $t - 1$ to deliver to the worker from this period t on. So the firm solves

$$\Pi_t(\bar{V}, L_t) = \sup_{\substack{w_t \geq 0 \\ h_{t+1} \geq 0 \\ W_{t+1} \geq U_{t+1}}} \left\langle (\omega_t p - w_t) L_t + \beta \mathbf{E}_t [-c(h_{t+1}) + \Pi_{t+1}(W_{t+1}, L_{t+1})] \right\rangle \quad (8)$$

subject to the law of motion (6) of firm size and a Promise-Keeping (PK) constraint to deliver the promised \bar{V} :

$$\bar{V} = w_t + \beta \mathbf{E}_t \left[\delta_{t+1} U_{t+1} + (1 - \delta_{t+1}) \left(1 - s \lambda_{t+1} \bar{F}_{t+1}(W_{t+1}) \right) W_{t+1} + (1 - \delta_{t+1}) s \lambda_{t+1} \int_{W_{t+1}}^{+\infty} v dF_{t+1}(v) \right]. \quad (9)$$

Expectations are taken with respect to the future realization ω_{t+1} of aggregate productivity and consequent employment distribution N_{t+1} , conditional on the date- t state variable (ω_t, N_t) .

To characterize the best response contract, we first describe an equivalent unconstrained recursive formulation of the contract-posting problem. We define the joint value of the firm and its *existing* workers:

$$S_t = \Pi_t + \bar{V} L_t.$$

Solving for the wage w_t from (9), replacing it into the firm's Bellman equation (8), and using (6) to replace the expression for L_{t+1} into $S_{t+1} = \Pi_{t+1} + W_{t+1} L_{t+1}$ in each future state, the joint value function S_t solves the unconstrained, recursive maximization of the joint value of the firm-worker collective:

$$S_t(p, L_t) = \omega_t p L_t + \beta \mathbf{E}_t \left[\delta_{t+1} L_t U_{t+1} + \sup_{\substack{h_{t+1} \geq 0 \\ W_{t+1} \geq U_{t+1}}} \left\langle (1 - \delta_t) s \lambda_t L_t \int_{W_{t+1}}^{+\infty} v dF_{t+1}(v) - c(h_{t+1}) + S_{t+1}(p, L_{t+1}) - W_{t+1} h_{t+1} \right\rangle \right] \quad (10)$$

The joint value S_t to the firm and its existing L_t employees equals flow output, $\omega_t p L_t$, plus the discounted expected continuation value. This includes (in order) the value of unemployment for those employees who are displaced exogenously, the value of a new job for those who are not displaced and find a better offer than the one extended by the current firm, minus the

cost of hiring new workers, and, on the second line, the joint continuation value of the firm and of its current (time t) employees. In turn, the latter equals the joint continuation value S_{t+1} of the firm and its future workforce — made up of stayers among the current (date- t) workforce plus next-period (date- $t+1$) hires — minus the value to be paid to new hires, either from unemployment or from other firms.⁶ The optimal policy solving the unconstrained DP problem (10) also solves (8) subject to (9). We therefore focus on the analysis of the simpler problem (10).

Definition 1 A MARKOV (CONTRACT-POSTING) EQUILIBRIUM is a pair of measurable functions (V, H) of firm-specific productivity p , firm size L , aggregate productivity ω , and aggregate distribution of employment N such that, for every firm type p , if all other firms of type x play (V, H) so that (3), (4) and (6) hold with $W_t = V_t(x)$ where $V_t(x) = V(x, L_{t-1}(x), \omega_t, N_{t-1})$ and $h_t = H_t(x)$ where $H_t(x) = H(x, L_{t-1}(x), \omega_t, N_{t-1})$, the value-posting function $V_t(p)$, the wage function that implements it (i.e. solves (9) with $\bar{V} = V_t(p)$) and the hiring function $H_t(p)$ are the optimal policies of the contracting problem (8).

An equilibrium is a fixed point, a solution (V, H) to this DP problem that coincides with the strategy followed by the other firms. The reason why strategies depend on the employment distribution N_t , rather than the offer distribution F_t , is simple. Given the equilibrium value-offer strategy V and the employment distribution N_t , any firm of type p can compute for every competitor x the probability $G_t(V_t(x))$ from (4) and thus the acceptance probability $P_t(V_t(x))$ from (7). Given the equilibrium hiring strategy H , firm p can then derive the required vacancy posting $a_t(x) = H_t(x) / \eta_t P_t(V_t(x))$ of any other firms x . Finally, F_t derives from (3). So knowledge of equilibrium strategies (a requirement of equilibrium) and of the employment distribution N_t suffice to characterize all other infinite-dimensional state variables.

2.3 Rank-Preserving Equilibrium (RPE)

2.3.1 Definition of RPE

The main difficulty in characterizing equilibrium in BM's environment in the presence of aggregate uncertainty, which hampered progress of this approach to the analysis of business cycles, is now clear. The offer distribution F_t is relevant to the firm maximization problem,

⁶If we did not subtract this cost of employing new hires, this Bellman equation would generate the joint value of the firm and *all* of its workers, current *and* future. If this were the object maximized by the firm, it would optimally offer its workers the maximum value, i.e. pay a wage equal to productivity (the proof, omitted, is available upon request). As is standard, the efficient solution to a moral hazard problem is to "sell the firm to the workers". In our economy, however, firms do not pursue efficiency, but maximize profits. Therefore, the optimal value-offer policy is an interior solution.

because it determines the retention rate $F_t(W)$ of the firm's employees who are offered a contract of value W . This is a defining property of a random search, wage posting model. Hence F_t , or alternatively the employment distribution N_t , is a state variable of the firm problem. Even if the firm does not directly control this aggregate state variable, keeping track of it is an infinitely complex problem.

A particularly simple class of equilibria circumvents this hurdle.

Definition 2 *A RANK-PRESERVING EQUILIBRIUM (RPE) is a Markov equilibrium such that, on the equilibrium path, a more productive firm always offers its workers a higher continuation value: $V_t(p) = V(p, L_{t-1}(p), \omega_t, N_{t-1})$ is increasing in p , including the effect of p on firm size $L_{t-1}(p)$.*

As a direct consequence of the above definition, in a RPE, workers rank their preferences to work for different firms according to firm productivity at all dates. The RP property holds in the unique steady-state equilibrium characterized by BM. Definition 2 proposes to extend it to a dynamic stochastic equilibrium with endogenous hiring.

2.3.2 A characterization result

The key contribution of MPV13 is the proof that, in BM's original environment with exogenous (but stochastic) contact rates, any Markov equilibrium must be Rank-Preserving under the weak sufficient condition that more productive firms start weakly larger at time 0. Thus, for any history of realizations of aggregate shocks, more productive and larger firms always offer their workers a higher value, which is the only available tool to control both hiring and retention. Therefore, the size ranking of firms is preserved over time: if $L_t(p) \geq L_t(p')$ for $p > p'$ at date t , then $L_{t+1}(p) > L_{t+1}(p')$. Thus, by induction, $L_{t+n}(p) > L_{t+n}(p')$ for all positive integers n : if a more productive firm is ever larger, it will remain larger for ever no matter what happens to aggregate productivity, because it will offer a higher value to retain its workers and hire more workers. Furthermore, the equilibrium evolution of the employment distribution is the same in any RPE, so it is uniquely pinned down.

In this paper, we add a hiring choice h to the firm's problem. The RP property alone therefore no longer suffices to ensure that more productive firms also hire more and, consequently, always remain larger. Offering a higher value only guarantees a better retention of existing employees, a reduction in the outflow, but firm size also depends on the gross hiring inflow h . Therefore, the conditions for the unique equilibrium of any kind to be RP as in MPV13, are slightly stronger, and the argument for the preservation of size ranking is not as straightforward. Yet we show in the following proposition that preservation of size ranking continues to hold in that case.

Proposition 1 (RPE with Hiring Costs) *Assume firms incur hiring costs $c(h)$ to hire h workers, where $c(\cdot)$ is a \mathcal{C}^2 function with $c'(h) > 0$ for all $h > 0$. Then any Markov contract-posting Equilibrium such that the sampling c.d.f. $F_t(\cdot)$ is everywhere differentiable at all dates is Rank-Preserving if:*

1. *more productive firms are weakly larger at the initial date, i.e. $L_0(p)$ is non-decreasing*
2. *the elasticity of $c'(\cdot)$, $\varepsilon(c') = hc''(h)/c'(h)$, is everywhere larger than 1.*

Moreover, in any such RPE, size ranking is preserved at all dates, i.e. $L_t(p)$ is non-decreasing at all t .

Proof. See Appendix A. □

Proposition 1 extends the characterization result of MPV13 to the case of endogenous hiring, where firms face recruitment costs that are a function of their gross hiring inflow. It states (a pair of) sufficient conditions for any Markov equilibrium of our model, within the class of equilibria with differentiable sampling c.d.f.'s, to be RP.⁷

2.3.3 Some useful properties of RPE

Three properties thus hold true in any RPE. First, the share of firms that offer less than $V_t(p)$ is simply the proportion of firms that are less productive than p weighted by their relative vacancy postings

$$F_t(V_t(p)) = \frac{1}{A_t} \int_p^p a_t(x) d\Gamma(x). \quad (11)$$

Second, the share of employed workers who earn a value that is lower than that offered by p equals the share of employment at firms less productive than p :

$$G_t(V_t(p)) = \frac{N_{t-1}(p)}{N_{t-1}(\bar{p})}. \quad (12)$$

Third, the probability that an offer is accepted equals

$$P_{t+1}(V_{t+1}(p)) = \frac{1 - N_t(\bar{p}) + s(1 - \delta_{t+1})N_t(p)}{1 - N_t(\bar{p}) + s(1 - \delta_{t+1})N_t(\bar{p})} := Y_t(p). \quad (13)$$

Therefore $a_t(p) = H_t(p)/(\eta_t Y_t(p))$, and we can write the sampling weight as a function of the chosen hiring flow. As we will see shortly, these restrictions will drastically simplify the computation of equilibrium in the stochastic model.

⁷The restriction to differentiable c.d.f.'s simplifies the proof, but we suspect it is inessential - it is not imposed in MPV13.

In MPV13 we further establish that on the RPE path the value $V_t(p)$ offered to workers is continuous in p and the joint value $S(p, L, \omega_t, N_t)$ of a firm of type p and of its L employees is differentiable in L . Our focus in this paper is on computation. Therefore, we restrict attention to differentiable equilibria without formally extending MPV13's differentiability proof to the case of endogenous hiring. Exploiting the differentiability properties of the equilibrium we are seeking to characterize, we define the costate variable

$$\mu_t(p) = \frac{\partial S_t}{\partial L}(p, L(p)) = \frac{\partial S}{\partial L}(p, L(p), \omega_t, N_t)$$

which is the shadow marginal value of an additional worker to the firm-employees collective. Using the RP implications (11) and (12) in the Bellman equation (10), differentiating the latter on both sides with respect to L_t , and invoking the Envelope theorem:

$$\begin{aligned} \mu_t(p) = \omega_t p + \beta \mathbf{E}_t \left[\delta_{t+1} U_{t+1} + (1 - \delta_{t+1}) \left(1 - \frac{s\lambda_{t+1}}{A_{t+1}} \int_p^{\bar{p}} a_{t+1}(x) d\Gamma(x) \right) \mu_{t+1}(p) \right. \\ \left. + (1 - \delta_{t+1}) \frac{s\lambda_{t+1}}{A_{t+1}} \int_p^{\bar{p}} a_{t+1}(x) V_{t+1}(x) d\Gamma(x) \right] \end{aligned} \quad (14)$$

where in the last integral we changed variable from value $V \sim F_t$ to productivity $p \sim \Gamma$.

