Online Appendices

A Description of the data

We use wage data for individual workers in the CPS outgoing rotation groups from 1979 to 2006. We match these workers to the three preceding basic monthly datafiles in order to construct four months (one quarter) of employment history, which we use to identify newly hired workers.

A1. Wages from the CPS outgoing rotation groups

We consider only wage and salary workers that are not self-employed and report non-zero earnings and hours worked. Both genders and all ages are included in our baseline sample. Our wage measure is hourly earnings (on the main job) for hourly workers and weekly earnings divided by usual weekly hours for weekly workers. For weekly workers who report that their hours vary (from 1994 onwards), we use hours worked last week. Top-coded weekly earnings are imputed assuming a log-normal cross-sectional distribution for earnings, following Schmitt (2003), who finds that this method better replicates aggregate wage series than multiplying by a fixed factor or imputing using different distributions. Notice that the imputation of top-coded earnings affects the mean, but not the median wage.

Outliers introduce extra sampling variation. Therefore, we apply mild trimming to the cross-sectional distribution of hours worked (lowest and highest 0.5 percentile) and hourly wages (0.3 percentiles). These values roughly correspond to USD 1 per hour and USD 100 per hour at constant 2002 dollars, the values recommended by Schmitt (2003). We prefer trimming by quantiles rather than absolute levels because (*i*) it is symmetric and therefore does not affect the median, (*ii*) it is not affected by real wage growth and (*iii*) it is not affected by increased wage dispersion over the sample period. We also check that our results are robust to using median wages, which are less affected by outliers.

We do not correct wages for overtime, tips and commissions, because (*i*) the relevant wage for our purposes is the wage paid by employers, which includes these secondary benefits, (*ii*) the data necessary to do this are not available over the whole sample period, and (*iii*) this correction has very little effect on the average wage (Schmitt 2003). We also do not exclude allocated earnings because (*i*) doing so might bias our estimate for the average wage and (*ii*) allocation flags are not available for all years and (*iii*) even if they are only about 25% of allocated observations are flagged as such (Hirsch and Schumacher 2004).

Mean and median wages in a given month are weighted by the appropriate sampling weights (the earnings weights for the outgoing rotation groups) and by hours worked, following Abraham et al. (1999) and Schmitt (2003). We explore robustness to the weights and confirm the finding of these papers that hours weighted series better replicate the aggregate wage. Average mean or median wages in a quarter are simple averages of the monthly mean or median wages. Consistent with the literature, we consider mean log wages rather than log mean wages.

In order to correct the business cycle statistics for the wage for sampling error (see Online Appendix B), we calculate standard errors for mean and median wages. Standard errors for the mean are simply the standard deviation of the wage divided by the square root of the number of observations. Medians are also asymptotically normal, but their variance is downward biased in small samples. Therefore, we bootstrap these standard errors.

We seasonally adjust our wage series by regressing the log wage on quarter dummies. Nominal wages are deflated by the implicit deflator for hourly earnings in the private non-farm business sector (chain-weighted) from the BLS productivity and costs program. Using different deflators affects the results very little, but decreases the correlation of our wage series with the aggregate wage.

Our baseline sample includes non-supervisory workers in the private non-farm business sector. This subsample of workers gives the best replication of the aggregate wage in terms of its correlation with hourly compensation from the establishment survey and in terms of its volatility, persistence and comovement with other variables.[26] We identify private sector workers using reported 'class of worker'. We construct an industry classification that is consistent over the whole sample period (building on the NBER consistent industry classification but extending it for data from 2003 onwards) and use it to identify farm workers. Similarly, we identify supervisory workers using reported occupation. Because of the change in the BLS occupation classification in 2003, there is a slight jump in the fraction of supervisory workers from 2002:IV to 2003:I. It is not possible to distinguish supervisory workers in agriculture or the military, so all workers in these sectors are excluded in the wage series for non-supervisory workers.