The equilibrium policies $V_t(p)$ and $H_t(p)$ solve the NFOCs for the maximization of (10):

$$\begin{aligned} [W] : \quad H_t(p) &= [\mu_t(p) - V_t(p)] (1 - \delta_t) s\lambda_t L_{t-1} f_t(V_t(p)) \\ [h] : \quad c'(H_t(p)) &= \mu_t(p) - V_t(p). \end{aligned} \quad (15)$$

Next, noticing that (11) implies

$$f_t(V_t(p)) = \frac{dF_t(V_t(p))}{dV_t(p)} = \frac{dF_t(V_t(p))}{dp} \frac{1}{V_t'(p)} = \frac{a_t(p)}{A_t} \frac{\gamma(p)}{V_t'(p)},$$

we can rewrite the FOC for W as

$$H_t(p) = [\mu_t(p) - V_t(p)] (1 - \delta_t) s\lambda_t L_{t-1}(p) \frac{a_t(p)}{A_t} \frac{\gamma(p)}{V_t'(p)} \quad (16)$$

Together with the law of motion of employment, equations (14), (16) and (15) are the backbone of an algorithm to describe the evolution of the equilibrium of this economy subject to aggregate shocks. We provide the details of this algorithm in Appendix C

3 Steady state and comparative statics

Before moving on to the quantitative results, in order to gain some intuition on the complex effects of shocks to aggregate productivity ω on average employment and wage, and

distributions thereof, we compare steady-state equilibria with different values of ω . More precisely, we assume that ω_t is constant for ever at ω , the economy is in steady state, so any transitional dynamics of the employment distribution N_t already played out, we describe equilibrium, which extends BM to endogenous contact rates, and study the effects of changes in ω on the steady state N and wages.

In our previous theoretical analysis, in equilibrium a firm of productivity $\omega_t p$ may decide to post a value $V_t(p)$ below the value of unemployment U_t . Because any equilibrium must be RP, i.e. $V_t(\cdot)$ must be increasing, if this happens then generically it will be true for all firms with productivity type below some threshold \underline{p}_t , and a measure $\Gamma_0(\underline{p}_t)$ of firms will be temporarily inactive. So the model allows for entry and exit “at the bottom”. As ω_t evolves stochastically, so will the entry threshold \underline{p}_t . To understand comparative statics, we need to study this effect carefully.

So let

$$\underline{p}(\omega) = \max \left\langle \underline{p}, \frac{\phi}{\omega}, \frac{w_{\min}}{\omega} \right\rangle \quad (17)$$

denote the lowest-productivity firm active in the market, where \underline{p} is the lowest existing (potentially active) firm type, $\phi = \phi(\omega)$ is the reservation wage, the minimum wage acceptable by unemployed workers, to be found later, and w_{\min} is a mandated minimum wage. Let Γ_0 be the population distribution of firm types, and $\Gamma(p) = \Gamma_0(p) / \bar{\Gamma}_0(\underline{p}(\omega))$ its normalized cdf. Let $L(w, a)$ denote the steady state size when posting a constant wage w and a adverts.

The firm maximizes steady-state flow profits, which equal flow output minus wage per worker, times firm size, minus the cost of hiring an inflow given by posted adverts a times the stationary contact probability of each advert η times the acceptance probability $P(w)$.⁸ Denote the optimal level of profit attained by a type- p firm by:

$$\pi(p) = \max_{\substack{a \geq 0 \\ w \geq \max(\phi, w_{\min})}} L(w) (\omega p - w) - c(a\eta P(w)). \quad (18)$$

This program gives rise to an optimal policy $(a(p), w(p))$. Proceeding as in Bontemps et al. (2000), in Appendix B we show that the steady state equilibrium firm size with endogenous sampling weights equals

$$L(p) = \left(1 + s\lambda \frac{1 - \delta}{\delta} \right) \frac{a(p)}{A} \left[1 + s\lambda \frac{1 - \delta}{\delta} \int_p^{\bar{p}} \frac{a(y)}{A} d\Gamma(y) \right]^{-2}$$

and the wage policy $w = w(p)$ is the implicit function defined by

$$w = \omega \left[p - \int_{\underline{p}}^p \frac{L(x)}{L(p)} dx \right] - \Psi(w, p, s\lambda) \quad (19)$$

⁸Maximization of steady-state profit flows coincides with full dynamic maximization of the PDV of future profits when firms are infinitely patient. If not, the two problems yield solutions that produce the same steady-state allocation. See MPV13 for a detailed comparison.

where

$$\Psi(w, p, s\lambda) = \frac{\underline{\pi} + c(\eta a(p) P(w))}{L(p)},$$

and $\underline{\pi}$ are the total profits earned by the marginal firm in the market. Such profits are zero unless neither the reservation wage ϕ nor the statutory minimum wage are binding, in which case the least productive firm is active: $\underline{p}(\omega) = \underline{p}$. Since Ψ is increasing in w , Equation (19) uniquely defines $w(p)$. Intuitively, the firm's mark-up comprises two terms: the integral in (19) captures market power due to frictions, as in the BM model and in Bontemps et al. (2000), while Ψ compensates the firm for hiring costs (which are new to this wage-posting model) and for the minimum rent it can earn in the market, $\underline{\pi} \geq 0$, per worker.

Furthermore, the reservation wage ϕ solves:

$$\phi = b + \beta \int_{\phi}^{\bar{w}} \frac{[1 - (1 - \delta)s] \frac{\lambda}{\delta} \int_x^{\bar{p}} \frac{a(y)}{A} d\Gamma(y)}{1 - \beta(1 - \delta) \left(1 - s\lambda \int_x^{\bar{p}} \frac{a(y)}{A} d\Gamma(y)\right)} dx. \quad (20)$$

The steady-state equilibrium wage function thus solves (19) with ϕ defined in (20) and $\underline{p}(\omega)$ in (17).

We are now ready to take a total derivative of the wage function (19) w.r. to ω to study comparative statics. Note that the value of leisure b , the arrival rate of separation shocks δ , of job offers λ and of job applications η , as well as the advertising policy $a(\cdot)$ and their aggregate A , hence the acceptance probability $P(\cdot)$, all depend on aggregate productivity ω . All of them, except b , influence the wage policy for a given productivity and reservation wage. We can then decompose the response as follows:

$$\frac{dw(p)}{d\omega} = \underbrace{\frac{\partial w(p)}{\partial \omega}}_{\text{Opportunity Cost}} + \underbrace{\omega \frac{L(\underline{p})}{L(p)} \frac{d\underline{p}(\omega)}{d\omega}}_{\text{Entry/Exit}} + \underbrace{\sum_{x \in \{\lambda, \delta, \eta, a, A, b\}} \frac{\partial w(p)}{\partial x} \frac{dx(\omega)}{d\omega}}_{\text{Competition}} \quad (21)$$

where the minimum productivity responds as follows:

$$\frac{d\underline{p}(\omega)}{d\omega} = -\frac{w_{\min}}{\omega^2} \mathbb{I}\{w_{\min} > \phi, \omega \underline{p}\} + \frac{\phi}{\omega^2} \left(\frac{d \log \phi}{d \log \omega} - 1 \right) \mathbb{I}\{\phi > w_{\min}, \omega \underline{p}\} \quad (22)$$

We study the three pieces of (21) in turn.

Opportunity cost. The first effect is the direct impact of the productivity scale: some algebra yields

$$\frac{\partial w(p)}{\partial \omega} \frac{\omega}{w(p)} = \frac{1 + \frac{\Psi(w(p), p, s\lambda)}{w}}{1 + \frac{\Psi(w(p), p, s\lambda)}{w} \frac{\partial \log \Psi(w(p), p, s\lambda)}{\partial \log w}}$$

This reflects the higher opportunity cost (due to a loss of output) of not hiring/retaining workers that a firm pays when ω increases. Suppose all firms keep their wage offers fixed

as ω rises to $\omega' > \omega$. Then firm p will go from productivity ωp to $\omega' p$. Given the strategy of other firms, value of leisure, and arrival rates, this firm will optimize by mimicking firm $p' = p\omega'/\omega$: whether the firm is more productive for idiosyncratic or aggregate reasons is immaterial to its choice, given a wage offer distribution and arrival rates. So firm p will raise its wage offer to $w(p') > w(p)$, and, in equilibrium, all firms will raise their wages.

The elasticity of the wage function with respect to aggregate productivity ω is less than one, so wages are “rigid”, if and only if the elasticity of the mark-up term Ψ with respect to the wage, evaluated at equilibrium wage $w(p)$, is larger than one. In the important and simple special case analyzed by BM, hiring is costless, contact rates are exogenous, and either the minimum wage or the reservation wage are binding, so that a positive measure $\Gamma_0(\underline{p}(\omega))$ of firms is inactive and profits of the marginal firm are zero: $\underline{\pi} = 0$. In this case, the mark-up term Ψ is zero, and all wages rise proportionally to aggregate productivity. If $\underline{\pi} = 0$ but hiring is costly, the mark-up term Ψ equals hiring costs per worker, so the wage $w(p)$ is rigid if the wage elasticity of total hiring costs exceeds one. Since firms pay hiring costs only if they hire ($c(0) = 0, c' > 0$) and the marginal hiring cost is increasing (which is implied by the condition for RP in Proposition 1), the elasticity of $c(\cdot)$ exceeds one. Therefore, the equilibrium wage $w(p)$ is rigid if the wage elasticity of the acceptance probability $P(w)$ exceeds one at $w = w(p)$. Intuitively, if an increase in ω leads firm to offer a higher wage, and this offer poaches many workers at the margin, the firm needs to pay for these additional hires. So the rising marginal hiring cost moderates the incentives to raise the wage in the first place. This is a novel effect that arises from endogenizing hiring choices in a job ladder model.

Entry and exit. The second effect is the impact of aggregate productivity on the set of active firms, i.e. the entry/exit effect. This effect is zero if neither minimum wage nor reservation wage bind, and all firms make positive profits. Otherwise, an increase in ω will pull some relatively unproductive firms into the market. The resulting negative effect on wages depends on $L(\underline{p})/L(p)$, which is exactly equal to one at the bottom and decreasing in p . So the impact of aggregate productivity on wages through the entry/exit margin dissipates as we move up and away from that margin, and top wages are fairly insulated.

Competition. The third effect is indirect impact of aggregate productivity on wage competition through the pace of separations $\delta(\omega)$ and the general equilibrium response of the arrival of outside offers to employed and unemployed workers λ and of job applications η , as well as the advertising policy $a(\cdot)$ and their aggregate A and the acceptance probability $P(\cdot)$, given job productivities and set of active firms. Here we can make two remarks on the

effects of λ and δ on the slope $w'(p)$ of the equilibrium wage function, given the reservation wage ϕ .

First, λ , which depends in equilibrium on job market tightness, only enters the wage function as $s\lambda$, i.e. this effect is

$$\frac{\partial w(p)}{\partial \lambda} = s \frac{\partial w(p)}{\partial (s\lambda)}.$$

This implies that a change in the arrival rate λ of offers to the unemployed has no direct effect on equilibrium wages, and will work only through the reservation wage ϕ . What matters for wage competition above ϕ is the arrival rate of offers to employed workers, which is the true index of competition. This is a sense in which firm's commitment insulates wages from a direct influence of the value of unemployment, as in (but for very different reasons than) Hall and Milgrom (2008)'s credible bargaining.

Second, the destruction rate δ has the opposite effect to λ on the slope of the wage policy, plus an extra negative effect through the term Ψ . Given $\lambda(1-\delta)/\delta$, an increase in δ raises the minimum profits that the firm must make to match the minimum rent in the market and to compensate the upfront hiring cost, hence reduces wages across the board. In BM's steady state analysis, as mentioned $\Psi = 0$, and $\lambda(1-\delta)/\delta$ is a summary index of competition. If trading frictions are extreme $s\lambda(1-\delta)/\delta = 0$, we can easily verify from the equations that the unique solution is $w(p) = \phi = b$ for all p , the Diamond (1971) paradox. Since $\lambda(1-\delta)/\delta$ is increasing in aggregate productivity ω , so is the slope of the wage function and, given ϕ , equilibrium wages, purely due to this competition effect.

4 Data

4.1 Labor market and productivity

Most of the aggregate time series that we exploit to either calibrate or test the model are available at monthly frequency. Nonetheless, we sample and filter them at quarterly frequency, because TFP and Average Labor Productivity (ALP) series are available only quarterly. We will calibrate and simulate the model at monthly frequency, then aggregate its output to quarterly frequency for consistency with the data.

We take logs of and HP-filter all quarterly time series. Labor market stocks and flows are filtered with parameter 100,000, while productivity series with parameter 1,600. The unemployment rate is the civilian unemployment rate of the US population aged 16 and up from the BLS, starting in 1948:Q1 and currently available through 2014:Q3. We calculate transition rates between labor market states. Employment and Unemployment using stocks of E and U and flows EU and UE between them from the BLS. Flows are from monthly CPS

	U	UE	EU	V/U	TFP	ALP
unemployment rate	.204					
UE rate	−.973	.124				
EU rate	.887	−.878	.100			
V/U ratio	−.963	.962	−.921	.357		
Total Factor Productivity (TFP)	.082	.154	.054	−.188	.009	
Average Labor Productivity (ALP)	−.210	−.063	−.239	−.010	.197	.013

Table 1: Correlations and standard deviations of quarterly cyclical components

matched files (BLS “Research series on labor force status flows from the CPS”) beginning in 1990. We do not use unemployment-duration based measures, nor do we correct for time aggregation (Shimer, 2012) and we treat any short spell of unemployment that completes within a month, and is missed by monthly CPS interviews, as a continuous employment spell. We make this choice for two reasons: the model is calibrated monthly, so workers there have no chance to make more than one transition per month, and in the data a majority of very short unemployment spells end in recall by the same employer (Fujita and Moscarini, 2013).

The V/U ratio is the ratio between the JOLTS vacancy rate (currently available for 2010:Q4-2014:Q1) for the US private nonfarm sector and the civilian unemployment rate.⁹

For TFP we use quarterly estimates of the Solow residual corrected for capacity utilization by Fernald (2012), which are updated by the author and currently cover the period 1947:Q1-2014:Q2. For ALP, we take Output per job in the Nonfarm business sector from the BLS (series PRS85006163), currently available for 1947:Q1-2014:Q3.

We report in Table 1 standard deviations (last number on each row) and correlation coefficients of the filtered components of each series. Because the series are available for different time spans, the second moments in the table refer to different time periods. We prefer to take this approach rather than to shorten our sample to the longest common period when they are all available. We remark that TFP and ALP have similar volatilities, much smaller than that of (log) unemployment, but TFP has a stronger negative correlation than ALP with the unemployment rate. This reflects in part the fact that our filtered ALP measure turned slightly countercyclical since the 1980s.

⁹This ratio is routinely referred to as labor market tightness in the context of the DMP model. However, the relevant concept of market tightness in our model, which features on-the-job search, is not the V/U ratio: rather, it is the ratio of vacancies to total worker search effort, $U + s(1 - \delta)(1 - U)$. Because the V/U ratio has been the focus of much attention in the literature, we report statistics pertaining to that variable, even though it doesn’t have a straightforward interpretation within the context of our model.

log earnings	monthly		quarterly	
	(1)	(2)	(3)	(4)
unemployment rate	-.60 (.11)	-.50 (.11)	-.58 (.11)	-.49 (.11)
25-99 employees		.065 (.004)		.062 (.004)
100+ employees		.078 (.005)		.075 (.005)
observations (i, t)	2,910,524		986,149	
# workers (i)	82,036		81,457	
$\text{Var}_{it}(\text{errors})$.18	.18	.16	.16
$E_t[\text{Var}_i(\text{errors})]$.18	.18	.16	.16
R^2	.22	.23	.23	.24

Notes: Source: SIPP. Monthly unemployment rate is HP-filtered with parameter 10E5. Quarterly log earnings and unemployment rate are averages of monthly series. $E_t[\text{Var}_i(\text{errors})]$ is time average of the cross-sectional variance of errors in each period. Robust standard errors in parentheses. All specifications include quartic time trend and demographic controls.

Table 2: Fixed-effect regression of individual real log earnings

4.2 Wages

We use the 1996, 2000, 2004 and 2008 panels of the Survey of Income and Program Participation (SIPP). This is a collection of panels of workers, on average about 40,000 workers per panel. Each worker is surveyed every four months about events that occurred over the past four months (‘wave’). Workers are divided in rotation groups, so each calendar month roughly one quarter of the panel completes a wave and is re-interviewed. The SIPP has significant longitudinal and time dimensions, provides information about a large number of individual outcomes, and is nationally representative.

As the model has no intensive margin of labor supply, for our measure of wages we take monthly earnings (TPMSUM), which are also less contaminated by measurement error than hourly pay rates, especially those constructed as the ratio of earnings to usual hours worked per month. We deflate with monthly CPI and take logs. A “seam bias” arises from a tendency of SIPP respondents to bunch reported changes in earnings and labor market status around interview times. Rotation groups are staggered, however, so seams are evenly distributed over time. The only exceptions are the first four and last three months of each panel, when no rotation group has a seam, but excluding those months has little impact on the results.

We aim to extract information about the cyclical component of earnings for comparable workers, as our model has homogeneous workers. To this purpose, we run, with both

monthly and quarterly data, a fixed effect regression of individual log earnings, constructed as described above, on the detrended unemployment rate, a common linear trend, a cubic in age, and dummies for the four available worker characteristics that may change over time: class of worker (private, government, no profit etc.), education (17 categories), marital status and union coverage. In columns 2 and 4 of Table 2 we also control for a coarse measure of employed size available in the SIPP data.¹⁰ Race, gender and any other time-invariant worker characteristics are absorbed by the individual fixed effect. The regression uses the SIPP longitudinal weights to maintain the representative character of the sample.

The results in Table 2 indicate a negative semi-elasticity, significant at conventional levels. Looking at column 2 and 4, we see that employer size comes out positive and significant, as predicted by the theory. The addition of employer size in the set of regressors reduces (the point estimate of) the unemployment semi-elasticity of wages a bit, likely owing to a composition effect, also consistent with the theory: periods of low (high) unemployment are periods in which the workforce is more concentrated among higher-paying and larger (lower-paying and smaller) employers (we return to this point in the next section). The first row of Table 2, specifically column 3, provides our empirical target for the cyclicity of wages produced by the model.

5 Calibration and simulation

5.1 Calibration

We specify the aggregate shock process as a discrete 20-state first-order Markov chain that approximates a monthly AR(1) process for $\ln \omega$. The AR coefficient (0.94) and the variance of innovations (0.006) are set such that the model replicates the observed variance and first-order autocorrelation of HP-detrended output per worker.¹¹

The job destruction rate is allowed to fluctuate deterministically with ω as follows:

$$\delta(\omega) = 0.0114 + 1.894 \times (\ln \omega_{\max} - \ln \omega)^{2.5}$$

where the intercept, slope and exponents in this definition are chosen to approximate the mean, standard deviation, (positive) skewness and kurtosis of the observed EU rate.

All other parameter values and functional form assumptions are summarized in Table 3. The discount factor is the monthly equivalent of a 5% discount rate per annum. The

¹⁰This measure is an indicator of whether the worker's current firm employs up to 24, between 25 and 99, or more than 100 people at the time of interview.

¹¹Output per worker does not exactly coincide with ω in the model because of the gradual selection of workers up the job ladder, but it turns out to be very close.

monthly discount factor:	$\beta = 0.95^{1/12}$
employed search intensity:	$s = 0.13$
matching function:	$m(A, Z) = 0.059 \times A$
hiring cost:	$c(h) = (43.47 \times h)^{100} / 100$
unemployment flow income:	$b = 0$
firm type distribution:	$\Gamma(p) = \frac{1-p^{-2.6}}{1-\bar{p}^{-2.6}}$, with $[p, \bar{p}] = [1, 100]$

Table 3: Calibration

employed search intensity s is set to match an average monthly EE transition rate of 1.2%, in line with SIPP data.¹² The matching function is simply a linear function of the aggregate number of job adverts. This differs from the more commonly used Cobb-Douglas specification with a vacancy elasticity of 0.5, and as such calls for some comments. First, there is some evidence that the estimated vacancy elasticity of the matching function is substantially higher than 0.5 when employed job seekers are counted as inputs into the matching process (Petrongolo and Pissarides, 2001).¹³ Second, taking simulated series from our model, and running an OLS regression of the log of the job finding rate on the log of the V/U ratio (the standard procedure in the empirical literature on matching functions that produced this “standard value” of 0.5 for the matching function elasticity), we obtain a “matching function elasticity” of 0.47, well within the range of standard estimates. Finally, results are not very sensitive to variations of the matching function elasticity within the range $[0.7, 1]$.

The hiring cost function is iso-elastic, and highly convex. The need for a highly convex hiring cost function is suggested by the sufficient condition stated in Proposition 1. In the simulation, more convex hiring costs tend to (moderately) increase the volatility of the job finding rate, for reasons discussed in Section 5.3.2. The scaling factor in the hiring cost function is set so that the aggregate measure of job adverts, A , equals one on average over the ergodic distribution of ω — a mere normalization.

The unemployment flow income is set equal to 0. Low estimates of b are common in empirical applications of BM-type models, for the reasons discussed in Hornstein et al. (2011), and are sometimes justified by appealing to a disutility, or “stigma”, attached to the

¹²There are many different ways to define EE transitions in the SIPP (or any other) data. Here we allow for up to seven days of non-employment between jobs, and up to a month for “voluntary” EE transitions, i.e. when respondents stated that they quit their past job to take another one. Adding “involuntary” EE transitions (in the form of forced job-to-job reallocations, as for instance in Jolivet et al., 2006) would be a straightforward extension.

¹³Even when on-the-job search is ignored, recent estimates by Borowczyk-Martins et al. (2013) that account for the endogeneity of vacancies in the presence of reallocation shocks suggest that the matching function elasticity may be higher than commonly assumed, in the region of 0.8.

state of unemployment which has to be subtracted from what is literally income in the flow value of unemployment. However the main reason why we adopt this low value is tractability: a higher value of b would induce some low-productivity firms to shut down in bad aggregate states. Working the possibility of firm exit into the simulation is not an easy task, and we leave it for future research.¹⁴

Finally, the underlying distribution of firm types is a truncated Pareto distribution, the most productive firm being 100 times more productive than the worst one in the economy.

5.2 Main quantitative results and model fit

We simulate the model over a window of 70 years (a length comparable to the times series described in Section 4) using the algorithm described in Appendix C. We now discuss various aspects of the fit and some predictions of the model.

5.2.1 Labor market variables and ALP

Table 4 replicates the data moments displayed in Table 1, together with their model counterparts.¹⁵

Although the model tends to understate slightly the volatility of unemployment and the V/U ratio, it still captures the fact that those labor market indicators are an order of magnitude more volatile than ALP. This demonstrates the model’s potential to generate a plausible level of amplification of TFP shocks. The model also replicates the correlations between the various labor market indicators well enough, despite a tendency to overstate the correlations between the V/U ratio and the job finding and unemployment rates. Overall, though, the model comes close enough to replicating the moments relating to labor market variables, including the volatility of the job finding rate, which has been the focus of attention in much of the recent literature.

Where the model does the least well is in predicting the correlations between ALP and the various labor market indicators: the magnitudes of all correlations are grossly overstated, and the model even gets the sign wrong for the correlations of ALP with the job finding rate and the V/U ratio. This, however, is hardly surprising in a model where the only impulse driving aggregate fluctuations is a TFP shock. Allowing for additional uncorrelated shocks is like to help with increasing the predicted volatility of the job finding rate, decoupling labor

¹⁴We further argue in MPV13 that ignoring entry and exit of firms is not such a big oversight when analyzing the cyclical volatility of employment, as the contribution of entry and exit to the cyclical volatility of employment is modest.

¹⁵The quarterly autocorrelation of detrended ALP, not shown in the tables, is 0.77, both in the data and in the simulation.