Finally, in order to control for composition bias because of heterogeneous workers (see section 2.2.), we need additional worker characteristics to use in a Mincerian earnings regression. Dummies for females, blacks, hispanics and married workers (with spouse present) are, or can be made, consistent over the sample period. We construct a consistent education variable in five categories as well as an almost consistent measure for years of schooling following Jaeger (1997) and calculate potential experience as age minus years of schooling minus six.

A2. Identifying newly hired workers

We match the individuals in the outgoing rotation groups to the three preceding basic monthly data files using the household identifier, household number (for multiple households on one address), person line number (for multiple wage earners in one household), month-in-sample and state. To identify mismatches, we use the s|r|a criterion proposed by Madrian and Lefgren (2000): a worker is flagged as a mismatch if gender or race changes between two subsequent months or if the difference

[26] Detailed results for this replication exercise are available in a previous version of this paper (July 2007), available from our websites.

in age is less than 0 or greater than 2 (to allow for some measurement error in the reported age). Madrian and Lefgren show that this criterion performs well in the trade-off between false matches and false mismatches. Within the set of measures that they find to perform well, s|r|a is the strictest. We choose a strict criterion because mismatches are more likely to be classified as newly hired workers (see below) and are therefore likely to affect our results substantially.

We can credibly match about 80% of workers in the outgoing rotation group to all three preceding monthly files. Because of changes in the sample design, we cannot match sufficiently many individuals to the preceding four months in the third and fourth quarter of 1985 and in the third and fourth quarter of 1995, so that the wage series for validly matched workers, job stayers and new hires have missing values in those quarters. In our regressions, we weight quarters by the variance of the estimate for the mean or median wage so that quarters with less than average number of observations automatically get less weight.

Including the outgoing rotation group itself, the matched data include four months employment history (employed, unemployed or not-in-the-labor-force), which we obtain from the BLS labor force status recode variable. We use this employment history to identify newly hired workers and workers in ongoing job relationships. New hires are defined as workers that were either unemployed or not in the labor force for any of the preceding three months. Job stayers are identified as workers that were employed for all four months. Notice that the two groups are not comprehensive for the group of all workers, because workers that cannot be matched to all preceding months can not always be classified.

B Correcting business cycle statistics for sampling error

We estimate wages for all workers, job stayers and new hires from an underlying micro-data survey. Therefore, our wage series are subject to sampling error. Given the way we construct these series, we know three things about the sampling error. First, because there is no overlap between individuals included in the outgoing rotation groups in two subsequent quarters, the sampling error is uncorrelated over time.[27] Second, because the sampling error in each period is the error associated with estimating a mean (or median), it is asymptotically normally distributed. Third, we have an estimate for the standard deviation of the sampling error in each quarter, which is given by the standard error of the mean (or median) wage in that quarter. Notice that taking first difference exacerbates the measurement error, increasing the standard deviation by a factor $\sqrt{2}$. Because of these three properties, and because the estimated standard errors are stable over time, we can treat the sampling error as classical measurement error, which is independent and

[27] Individuals in the CPS are interviewed four months in a row, the last one of which is an outgoing rotation group, then leave the sample for eight months, after which they are interviewed another four months, the last one of which is again an outgoing rotation group. Therefore, about half of the sample in quarter t (individuals in rotation group 8) is also included in the sample in quarter $t-4$ (when they were in rotation group 4) and the other half is included in the sample in quarter $t+4$. Thus, the sampling error may be correlated with a four quarter lag, but not between subsequent quarters. We ignore this correlation structure and treat the sampling error as uncorrelated over time.

identically distributed.