		U	UE	EU	V/U	ALP
unempl. rate	(M)	.196				
	(D)	.204				
UE rate	(M)	-.991	.126			
	(D)	-.973	.124			
EU rate	(M)	.828	-.751	.111		
	(D)	.887	-.878	.100		
V/U ratio	(M)	-.999	.996	-.806	.270	
	(D)	-.963	.962	-.921	.357	
ALP	(M)	-.726	.659	-.865	.709	.013
	(D)	-.210	-.063	-.239	-.01	.013

Table 4: Fit to correlations and standard deviations of quarterly cyclical components

market indicators from ALP.¹⁶

5.2.2 Wages

We construct a model-predicted semi-elasticity of wages to the unemployment rate by running a weighted OLS regression of firm-level log wages on a constant and on the detrended unemployment rate, with weights equal to $L_t(p)\gamma(p)$. This is the closest we can get, using the model, to the worker-level regressions reported in Table 2. The predicted semi-elasticity thus computed (at quarterly frequency) is -6.35 . Comparison of this number with the value of -0.6 reported in the last column of Table 2 suggests that wages are much more procyclical in the model than in the data (where they are fairly sticky). The predicted semi-elasticity drops a little to -5.72 when firm size is included among the covariates in the aforementioned regression. While this drop parallels the drop seen between columns 3 and 4 of Table 2, the predicted semi-elasticity is still far too large.

As another point of comparison, the quarterly unit labor cost series produced by the BLS over the period 1947-2013 has a cyclical volatility of 0.016 (after taking logs and HP-filtering with parameter 1,600). The model counterpart of this number is 0.062, further corroborating the conclusion that predicted mean wages are too cyclical. Excess wage volatility is a feature that this model shares with Nash bargaining models, although, as discussed in Section 3, it arises from very different causes in the two models. And contrary to the Nash bargaining model, the contract posting model does generate plausible amplification of TFP shocks *despite* flexible wages. Perhaps more informative than wages in this regard, posted *worker*

¹⁶Shocks to the matching technology or to the job destruction rate come to mind as natural candidates. A recent paper by Hall (2014) makes a case for a countercyclical discount factor.

values are substantially less volatile than wages, their standard deviation being 0.038 at quarterly frequency, and their semi-elasticity to unemployment being -1.82 . This partly explains the degree of amplification generated by the model: although wages in the model are allocative, they merely serve to implement values that drive the allocation of employment in the model.

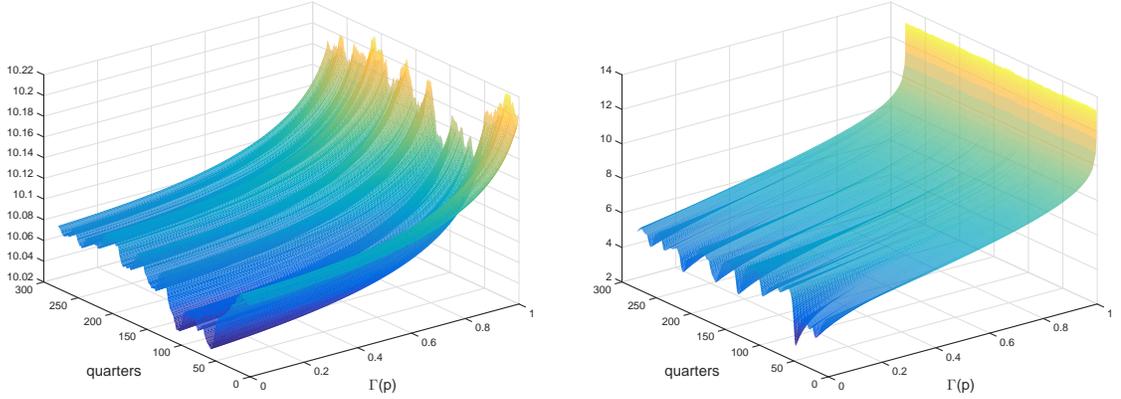
As discussed in the Introduction, one of the key contributions of our SBM model is that it generates endogenous cross-sectional wage dispersion and as such makes predictions on second- and higher-order moments of the wage distribution. Focusing on cross-sectional wage dispersion, the pooled cross-section wage variance predicted by the model is 0.03. This falls far short of the residual variance of 0.18 found in the SIPP data and reported in Table 2. While some of that 0.18 variance is arguably attributable to measurement error, it can safely be concluded that the model understates wage dispersion.¹⁷ The failure to replicate the observed amount of residual wage dispersion — at least under “standard” parameterizations — is a well-known and recurring problem of wage posting models (see Hornstein et al., 2011, and the discussion in the Introduction of this paper), from which the dynamic contract posting model of this paper is not immune. Just as in the class of steady-state models discussed by Hornstein et al. (2011), this model could be made to predict levels cross-section wage dispersion more in line with the data under some extreme parameterizations involving a large negative flow value of unemployment, or a very high discount rate.¹⁸

Without entering the debate about whether such unconventional parameterizations make sense or not, we take note of the fact that our model under-predicts the level of wage dispersion, and move on to its cyclical behavior. A number of recent papers have documented the fact that idiosyncratic wage risk, broadly defined as the cross-sectional standard deviation of wage growth, is countercyclical (Storesletten, Tamer and Yaron, 2004). The predicted correlation between the cross-section standard deviation of log wage growth and the (detrended) unemployment rate is 0.21.¹⁹ The model thus replicates, at least qualitatively, the counter-cyclicality of wage risk found in the data.

¹⁷Attributing the entire discrepancy between the model-based variance of 0.032 and the data-based residual variance of 0.18 to classical measurement error would imply a measurement error variance of $0.18 - 0.03 = 0.15$, roughly 83% of the total residual wage variance. This seems high: Lemieux (2006) reports a corresponding share of around a third based on CPS data.

¹⁸For example, assuming $b = -93$ (which amounts to about -100% of the mean wage) and leaving the rest of the calibration unchanged, the model predicts a cross-section wage variance of 0.105, which is much closer to the data, at least if one allows for a reasonable amount of measurement error. Its predictions about amplification and volatility are virtually unaffected by b .

¹⁹The model-predicted standard deviation of wage growth is constructed by taking the standard deviation of firm-level wages weighted by initial employment at each firm. Thus, by design, it measures the wage risk faced by job stayers.



(a) Worker values

(b) Profits

Fig. 1: Marginal profits and worker values over time

5.3 Additional quantitative results

5.3.1 Profits and worker values

Figure 1a plots the profile of (log) values posted by each firm, $\ln V_t(p)$, as a function of the quantiles of p in the distribution Γ , over time. Figure 1b does the same for the (log) profitability of the marginal job in any given firm type, $\ln [\mu_t(p) - V_t(p)]$. As per the properties of RPE, $V_t(p)$ is increasing in p at all dates. Figure 1b suggests that this is also the case for marginal profitability, although this result is not guaranteed in theory.

Perhaps the most striking aspect of Figure 1 is the relative range of $\ln V_t(p)$ vs. $\ln [\mu_t(p) - V_t(p)]$. Figure 1b indicates that the marginal job at the most productive firm is over 1,000 times more profitable than the marginal job at the least productive firm. By contrast, the highest posted value is a mere 11.9% higher than the lowest posted value (which equals U , the value of unemployment). This is another rendition of the model's failure to generate enough cross-sectional dispersion in wages (or, in this case, job values): firms at the top of the productivity ladder face very little competition, as there are only very few firms that can poach their workers. As a result, they enjoy a very large amount of market power.

A closer look at Figure 1b further reveals that job profitability is much more cyclically sensitive at low-productivity firms than at the top end of the productivity distribution. This is because jobs at low-productivity firms generate smaller surplus than jobs at high-productivity firms, and their value is therefore more elastic to TFP shocks. Interestingly, 1a suggests that this difference in cyclicity across firm types is not passed on to worker values.

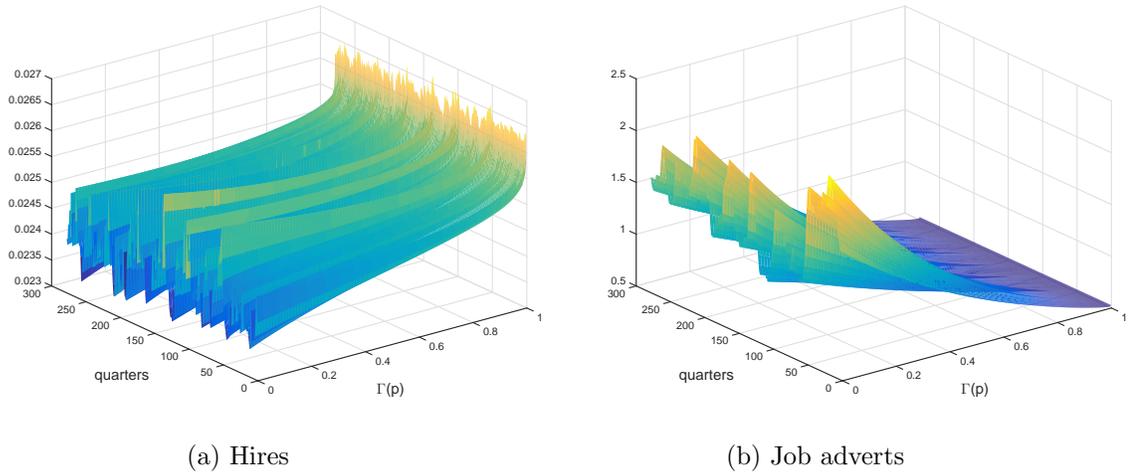


Fig. 2: Firm-level job adverts and hiring flows over time

5.3.2 Job adverts, hiring flows, and firm size

Proposition 1 states that in RPE, a firm that is more productive and initially larger than another stays larger period after period. But it says nothing about whether more productive firms actually *hire* more, or if they merely retain a larger fraction of their initial workforce by posting more attractive job values. Figure 2 provides an answer, under the parameterization discussed in the previous subsection. Figure 2a plots the profile of firm-level hiring flows, $H_t(p)$ (still as functions of the quantiles of p in the distribution Γ) over time, while Figure 2b plots the profile of firm-level job adverts, $a_t(p)$. It appears that $H_t(p)$ is mostly, although not always increasing in p , while $a_t(p)$ is frankly decreasing: more productive firms tend to hire a larger inflow of workers at most dates, but manage to do so by spending *less* effort on hiring, simply through the fact that they offer more attractive job values, thus enjoying a higher job-ad yield. Figure 2 further shows that job adverts, and therefore hires, are much more cyclically sensitive at low-productivity than at high-productivity firms, echoing the similar remark made above about profitability of the marginal job. The firm's hiring decision is indeed governed by the FOC (15), which imposes equality between the cost of the marginal hire ($c'(h)$, an increasing function of h), and its profitability, $\mu(p) - V(p)$.

Another remarkable feature of Figure 2a is the lack of dispersion in the level of gross hiring inflows across firm types. On average over the simulated sample, the most productive firm in the market only hires 8.35% more workers per quarter than the least productive one. This is in spite of a 100-fold difference in productivity and a 1,000-fold difference in profitability between the highest- and lowest-productivity employer in this economy (see previous paragraph). As a further consequence, the model vastly under-predicts dispersion in firm sizes: the predicted size ratio between the smallest and largest employer in the economy

is 3.9 on average, nowhere near the levels observed in the data.

The reason for this lack of dispersion in employer size and hiring inflow is the very high degree of convexity of the hiring cost function $c(\cdot)$. As mentioned Sub-section 5.1, this high degree of convexity is needed for the model to amplify TFP shocks to the extent observed in the data. In that sense, the model carries a tradeoff between amplification and size dispersion, the terms of which are as follows. As seen on Figures 1 and 2, profitability and (consequently) hiring and job-ad posting are much more volatile at low-productivity than at high-productivity firms, as the former generate a much smaller surplus than the latter. In order to inflate the volatility of *aggregate* job adverts, the model needs to concentrate hiring among low-productivity firms. The highly convex training cost does exactly that: it curtails hiring at high- p firms, which are much more profitable than low- p firms but face prohibitive marginal training costs beyond a modest level of hiring.

5.3.3 The Beveridge curve

Figure 3 shows two different versions of the Beveridge curve produced by the model. The blue plot is the “empirical Beveridge curve”, or VU -curve, i.e. a plot of the unemployment rate against aggregate job adverts. While this VU -curve makes no particular sense in a model with on-the-job search, it is interesting for being the object of keen empirical investigation, and the basis for most of the existing inference on matching functions. The red plot on Figure 3 is the Beveridge curve that is consistent with our model, i.e. a plot of aggregate worker search effort, $1 - N_t(\bar{p}) + s(1 - \delta_t) N_t(\bar{p})$, against aggregate job adverts.²⁰

Both versions of the Beveridge curve look as a Beveridge curve should, i.e. are downward sloping and convex, despite a perfectly countercyclical job destruction rate. Interestingly, however, the “empirical” Beveridge curve is slightly flatter than the “consistent” one. In other words, if one were to use any of these plots for prediction, one might conclude that a decrease in labor market tightness by a given amount will cause a larger increase in the mass of job seekers, hence a larger drop in the job finding rate, under the empirical Beveridge curve than under the consistent one, which takes employed job seekers into account. This tallies with the remark made earlier, when discussing calibration: although the true matching function has a unit elasticity in the model, estimating a misspecified job finding rate based on the empirical version of the Beveridge curve — i.e. regressing this rate on the V/U ratio — will understate the sensitivity of the job finding rate to changes in market tightness. The magnitude of the bias suggested by the model (estimated elasticity of 0.46 instead of 1) is quite remarkable. Where does the bias come from? The job finding rate is a function not of

²⁰Aggregate worker search effort has been rescaled to produce Figure 3, in such a way that both series on the x -axis of the figure have equal means.

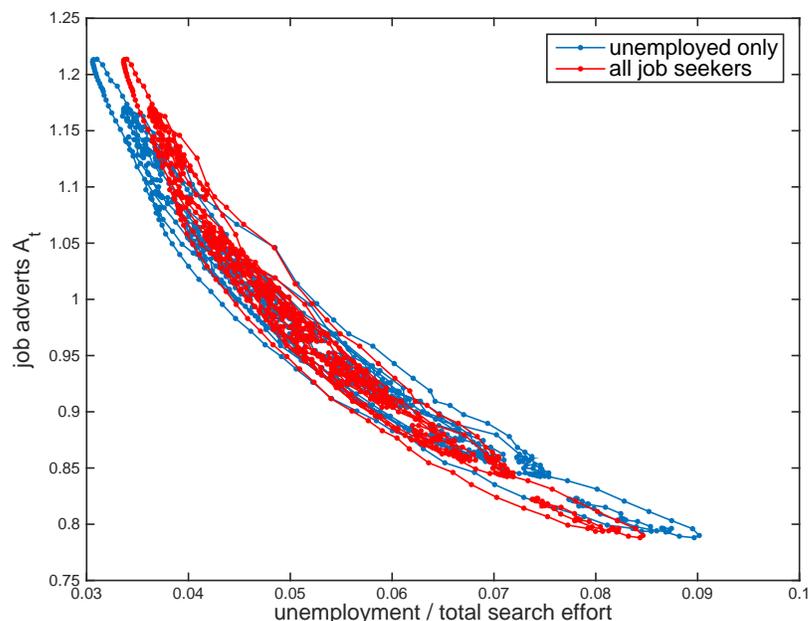


Fig. 3: The Beveridge curve

the V/U ratio, but of labor market tightness which equals the ratio of aggregate job adverts (or vacancies) to aggregate job search effort, $1 - N(\bar{p}) + s(1 - \delta)N(\bar{p})$, which accounts for the competition that unemployed workers face from employed job seekers. Because employed job seekers are fewer in recessions and more in expansions, actual tightness varies relatively less than the V/U ratio over the cycle, which explains the lower estimated elasticity of the job finding rate to the V/U ratio than to tightness.

5.3.4 The cyclicity of net job creation

As pointed out earlier, *gross* job creation is much *more* cyclical among small, low-productivity employers than at the top end of the productivity distribution. Yet in MPV12 we highlight that, in US data, exactly the opposite is true of *net* job creation: as a group, small employers fare relatively better in bad times of high unemployment, and worse in good times of low unemployment. As we propose in other work (MPV09, MPV14), this empirical pattern is qualitatively consistent with job ladder models, of which the model in this paper is an example: in a tight labor market, high-paying, large employers overcome the scarcity of unemployed job applicants by poaching employees from smaller, less productive and lower-paying competitors, whose employment share then shrinks in relative terms. When the expansion ends, large employers, that were less constrained, have more employment to shed than small ones. In addition, the resulting high unemployment relaxes hiring constraints on all employers, particularly the small ones that are less capable of poaching from other firms.

As a result, small employers downsize less in the recession and grow (relatively) faster in the recovery.

This mechanism is at work here. To put a number on its magnitude, we construct series of total net job creation among two groups of firms, one at the top and one at the bottom of the productivity ladder, both classes being re-designed at each date so that they each have an employment share of 25%. We then subtract net job creation in the group of low- p firms from net job creation in the group of high- p firms, and HP-detrend this difference to obtain a measure of relative net job creation similar to the series we analyzed in MPV13. The correlation of this simulated series of relative net job creation with the simulated unemployment rate is -0.87 : as expected, the relative net job creation of large vs. small employers is procyclical. How plausible the magnitude of that predicted correlation is hard to assess, for lack of immediately comparable data. In MPV13 we report numbers for U.S. data that vary between -0.25 and -0.6 , depending on the data set, period, sector, etc. Again, the correlation produced by the model is on the high side of that range, something that is to be expected as the outcome of a one-shock model.

6 Concluding remarks

The wage-posting model of the labor market with random search is one of the most influential and enduring contributions of Dale Mortensen and his students. The main goal of this paper is to show that this canonical model of the labor market naturally belongs in the toolbox of business cycle theorists. We study its quantitative performance and confront it with empirical evidence from the US. This is our first step in this direction and, as such, is truly exploratory in nature.

Our exercise reveals dimensions on which the model performs well, and also a fundamental tension between its predictions for cross-sectional inequality in wages and firm size, and aggregate volatility. We emphasize that this tension emerges precisely because the model is unusually equipped to address both sides of this trade-off, inequality and cyclical volatility. In this regard, the model generates business cycles of the correct magnitude by exploiting firm heterogeneity, in particular, the concentration of gross hires among unproductive, small, cyclically sensitive firms. This is more than offset by a strong cyclical response of job to job quits, so that net job creation is more volatile among large firms, a prediction that in previous work we argued to be shared by many job ladder models, and that we validated empirically. We also find promise in the model's ability to generate cyclical comovement of the empirically correct sign between unemployment and wage levels, wage inequality, worker turnover. Much work lies ahead to identify suitable specifications and parameterizations of

hiring costs, state-contingent contracts, search effort that may further improve the model's quantitative performance along these dimensions.

We believe that additional insights from this stochastic job ladder model lie ahead. An especially promising direction is understanding individual earnings dynamics, the focus of the incomplete market literature, their relationship to labor market transitions, up and down the job ladder, and their cyclicalities.

References

- [1] Bilal, Mark, Y. Chang and S. Kim, (2012), "Comparative advantage and unemployment", *Journal of Monetary Economics*, 59(2), 150-165.
- [2] Burdett, K. and D. T. Mortensen (1998) "Wage Differentials, Employer Size and Unemployment", *International Economic Review*, 39, 257-73. [BM]
- [3] Bontemps, C., J.-M. Robin and G. J. Van den Berg (2000) "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation", *International Economic Review*, 41(2), 305-358.
- [4] Borowczyk-Martins, D., G. Jolivet and F. Postel-Vinay (2013) "Accounting for Endogeneity in Matching Function Estimation", *Review of Economic Dynamics*, 16, 440-51.
- [5] Christiano, L., M. Eichenbaum and M. Trabandt (2013) "Unemployment and Business Cycles", Northwestern University mimeo.
- [6] Coase, R. (1972) "Durability and Monopoly", *Journal of Law and Economics*, 15(1), 143-49.
- [7] Coles, M. and D. T. Mortensen (2013) "Equilibrium Wage and Employment Dynamics in a Model of Wage Posting without Commitment" NBER WP 17284.
- [8] Diamond, P. A. (1971) "A Model of Price Adjustment", *Journal of Economic Theory*, 3, 156-168.
- [9] Diamond, P. A. (1982) "Aggregate Demand Management in Search Equilibrium", *Journal of Political Economy*, 90(5), 881-894.
- [10] Fernald, J. (2012) "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity", Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- [11] Fujita, S. and G. Moscarini (2013) "Recall and Unemployment", NBER WP No. 19640.

- [12] Hagedorn, M. and I. Manovskii (2008) “The Cyclical Behavior Of Equilibrium Unemployment and Vacancies Revisited” *American Economic Review*, September, 98(4), 1692-1706.
- [13] Hall, R. E. (2014) “High Discounts and High Unemployment”, NBER WP 19871.
- [14] Hall, R. E. and P. Milgrom (2008) “The Limited Influence of Unemployment on the Wage Bargain”, *American Economic Review*, September, 98(4), 1653-74.
- [15] Hornstein, A., P. Krusell and G. L. Violante (2011) “Frictional Wage Dispersion in Search Models: A Quantitative Assessment”, *American Economic Review*, December, 101(7), 2873-98.
- [16] Jolivet, G., F. Postel-Vinay and J.-M. Robin (2006), “The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US”, *European Economic Review*, 50(4), 877-907.
- [17] Kahn, L. and E. McEntarfer (2014), “Employment Cyclicalities and Firm Quality”, NBER Working Paper 20698.
- [18] Lemieux, T. (2006) “Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?”, *American Economic Review*, June, 96(3), 461-498.
- [19] Menzio, G. and S. Shi (2011) “Efficient Search on the Job and the Business Cycle”, *Journal of Political Economy*, 119(3), 468-510.
- [20] Mortensen, D. T. (2003) *Wage Dispersion: Why Are Similar Workers Paid Differently?*, MIT Press.
- [21] Mortensen, D. T. and C. A. Pissarides (1994) “Job Creation and Job Destruction in the Theory of Unemployment”, *Review of Economic Studies*, 61, 397-415.
- [22] Moscarini, G. and F. Postel-Vinay (2009) “The Timing of Labor Market Expansions: New Facts and a New Hypothesis”, in D. Acemoglu, K. Rogoff and M. Woodford, eds., *NBER Macroeconomics Annual*, 23, 1-51. [MPV09]
- [23] Moscarini, G. and F. Postel-Vinay (2012) “The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment”, *American Economic Review*, October, 102(6), 2509-2539. [MPV12]
- [24] Moscarini, G. and F. Postel-Vinay (2013) “Stochastic Search Equilibrium”, *Review of Economic Studies*, 80(4), 1545-1581 [MPV13].

- [25] Moscarini, G. and F. Postel-Vinay (2014) “Did the Job Ladder Fail After the Great Recession?”, CfM discussion paper No.1 [MPV14].
- [26] Petrongolo, B, and C. A. Pissarides (2001) “Looking Into the Black Box: A Survey of the Matching Function”, *Journal of Economic Literature*, 39, 390-431.
- [27] Pissarides, C. A. (1985) “Short-run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages”, *American Economic Review*, September, 75(4), 676-690.
- [28] Pissarides, C. A (2009) “The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?”, *Econometrica*, 77(5), 1339-1369.
- [29] Postel-Vinay, F. and J.-M. Robin (2002) “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity”, *Econometrica*, 70(6), 2295-350.
- [30] Robin, J.-M. (2011) “On the Dynamics of Unemployment and Wage Distributions”, *Econometrica*, 79(5), 1327-1355.
- [31] Shimer, R. (2005) “The Cyclical Behavior of Unemployment and Vacancies”, *American Economic Review*, March, 95(1), 25-49.
- [32] Shimer, R. (2012) “Reassessing the Ins and Outs of Unemployment”, *Review of Economic Dynamics*, 15(2), 127-148.
- [33] Spear, S. and S. Srivastava, (1987) “On Repeated Moral Hazard With Discounting”, *Review of Economic Studies*, 54, 599-617.
- [34] Storesletten, K., C. Telmer and A. Yaron (2004), ”Cyclical dynamics in idiosyncratic labor market risk”, *Journal of Political Economy*, 112(3), 695-717.