Let w_t denote an estimated wage series, $w_t = w_t^* + \varepsilon_t$, where w_t^* is the true wage and ε_t is the sampling error in the wage, which is uncorrelated over time and with w_t^* and has a known variance σ^2. The business cycle statistics we consider are the standard deviation of w_t^*, the autocorrelation of w_t^* and the correlation of w_t^* with x_t, an aggregate variable that is not subject to measurement error. These statistics can be calculated from the estimated wage series w_t and the estimated standard deviation of the sampling error σ as follows.

$$var\left(w_t\right) = var\left(w_t^*\right) + \sigma^2 \Rightarrow sd\left(w_t^*\right) = \sqrt{R} \cdot sd\left(w_t\right) \tag{8}$$

$$cov\left(w_t, w_{t-1}\right) = cov\left(w_t^*, w_{t-1}^*\right) \Rightarrow corr\left(w_t^*, w_{t-1}^*\right) = \frac{corr\left(w_t, w_{t-1}\right)}{R} \tag{9}$$

$$cov\left(w_t, x_t\right) = cov\left(w_t^*, x_t\right) \Rightarrow corr\left(w_t^*, x_t\right) = \frac{corr\left(w_t, x_t\right)}{\sqrt{R}} \tag{10}$$

where $R = \left(var\left(w_t\right) - \sigma^2\right)/var\left(w_t\right) \in (0,1)$ is the fraction of signal in the variance of w_t. Unless explicitly specified, we use the correction factors $\sqrt{R}$, $1/R$ and $1/\sqrt{R}$ for all reported business cycle statistics. This bias correction is small for the wages of all workers and job stayers, because sample sizes are large and therefore σ^2 is small, but substantial for the wage of new hires. Notice that the bias correction decreases the reported standard deviations towards zero but increases the reported autocovariances and correlation coefficients away from zero. For bandpass filtered series no correction is necessary because the filter removes the high-frequency fluctuations due to measurement error from the data. Regression coefficients for the wage on labor productivity are not biased in the presence of classical measurement error in the dependent variable so no correction is necessary.

C Derivation of job creation condition (5)

In a search and matching model, e.g. as in Pissarides (1985, 2000) or Shimer (2005), free entry drives the value of a vacancy to zero, which implies that the period cost $c\left(q_t\right)$ must equal the probability that the vacancy transforms in a match times the expected value of that match.

$$c\left(q_t\right) = E_t J_{t+1} \tag{11}$$

The value to the firm of having a filled job J_t, is given by the following Bellman equation.[28]

$$(1+r)J_t = y_t - w_t + (1-\delta)E_t J_{t+1}. \tag{12}$$

[28] We write the model in discrete time but assume that all payments are made at the end of the period, so that the expressions look similar to the continuous time representation.

Solving equation (12) forward gives an expression for the value of a filled job.

$$E_t J_{t+1} = \frac{\bar{y}_t - \bar{w}_t}{r + \delta} \tag{13}$$

Substituting (13) into (11) gives the job creation equation in the main text.

For many other models, some details may be different, but the condition will still look very similar and the results in the main text will go through. For example, the separation probability δ may be time-varying as in Mortensen and Pissarides (1994), productivity y_t may represent the marginal product of labor and depend on capital as in Merz (1995) and Andolfatto (1996), firms may have multiple workers as in Rotemberg (2008) or Ebell and Haefke (2009), participation may be endogenous as in Haefke and Reiter (2011) or expectations about future productivity and wages may include the option value of moving into a different job if there is on-the-job search as in Menzio and Shi (2010). An identical job creation condition can also be derived in a model without search frictions but with worker heterogeneity, as in Merkl and van Rens (2012).

D Long-term wage contracting and job creation

In the main text of the paper, we interpret our estimates for the cyclicality of the wages of newly hired workers and workers in ongoing matches as the cyclicality of wages at the start and over the duration of individual wage contracts. This interpretation is only approximately correct, because of compositional changes in our dataset. The pool of new hires in a given quarter does not include the same workers as new hires in the quarter before. And the pool of workers in ongoing matches includes workers that were newly hired only last quarter as well as workers that have been in their current job for a long time. Nevertheless, our estimates are of course informative about the cyclicality of individual wage contracts. Here, we formalize that link.