Appendix

A Proof of Proposition 1

This proof is adapted from the proof of Proposition 2 presented in Appendix A of MPV13, and follows the same general strategy. For clarity, we amend the notation slightly in this proof by restoring the explicit dependence of various variables on the aggregate state of the economy (ω, N) , rather than subsuming it into a time index as we do in the main text.

We begin with the following preliminary remark. The firms' maximization problem (10) is couched in terms of a joint choice of posted worker values W and hires h . This joint choice yields a next-period, state-contingent size:

$$\mathcal{L}(L, W, h, \omega', N) := L \left(1 - \delta^{\omega'}\right) \left(1 - s\lambda^{\omega'} \bar{F}(W | \omega', N)\right) + h,$$

so that the hire flow $h = \mathcal{L} - L(1 - \delta^{\omega'}) (1 - s\lambda^{\omega'} \bar{F}(W | \omega', N))$. Problem (10) can therefore be recast as an equivalent choice of posted value W and future size \mathcal{L} . We shall work with this alternate formulation of the problem in this proof.

We now introduce some more notation. First, we define the retention rate of a firm posting value W in state ω :

$$R(W | \omega, N) := \left(1 - \delta^{\omega'}\right) \left(1 - s\lambda^{\omega'} \bar{F}(W | \omega', N)\right).$$

For any choice of (W, \mathcal{L}) , the firm makes corresponding hires $h = \mathcal{L} - R(W)L$. We further let:

$$\varphi(p, L, \omega, N) := \omega p L + \beta \int_{\Omega} \delta^{\omega'} U(\omega', \mathcal{N}(\omega', N)) L Q(d\omega' | \omega),$$

with:

$$\mathcal{N}(\omega', N) := \int_{\underline{p}}^p \mathcal{L}(x, L(x), \omega', N) d\Gamma(x).$$

and:

$$\begin{aligned} \Phi(L, W(\omega'), \mathcal{L}(\omega'), \omega', N) &:= -c[\mathcal{L}(\omega') - LR(W(\omega') | \omega')] \\ &+ L \int_{W(\omega')}^{+\infty} v dR(v | \omega', N) - W(\omega') [\mathcal{L}(\omega') - LR(W(\omega') | \omega')]. \end{aligned}$$

Fix N to be some given c.d.f. over $[\underline{p}, \bar{p}]$. Then, for any function $\mathcal{S}(p, L, \omega)$, we define the following operator \mathbf{M}^N :

$$\begin{aligned} \mathbf{M}^N \mathcal{S}(p, L, \omega) &:= \varphi(p, L, \omega, N) + \beta \int_{\Omega} \max_{W(\omega'), \mathcal{L}(\omega')} \left\langle \mathcal{S}[p, \mathcal{L}(\omega'), \omega'] \right. \\ &\quad \left. + \Phi(L, W(\omega'), \mathcal{L}(\omega'), \omega', N) \right\rangle Q(d\omega' | \omega). \quad (23) \end{aligned}$$

Importantly, as stated in the proposition, we make the following two assumptions: first, the sampling c.d.f. $F(W)$ is everywhere differentiable. This implies differentiability of $R(W)$. Second, the hiring cost function is such that $\varepsilon(c') := hc''(h)/c'(h) \geq 1$ for all $h \geq 0$.

We then start the proof with the following lemma:

Lemma 1 *Let $\mathcal{S}(p, L, \omega)$ be bounded, continuous in p and L , increasing and convex in L , and with increasing differences in (p, L) over $(\underline{p}, \bar{p}) \times (0, 1)$. Then:*

1. $\mathbf{M}^N \mathcal{S}$ is bounded and continuous in p and L ;
2. There exists a measurable selection $(V(p, L, \omega, N), \mathcal{L}(p, L, \omega, N))$ from the maximizing correspondence associated with $\mathbf{M}^N \mathcal{S}$;
3. Any such measurable selection is such that V and \mathcal{L} are both increasing in p and L ;
4. $\mathbf{M}^N \mathcal{S}$ is increasing and convex in L , and has increasing differences in (p, L) over $(\underline{p}, \bar{p}) \times (0, 1)$.

Proof. In this proof, wherever possible without causing confusion, we will make the dependence of all functions on aggregate state variables ω and N implicit to streamline the notation.

Points 1 and 2 of this proposition follow immediately from the same arguments as in the proof of Lemma 1 in MPV13, Appendix A. We now turn to the proof of point 3. In each future state ω' , the firm solves:

$$\max_{W, \mathcal{L}} \langle \mathcal{S}(p, \mathcal{L}) + \Phi(L, W, \mathcal{L}) \rangle, \quad (24)$$

where the dependence on ω' and N of the various functions involved was kept implicit to lighten notation. The FOC for W is:

$$0 = c' [\mathcal{L} - LR(W)] \times LR'(W) - [\mathcal{L} - LR(W)] \quad (25)$$

and that for \mathcal{L} is:

$$\mathcal{S}_{L,r}(p, \mathcal{L}) - W - c' [\mathcal{L} - LR(W)] \leq 0 \leq \mathcal{S}_{L,\ell}(p, \mathcal{L}) - W - c' [\mathcal{L} - LR(W)] \quad (26)$$

where $\mathcal{S}_{L,r}$ and $\mathcal{S}_{L,\ell}$ denote the right- and left-derivatives of \mathcal{S} w.r.t. L , which exist everywhere by convexity of \mathcal{S} .

Equation (26) implies that, at an optimal choice of \mathcal{L} , $\mathcal{S}_{L,r}(p, \mathcal{L}) \leq \mathcal{S}_{L,\ell}(p, \mathcal{L})$. But convexity of \mathcal{S} implies that the converse inequality also holds, which proves

that \mathcal{S} is differentiable w.r.t. L at any optimal choice of (W, \mathcal{L}) . The FOC for \mathcal{L} thus simply writes as:

$$0 = \mathcal{S}_L(p, \mathcal{L}) - W - c'[\mathcal{L} - LR(W)].$$

Next, from the properties of \mathcal{S} assumed in the statement of the lemma and from the convexity of $c(\cdot)$, the maximand has increasing differences in (\mathcal{L}, p) and (\mathcal{L}, L) , and “flat” differences in (W, p) as it is additively separable in (W, p) . Moreover, it has increasing differences in (\mathcal{L}, W) if:

$$c''[\mathcal{L} - LR(W)] \times LR'(W) - 1 \geq 0$$

which, substituting the FOC for W (25), holds around a selection from the optimal correspondence if:

$$\frac{[\mathcal{L} - LR(W)] c''[\mathcal{L} - LR(W)]}{c'[\mathcal{L} - LR(W)]} \geq 1,$$

which is true by the assumption $\varepsilon(c') \geq 1$.

Let $p_2 > p_1$, and denote corresponding selections from the optimal correspondence by (W_1, \mathcal{L}_1) and (W_2, \mathcal{L}_2) . Then, by revealed preferences:

$$\begin{aligned} & [\mathcal{S}(p_2, \mathcal{L}_2) + \Phi(L, W_2, \mathcal{L}_2)] - [\mathcal{S}(p_2, \mathcal{L}_1) + \Phi(L, W_1, \mathcal{L}_1)] \geq 0 \\ & \geq [\mathcal{S}(p_1, \mathcal{L}_2) + \Phi(L, W_2, \mathcal{L}_2)] - [\mathcal{S}(p_1, \mathcal{L}_1) + \Phi(L, W_1, \mathcal{L}_1)] \\ & \Leftrightarrow \mathcal{S}(p_2, \mathcal{L}_2) - \mathcal{S}(p_1, \mathcal{L}_2) \geq \mathcal{S}(p_2, \mathcal{L}_1) - \mathcal{S}(p_1, \mathcal{L}_1) \end{aligned}$$

which implies that $\mathcal{L}_2 \geq \mathcal{L}_1$ as \mathcal{S} has increasing differences in (p, L) . Similarly, by revealed preferences:

$$\begin{aligned} & [\mathcal{S}(p_2, \mathcal{L}_2) + \Phi(L, W_2, \mathcal{L}_2)] - [\mathcal{S}(p_2, \mathcal{L}_2) + \Phi(L, W_1, \mathcal{L}_2)] \geq 0 \\ & \geq [\mathcal{S}(p_1, \mathcal{L}_1) + \Phi(L, W_2, \mathcal{L}_1)] - [\mathcal{S}(p_1, \mathcal{L}_1) + \Phi(L, W_1, \mathcal{L}_1)] \\ & \Leftrightarrow \Phi(L, W_2, \mathcal{L}_2) - \Phi(L, W_2, \mathcal{L}_1) \geq \Phi(L, W_1, \mathcal{L}_2) - \Phi(L, W_1, \mathcal{L}_1) \end{aligned}$$

which, together with the facts that $\mathcal{L}_2 \geq \mathcal{L}_1$ and that Φ has increasing differences so long as $\varepsilon(c') \geq 1$, shows $W_2 > W_1$.

This establishes that any selection $(V(p, L), \mathcal{L}(p, L))$ from the optimal correspondence is such that V and \mathcal{L} are both increasing in p . We still need to show that they are also increasing in L . To this end, fix W and consider the optimal choice of \mathcal{L} given W and L (call it $\mathcal{L}^*(W, L)$). The corresponding FOC

writes as:

$$\begin{aligned} \mathcal{S}_{L,r}(p, \mathcal{L}^*(W, L)) - W - c'[\mathcal{L}^*(W, L) - LR(W)] \\ \leq 0 \leq \mathcal{S}_{L,\ell}(p, \mathcal{L}^*(W, L)) - W - c'[\mathcal{L}^*(W, L) - LR(W)]. \end{aligned}$$

As before, convexity of \mathcal{S} combined with the FOC above establishes differentiability of \mathcal{S} at $\mathcal{L}^*(W, L)$, implying that the FOC indeed holds with equality:

$$0 = \mathcal{S}_L(p, \mathcal{L}^*(W, L)) - W - c'[\mathcal{L}^*(W, L) - LR(W)]. \quad (27)$$

Next define the maximand:

$$\mathcal{M}(W, L) = \mathcal{S}(p, \mathcal{L}^*(W, L)) + \Phi(L, W, \mathcal{L}^*(W, L)).$$

This maximand is differentiable w.r.t. L with:

$$\mathcal{M}_L(W, L) = c'[\mathcal{L}^*(W, L) - LR(W)]R(W) + WR(W) + \int_W^{+\infty} v dR(v)$$

Next consider two posted values $W_2 > W_1$. Taking differences of $\mathcal{M}_L(W, L)$:

$$\begin{aligned} \mathcal{M}_L(W_2, L) - \mathcal{M}_L(W_1, L) \\ = c'[\mathcal{L}^*(W_2, L) - LR(W_2)]R(W_2) - c'[\mathcal{L}^*(W_1, L) - LR(W_1)]R(W_1) \\ + W_2R(W_2) - W_1R(W_1) - \int_{W_1}^{W_2} v dR(v) \\ = \{c'[\mathcal{L}^*(W_2, L) - LR(W_2)] - c'[\mathcal{L}^*(W_1, L) - LR(W_1)]\}R(W_1) \\ + c'[\mathcal{L}^*(W_2, L) - LR(W_2)] \times [R(W_2) - R(W_1)] + \int_{W_1}^{W_2} R(v)dv, \end{aligned}$$

where the second equality uses integration by parts. Substituting the FOC (27):

$$\begin{aligned} \mathcal{M}_L(W_2, L) - \mathcal{M}_L(W_1, L) = \{\mathcal{S}_L[p, \mathcal{L}^*(W_2, L)] - \mathcal{S}_L[p, \mathcal{L}^*(W_1, L)]\}R(W_1) \\ + c'[\mathcal{L}^*(W_2, L) - LR(W_2)] \times [R(W_2) - R(W_1)] + \int_{W_1}^{W_2} [R(v) - R(W_1)] dv. \end{aligned}$$

All terms in the r.h.s. of the above equation are positive: the first one by convexity of \mathcal{S} w.r.t. L , and the other two from the fact that $R(\cdot)$ is an increasing function. Now, a solution $(V(p, L), \mathcal{L}(p, L))$ to the initial problem (24) also solves $V = \operatorname{argmax}_W \mathcal{M}(L, W)$ and $\mathcal{L} = \mathcal{L}^*(V, L)$. Because the maximand in the latter problem has increasing differences in (W, L) , it must be the case that $V(p, L)$ is increasing in L . And because $\partial \mathcal{L}^* / \partial L$ and $\partial \mathcal{L}^* / \partial W$ are both positive

(provided that $\varepsilon(c') \geq 1$), it must be the case that $\mathcal{L}(p, L) = \mathcal{L}^*(V(p, L), L)$ is also increasing in L .

Now on to point 4. Inspection of (23), in combination with the fact that \mathcal{S} is differentiable w.r.t. L at any optimal solution, shows that $\mathbf{M}^{\mathbf{N}}\mathcal{S}$ is in fact differentiable w.r.t. L (a stronger property than claimed in the lemma). Furthermore:

$$\begin{aligned} \frac{\partial}{\partial L} \mathbf{M}^{\mathbf{N}} \mathcal{S}(p, L, \omega) &= \omega p + \beta \int_{\Omega} \delta^{\omega'} U(\omega') Q(d\omega' | \omega) \\ &+ \beta \int_{\Omega} \left\{ c' [\mathcal{L}(p, L, \omega') - LR(V(p, L, \omega') | \omega')] R(V(p, L, \omega') | \omega') \right. \\ &\left. + \int_{V(p, L, \omega')}^{+\infty} v dR(v | \omega') + V(p, L, \omega') R(V(p, L, \omega') | \omega') \right\} Q(d\omega' | \omega) \quad (28) \end{aligned}$$

This latter equation readily shows that $\mathbf{M}^{\mathbf{N}}\mathcal{S}$ is increasing in L .

What is left is for us to establish that $\mathbf{M}^{\mathbf{N}}\mathcal{S}$ is convex w.r.t. L and has increasing differences in (L, p) . To this end, first note that substitution of the FOC (26) into (28) yields an alternative expression for $\partial \mathbf{M}^{\mathbf{N}} \mathcal{S} / \partial L$:

$$\begin{aligned} \frac{\partial}{\partial L} \mathbf{M}^{\mathbf{N}} \mathcal{S}(p, L, \omega) &= \omega p + \beta \int_{\Omega} \delta^{\omega'} U(\omega') Q(d\omega' | \omega) \\ &+ \beta \int_{\Omega} \left\{ \mathcal{S}_L(p, \mathcal{L}(p, L, \omega')) \times R(V(p, L, \omega') | \omega') + \int_{V(p, L, \omega')}^{+\infty} v dR(v | \omega') \right\} Q(d\omega' | \omega) \end{aligned}$$

Now take $L_2 > L_1$:

$$\begin{aligned} &\frac{\partial}{\partial L} \mathbf{M}^{\mathbf{N}} \mathcal{S}(p, L_2, \omega) - \frac{\partial}{\partial L} \mathbf{M}^{\mathbf{N}} \mathcal{S}(p, L_1, \omega) \\ &= \beta \int_{\Omega} \left\{ \mathcal{S}_L(p, \mathcal{L}(p, L_2, \omega')) \times R(V(p, L_2, \omega') | \omega') \right. \\ &- \mathcal{S}_L(p, \mathcal{L}(p, L_1, \omega')) \times R(V(p, L_1, \omega') | \omega') - \int_{V(p, L_1, \omega')}^{V(p, L_2, \omega')} v dR(v | \omega') \left. \right\} Q(d\omega' | \omega) \\ &= \beta \int_{\Omega} \left\{ [\mathcal{S}_L(p, \mathcal{L}(p, L_2, \omega')) - \mathcal{S}_L(p, \mathcal{L}(p, L_1, \omega'))] \times R(V(p, L_1, \omega') | \omega') \right. \\ &\quad \left. + \int_{V(p, L_1, \omega')}^{V(p, L_2, \omega')} [\mathcal{S}_L(p, \mathcal{L}(p, L_2, \omega')) - v] dR(v | \omega') \right\} Q(d\omega' | \omega) \end{aligned}$$

which is positive, by convexity of \mathcal{S} , the fact that both V and \mathcal{L} are increasing in L , and the fact that $\mathcal{S}_L(p, \mathcal{L}(p, L_2, \omega')) \geq V(p, L_2, \omega')$. This establishes

convexity of $\mathbf{M}^N \mathcal{S}$ w.r.t. L . Increasing differences in (L, p) is follows from similar steps (only taking differences in p of $\partial \mathbf{M}^N \mathcal{S} / \partial L$), so we omit the details.

□

Now consider the set of functions defined over $[\underline{p}, \bar{p}] \times [0, 1] \times \Omega$ that are continuous in (p, L) and call it $\mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$. That set is a Banach space when endowed with the sup norm. As Lemma 1 suggests we will be interested in the properties of a subset $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega} \subset \mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ of functions that are increasing and convex in L and have increasing differences in (p, L) . We next state two ancillary lemmas, which will establish as a corollary (Corollary 1) that $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ is closed in $\mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ under the sup norm.

Lemma 2 *Let X be an interval in \mathbb{R} and $f_n : X \rightarrow \mathbb{R}$, $N \in \mathbb{N}$ such that $\{f_n\}$ converges uniformly to f . Then:*

1. *if f_n is nondecreasing for all n , so is f ;*
2. *if f_n is convex for all n , so is f .*

Proof. See the appendix of MPV13. □

Lemma 3 *Let $X \subset \mathbb{R}^2$ be a convex set and $f_n : X \rightarrow \mathbb{R}$, $N \in \mathbb{N}$ be functions with increasing differences such that $\{f_n\}$ converges uniformly to f . Then f has increasing differences.*

Proof. See the appendix of MPV13. □

Corollary 1 *The set $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ of functions defined over $[\underline{p}, \bar{p}] \times [0, 1] \times \Omega$ that are increasing and convex in L and have increasing differences in (p, L) is a closed subset of $\mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ under the sup norm.*

The latter corollary establishes that, given a fixed N , the functions that are relevant to Lemma 1 live in a closed subset of a Banach space of functions under the sup norm. The following lemma shows that the operator considered in Lemma 1 is a contraction under that same norm.

Lemma 4 *The operator \mathbf{M}^N defined in (23) maps $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ into itself and is a contraction of modulus β under the sup norm.*

Proof. That \mathbf{M}^N maps $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ into itself flows directly from part of the proof of Lemma 1. To prove that \mathbf{M}^N is a contraction, we can easily check using (23) that \mathbf{M}^N satisfies Blackwell's sufficient conditions with modulus β . □

We are now in a position to prove the proposition. Given the initially fixed N , the operator \mathbf{M}^N , which by Lemma 4 is a contraction from $\mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ into itself, and has a unique fixed point \mathcal{S}_N in that set (by the Contraction Mapping Theorem). Moreover, since $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ is a closed subset of $\mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ (Lemma 2) and since \mathbf{M}^N also maps $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$ into itself (Lemma 1), that fixed point \mathcal{S}_N belongs to $\mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$.

Summing up, what we have established thus far is that for any fixed $N \in \mathcal{C}_{[\underline{p}, \bar{p}]}$, the operator \mathbf{M}^N over functions of (p, L, ω) has a unique, bounded and continuous fixed point $\mathcal{S}_N^* = \mathbf{M}^N \mathcal{S}_N^* \in \mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega} \subset \mathcal{C}_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$.

We finally turn to the Bellman operator \mathbf{M} which is relevant to the firm's problem. That operator \mathbf{M} applies to functions $\overline{\mathcal{S}}$ defined on $[\underline{p}, \bar{p}] \times [0, 1] \times \Omega \times \mathcal{C}_{[\underline{p}, \bar{p}]}$ and is defined as the following "extension" of \mathbf{M}^N :

$$\begin{aligned} \mathbf{M}\overline{\mathcal{S}}(p, L, \omega, N) := & \varphi(p, L, \omega, N) + \beta \int_{\Omega} \max_{W(\omega'), \mathcal{L}(\omega')} \left\langle \overline{\mathcal{S}}[p, \mathcal{L}(\omega'), \omega'] \right. \\ & \left. + \Phi(L, W(\omega'), \mathcal{L}(\omega'), \omega', N) \right\rangle Q(d\omega' | \omega). \end{aligned}$$

If an equilibrium exists, then a firm has a best response and a value S which solves $S = \mathbf{M}S$. For every $N \in \mathcal{C}_{[\underline{p}, \bar{p}]}$, by definition of \mathbf{M} and \mathbf{M}^N this implies $S = \mathbf{M}^N S$. Since the fixed point of \mathbf{M}^N is unique, if $S = \mathbf{M}S$ exists then for every fixed $N \in \mathcal{C}_{[\underline{p}, \bar{p}]}$ we have for all $(p, L, \omega) \in [\underline{p}, \bar{p}] \times [0, 1] \times \Omega$: $S(p, L, \omega, N) = \mathcal{S}_N^*(p, L, \omega)$. Therefore, if the value function S and an equilibrium of the contract-posting game exist, then $S \in \mathcal{C}'_{[\underline{p}, \bar{p}] \times [0, 1] \times \Omega}$: the typical firm's value function is continuous in p and L , increasing and convex in L and has increasing differences in (p, L) . By the same standard comparative statics arguments that we invoked in the proof of Lemma 1, the maximizing correspondence is increasing in p and L in the strong set sense, hence all of its measurable selections are weakly increasing in p and L .

B Derivation of steady state wage policy

Let $F(w) = \int^w \frac{a(x)}{A} d\Gamma(x)$ and $G(w) = \int^w L(x) d\Gamma(x)$ denote the wage offer distribution and the cross-section (earned) wage distributions, respectively. In steady-state equilibrium the size of a firm offering a wage w is

$$L(w) = \frac{d \int^w L(x) d\Gamma(x)}{dw} = \frac{a(w)}{A} \frac{d \int^w L(x) d\Gamma(x)}{d \int^w \frac{a(x)}{A} d\Gamma(x)} = \frac{a(w)}{A} \frac{dG(w)}{dF(w)}.$$

To find an expression for dG/dF , we write the condition for balanced flows below any wage w : the inflow, which is only from unemployment, equals the outflow, to unemployment and to firms offering more than w :

$$\lambda [1 - N(\bar{p})] F(w) = [\delta + s(1 - \delta)\lambda \bar{F}(w)] N(\bar{p}) G(w)$$

Evaluating at the maximum wage \bar{w} offered and earned we obtain $\lambda [1 - N(\bar{p})] = \delta N(\bar{p})$, so we can solve for

$$G(w) = \frac{\delta F(w)}{\delta + s(1 - \delta)\lambda [1 - F(w)]}.$$

Taking the Radon-Nykodim derivative

$$L(w) = \frac{a(w)}{A} \frac{dG(w)}{dF(w)} = \frac{a(w)}{A} \frac{1 + s\kappa}{[1 + s\kappa \bar{F}(w)]^2}$$

where $\kappa = \lambda(1 - \delta)/\delta$ denotes is an inverse measure of frictions, the ratio between the probabilities with which an unemployed worker receives job offers and an employed worker is thrown off the ladder back to unemployment. Finally, the hiring inflow is $H(a, w) = a\eta P(w)$, where the chance that an offer is accepted equals

$$P(w) = \frac{1 - N(\bar{p}) + s(1 - \delta)N(\bar{p})G(w)}{Z}.$$

The firm maximizes steady-state flow profits, which equal flow output minus wage per worker, times firm size, minus search costs. Denoting the optimal level of profit achieved by a type- p firm as $\pi(p)$:

$$\pi(p) = \max_{\substack{a \geq 0 \\ w \geq \max\{\phi, w_{\min}\}}} L(w) (\omega p - w) - c(a\eta P(w)).$$

Denoting the solution by $w(p)$, $a(p) = a(w(p))$, the Envelope theorem ensures:

$$\pi'(p) = \omega L(w(p)) = \frac{a(p)}{A} \frac{\omega(1 + s\kappa)}{\left[1 + s\kappa \int_p^{\bar{p}} \frac{a(x)}{A} d\Gamma(x)\right]^2}$$

Integrating from $\underline{p}(\omega)$ to p :

$$\pi(p) = \pi(\underline{p}(\omega)) + (1 + s\kappa)\omega \int_{\underline{p}(\omega)}^p \frac{\frac{a(x)}{A}}{\left[1 + s\kappa \int_x^{\bar{p}} \frac{a(y)}{A} d\Gamma(y)\right]^2} dx$$

Equating this expression to the r.h.s. of (18) we can solve for the wage function (19), where

$$\underline{\pi} = \pi(\underline{p}(\omega), a(\underline{p}(\omega)), \underline{p}(\omega)) \mathbb{I}\{\omega \underline{p} > \phi, w_{\min}\} \quad (29)$$

are the profits of the marginal firm in the market, which are zero unless neither the reservation wage ϕ nor the statutory minimum wage are binding.