D1. Parameters of the long-term wage contracts

The wage w_{it}^a of a worker i in a match of age a at time t consists of four components: the initial wage this worker received at the time of hiring $w_{i,t-a}^0$, wage growth with job tenure, revisions to the wage in response to changes in aggregate economic conditions, and changes in the wage because of idiosyncratic circumstances. For simplicity, we assume the functional form of the wage contract is log-linear, like our estimation equation (3), so that the wage is given by,

$$\log w_{it}^a = \log w_{i,t-1}^{a-1} + \phi_{0,\text{stay}} + \phi_{1,\text{stay}} \left(\log y_t - \log y_{t-1}\right) + v_{it} \tag{14}$$

where $\phi_{0,\text{stay}}$ is average wage growth per period of tenure, $\phi_{1,\text{stay}}$ is the response of the wages in ongoing matches to aggregate productivity, and v_{it} is idiosyncratic wage growth, which averages zero over the cross-section in each period. The question is what values for $\phi_{0,\text{stay}}$ and $\phi_{1,\text{stay}}$ are consistent with our estimates.

We simulate wages using wage contract (14) for 1.5 million workers over 158 periods, dropping the first 70 periods to initialize the wage distribution so that our sample of simulated data, like the actual data, consists of 88 quarters. In these simulations, we assume the idiosyncratic component of wage growth v_{it} is normally distributed, so that cross-sectional distribution of wages is log normal, although this assumption does not matter for the result because invididual heterogeneity is averaged out. In order to replicate the compositional changes in the actual data, we also need to model when contracts start and end. To this end, we match the number of separations and new hires in each period. Notice that this strategy yields an employment rate that is consistent with the data as well. Finally, we assume stochastic processes for productivity and wages of new hires so that we can forecast both variables to compute expected values and backcast wages in order to initialize the wage distribution. We assume wages of new hires depend log-linearly on productivity, $\log w_{it}^0 = \phi_{0,\text{newh}} + \phi_{1,\text{newh}} \log y_t + v_{it}^0$, setting $\phi_{0,\text{newh}}$ and $\phi_{1,\text{newh}}$ to match the average wage of newly hired workers and the elasticity of the wage of new hires with respect to productivity. For productivity we assume a simple ARIMA$(1, 1, 0)$ process, $\log y_t = \log y_{t-1} + \psi_0 + \psi_1 (\log y_{t-1} - \log y_t) + v_t$, where ψ_0 and ψ_1 are estimated directly from the data.[29] Then, we vary $\phi_{0,\text{stay}}$ and $\phi_{1,\text{stay}}$ so that average wage growth α_{allw} and the elasticity of wages with respect to productivity η_{allw}, as in equation (3), estimated from the simulated data are the same as in the actual data.

The goal of the simulation is to obtain individual wage histories, which – when aggregated – yield the same aggregate behavior as the observed series for all workers. In order to simulate the individual wage histories, several data needs to be known:

1. Productivity: taken from the actual data, or otherwise based on coefficients estimated from actual data. When estimated we assume a specification as represented in equation (15).

2. Aggregate wage series of newly hired workers: taken from the actual data, or otherwise based on coefficients estimated from actual data. When estimated we assume a specification as represented in equation (16).

3. Form of the Wage Contract for Continuing Workers: We do not directly observe the individual wage path for continuing workers. There are two obvious components in the wage contract for continuing workers, one is how these wages react to productivity, the other one is a constant expected growth rate, the equation is (17). We pick $\phi_{1,\text{stay}}$ so that the regression coefficient for the simulated wage of all workers on productivity matches the one from the data. However, in the data the average wage of all workers is substantially larger than the average wage of newly hired workers. This difference in the two aggregate wage series comes from the growth

[29] Within the class of ARIMA$(p, 1, q)$ processes, the ARIMA$(1, 1, 0)$ specification fits the data best according to the Bayesian Information Criterion. Moreover, the estimate for ψ_1 is small, so that productivity is close to a random walk. As a robustness check, we repeat the exercise with actual data for the wages of new hires and productivity for the period these data are available, using simulated data only for the backcasting, and find the results are very similar.

of wages on the job independent of productivity. The role of $\phi_{0,\text{stay}}$ is to allow for this average wage difference.