To close the model, we find the reservation wage ϕ , thus the lowest productivity firm active in the market $\underline{p}(\omega) = \max\left\{\underline{p}, \frac{\phi}{\omega}, \frac{w_{\min}}{\omega}\right\}$. To ease notation, here we omit ω as an argument of the endogenous variables. The value of unemployment solves

$$U = b + \beta \left[U + \lambda \int_{\phi}^{\bar{w}} [W(x) - U] dF(x) \right]$$

and the value to the worker of earning a wage w is

$$W(w) = w + \beta \left[\delta U + (1 - \delta) W(w) + (1 - \delta) s \lambda \int_w^{\bar{w}} [W(x) - W(w)] dF(x) \right]$$

so that $W'(w) = 1 / [1 - \beta(1 - \delta)(1 - s\lambda\bar{F}(w))]$. Substituting those value functions into the definition of the reservation wage, $W(\phi) = U$, and using integration by parts, one obtains:

$$\begin{aligned} \phi &= b + \beta \lambda [1 - (1 - \delta) s] \int_{\phi}^{\bar{w}} [W(x) - U] dF(x) = b + \beta \lambda [1 - (1 - \delta) s] \int_{\phi}^{\bar{w}} W'(x) \bar{F}(x) dx \\ &= b + \beta \lambda [1 - (1 - \delta) s] \int_{\phi}^{\bar{w}} \frac{\bar{F}(x)}{1 - \beta(1 - \delta)(1 - s\lambda\bar{F}(x))} dx \\ &= b + \beta \kappa [1 - (1 - \delta) s] \int_{\phi}^{\bar{w}} \frac{\int_x^{\bar{p}} \frac{a(y)}{A} d\Gamma(y)}{1 - \beta(1 - \delta) \left(1 - s\delta\kappa \int_x^{\bar{p}} \frac{a(y)}{A} d\Gamma(y) \right)} dx. \end{aligned}$$

C The computation algorithm

We now describe the algorithm we implement to simulate the dynamic equilibrium of our model. We simulate the model over an observation window $t = 0, 1, \dots, T$ of chosen length T and a K -point grid of firm productivity values $\mathcal{P} = \{p_1, \dots, p_K\}$ approximating the continuous support of the distribution Γ_0 (the grid is such that $\underline{p} = p_1 < p_2 < \dots < p_K = \bar{p}$). We initiate the simulation by drawing a random path of the aggregate state $\{\omega_t^*\}_{t=0}^T$, which also determines the path of the job destruction rate, $\delta_t = \delta(\omega_t^*)$, and an initial condition for the distribution of employment across productivity types, $L_0^*(p_k)$.

Simulating the model involves solving for the paths of the following endogenous variables, for all grid points $p_k \in \mathcal{P}$: $L_t^*(p_k)$, $\mu_t^*(p_k)$, $V_t^*(p_k)$, $a_t^*(p_k)$, and $w_t^*(p_k)$.²¹ We use a star to denote “observed” quantities, i.e. quantities that are realized in the aggregate states $\{\omega_t^*\}_{t=0}^T$ that our simulated economy does visit. Those have to be distinguished from “latent” quantities, denoted by $\mu_t(p_k|\omega)$, $V_t(p_k|\omega)$, etc., which would be observed at date t if the economy’s aggregate state at date t were ω instead of ω_t^* . The relationship between observed and latent quantities for any variable X is $X_t^*(p_k) = X_t(p_k|\omega_t^*)$.

Simulation proceeds by iteration over successive guesses about the path of firm search effort, $a_t(p_k|\omega)$. The typical iteration n takes an initial guess $a_t^{(n)}(p_k|\omega)$ as input, and goes as follows:

²¹Aggregate variables, such as $N^*(\cdot)$, A_t^* , $F_t^*(p_k)$, Z_t^* , etc. are then constructed from those firm-level series by integration over the grid \mathcal{P} .

1. Simulating the employment distribution. Based on the initial guess $a_t^{(n)}(p_k|\omega)$, we construct the aggregate firm and offer sampling distributions as per the model's equations:

$$A_t^{(n)}(\omega) = \int_{\underline{p}}^{\bar{p}} a_t^{(n)}(x|\omega) d\Gamma(x) \quad \text{and} \quad F_t^{(n)}(\omega) = \frac{1}{A_t^{(n)}(\omega)} \int_{\underline{p}}^{\bar{p}} a_t^{(n)}(x|\omega) d\Gamma(x)$$

We then iterate forward in time, starting at date $t = 0$ from our chosen initial condition $L_0^*(\cdot)$, to construct the aggregate worker search effort that applies between periods t and $t + 1$, conditional on any given state ω :

$$Z_{t+1}^{(n)}(\omega) = 1 - N_t^{*(n)}(\bar{p}) + s(1 - \delta(\omega)) N_t^{*(n)}(\bar{p}),$$

the job finding and vacancy filling rates that apply between periods t and $t + 1$:

$$\lambda_{t+1}^{(n)}(\omega) = \frac{m \left(Z_{t+1}^{(n)}(\omega), A_{t+1}^{(n)}(\omega) \right)}{Z_{t+1}^{(n)}(\omega)} \quad \text{and} \quad \eta_{t+1}^{(n)}(\omega) = \lambda_{t+1}^{(n)}(\omega) \times \frac{Z_{t+1}^{(n)}(\omega)}{A_{t+1}^{(n)}(\omega)},$$

the advert yield using equation (13):

$$Y_{t+1}^{(n)}(p_k|\omega) = \frac{1 - N_t^{*(n)}(\bar{p}) + s(1 - \delta(\omega)) N_t^{*(n)}(p_k)}{Z_{t+1}^{(n)}(\omega)}$$

and the date- $t + 1$ observed employment distribution using the law of motion of firm-level employment:

$$L_{t+1}^{*(n)}(p_k) = L_t^{*(n)}(p_k) (1 - \delta(\omega_{t+1}^*)) \left[1 - s\lambda_{t+1}^{(n)}(\omega_{t+1}^*) \left(1 - F_{t+1}^{(n)}(p_k|\omega_{t+1}^*) \right) \right] + a_{t+1}^{(n)}(p_k|\omega_{t+1}^*) \eta_{t+1}^{(n)}(\omega) Y_{t+1}^{(n)}(p_k|\omega_{t+1}^*).$$

2. Simulating job values. After completion of stage 1, we simulate job values by backward recursion over time, starting from the final date T . More precisely, we construct the following objects: $\mu_t^{(n)}(p_k|\omega) - U_t^{(n)}(\omega)$ and $V_t^{(n)}(p_k|\omega) - U_t^{(n)}(\omega)$ given the terminal condition $\mu_T^{(n)}(p_k|\omega) - U_T^{(n)}(\omega) = \omega p_k - b(\omega)$. While obviously inaccurate, this terminal condition has a vanishing impact on the simulated solution at earlier dates, because of discounting, as $T \rightarrow +\infty$. We thus choose a distant final date T and discard the last $T - T^*$ periods of the simulation (which are affected by the approximate terminal condition), thus ending up with a “valid” observation window of length T^* .

The recursion then proceeds as follows. Starting from $t = T - 1$, we first solve the FOC (16) for worker values:

$$V_{t+1}^{(n)}(p_k|\omega) - U_{t+1}^{(n)}(\omega) = \frac{1}{Y_{t+1}^{(n)}(p_k|\omega)} \int_{\underline{p}}^{p_k} \left[\mu_{t+1}^{(n)}(x|\omega) - U_{t+1}^{(n)}(\omega) \right] dY_{t+1}^{(n)}(x|\omega).$$

We then construct $\mu_t^{(n)}(p_k|\omega) - U_t^{(n)}(\omega)$ combining the Euler equation (14) and the Bellman equation defining the value of unemployment, (5):

$$\begin{aligned} \mu_t^{(n)}(p_k|\omega) - U_t^{(n)}(\omega) &= \omega p - b(\omega) \\ &+ \beta \mathbf{E}_{\omega'|\omega} \left\{ (1 - \delta(\omega')) \left[1 - s\lambda_{t+1}^{(n)}(\omega') \left(1 - F_{t+1}^{(n)}(p_k|\omega') \right) \right] \left(\mu_{t+1}^{(n)}(p_k|\omega) - U_{t+1}^{(n)}(\omega) \right) \right. \\ &\quad + (1 - \delta(\omega')) s\lambda_{t+1}^{(n)}(\omega') \int_p^{\bar{p}} \left[V_{t+1}^{(n)}(x|\omega') - U_{t+1}^{(n)}(\omega') \right] dF_{t+1}^{(n)}(x|\omega') \\ &\quad \left. - \lambda_{t+1}^{(n)}(\omega') \int_p^{\bar{p}} \left[V_{t+1}^{(n)}(x|\omega') - U_{t+1}^{(n)}(\omega') \right] dF_{t+1}^{(n)}(x|\omega') \right\} \end{aligned}$$

3. Updating job adverts. The final step of the n^{th} iteration is to update the distribution of posted job adverts. This is done using the FOC (15)

$$a_t^{(\text{updated})}(p_k|\omega) = c'^{-1} \left(\mu_t^{(n)}(p_k|\omega) - V_t^{(n)}(p_k|\omega) \right) \cdot \frac{1}{\eta_t^{(n)}(\omega) Y_t^{(n)}(p_k|\omega)}$$

If the distance between $a_t^{(\text{updated})}(p_k|\omega)$ and $a_t^{(n)}(p_k|\omega)$ is less than a convergence tolerance, we stop iterating and use the current state of the model as our simulation output. If not, we continue iterating using $a_t^{(n+1)}(p_k|\omega) = a_t^{(\text{updated})}(p_k|\omega)$ as initial guess.²²

Some additional remarks about this simulation protocol are in order. First, the iterative procedure outlined above only produces a subset of the endogenous variables of interest — the subset of variables that are needed to update the job advert profile until convergence. The missing variables (chiefly, wages and the value of unemployment) are constructed only once in a final run of the loop described above, after convergence of the job advert profile.

Second, for clarity of exposition, we described the simulation protocol in the “simple” case where there is no binding minimum wage. Our computer code can handle this extra feature, however. The code is available from:

<https://sites.google.com/site/fabienpostelvinay/working-papers>

and details of the extra steps that must be taken in the simulation are available upon request.

Finally, execution time obviously depends on the fineness of the productivity grid \mathcal{P} , the size of the aggregate state space Ω , and the length of the simulation window T . For the values we use (a 100-point grid for \mathcal{P} , 20 different states for ω , and $T = 720$ months), execution time ranges from a few seconds when the simulation is started with an initial guess $a_t^{(0)}(p_k|\omega)$ which is “close” to the solution to a few minutes when starting from a coarse approximation.

²²In practice, we actually update the guess as $a_t^{(n+1)}(p_k|\omega) = r \times a_t^{(\text{updated})}(p_k|\omega) + (1 - r) \times a_t^{(n)}(p_k|\omega)$, where $r \in (0, 1)$ is a “relaxation parameter”. This helps the algorithm to converge faster and monotonically.