4. Separation Probabilities, δ_t: We simulate the data at a quarterly frequency and take the observed number of workers and number of separations from our CPS dataset to find the quarterly separation probability:

$$
\begin{aligned}
S_{t-1} &= E_{t-1} + N_t - E_t \\
\delta_t &= \frac{S_t}{E_t}
\end{aligned}
$$

5. Job Finding Probabilities: Quarterly job finding probabilities are very high. In order to avoid these probabilities to exceed one, we assume that workers who were separated can immediately search again. This is a completely innocuous assumption for our purposes because we only care about the evolution of wages on the job and not about what happens during unemployment. Picking job finding and separation probabilities in this way guarantee that the employment path in our simulations coincides with actual employment from the data.

6. Interest Rate: To compute the present value, an interest rate needs to be used, we take it from FRED, either the three month T-bill rate (TB3MS) or the bank prime loan rate (MPRIME).

We assume that productivity and wages are described by:

$$
\begin{aligned}
\log y_t &= \log y_{t-1} + \psi_0 + \psi_1 \left(\log y_{t-1} - \log y_{t-2}\right) + \upsilon_t & (15)\\
\log w_t^0 &= \phi_{0,\text{newh}} + \phi_{1,\text{newh}} \log y_t + \upsilon_t^0 & (16)\\
\log w_t^a &= \log w_{t-1}^{a-1} + \phi_{0,\text{stay}} + \phi_{1,\text{stay}} \left(\log y_t - \log y_{t-1}\right) + \nu_t & (17)\\
\upsilon_t^0 &\sim i.i.d\, \mathcal{N}(0, \sigma_0^2) \quad \forall t & (18)\\
\nu_t &\sim i.i.d\, \mathcal{N}(0, \sigma^2) \quad \forall t & (19)
\end{aligned}
$$

Denote by $\log w_t^a$ the wage of a worker in period t who was hired in period $t - a \leq t$. Thus the wage $\log w_t^0$ is the newly hired wage.

We also see that log-productivity follows an AR(1) in first differences or an AR(2) in levels. Wages for stayers are specified to grow for two reasons. $\phi_{0,\text{stay}}$ denotes a constant, autonomous growth rate, and $\phi_{1,\text{stay}}$ determines how strongly wages of stayers grow with productivity. The parameters of equation (17) will be determined by the simulation estimator, equations (15) and (16) are estimated directly from the data.

We simulate for the 88 quarters which we are also using for the original data analysis. All simulations are executed for 1 500 000 individuals. In the simulation we proceed in several steps. First all the data is processed and the empirical process for aggregate newly hired wages and labor productivity are estimated. To generate a wage distribution at the beginning of 1984, we start the simulation 70 quarters earlier, so that by the time 1984 is reached, we have a nice distribution of wages for ongoing workers. In these 70 quarters preceding 1984 productivity is taken from the actual data but newly hired wages are computed based on the estimated coefficients. From 1984 onwards we use actual data for newly hired wages. We then perform two exercises. First, we find the wage contract for continuing workers such that the the elasticity of the wage for all workers is matched. Second, for various, exogenously given values of $\phi_{1,\text{newh}}$ and $\phi_{1,\text{stay}}$, we simulate wage paths and estimate $\hat{\eta}_{\text{newh}} = \hat{\phi}_{1,\text{newh}}$, $\hat{\eta}_{\text{allw}}$ as well as the response of the expected present values.

Table 12 shows the results of the simulations for different values of the cyclicality of the wage of new hires $\phi_{1,\text{newh}}$ and the cyclicality of the contract wage $\phi_{1,\text{stay}}$. As expected, a wage contract like (14) drives a wedge between the cyclicality of wages of new hires and all workers, the former responding more to changes in productivity than the latter if $\phi_{1,\text{stay}} < \phi_{1,\text{newh}}$. The measured elasticity of wages of new hires with respect to productivity by construction equals $\phi_{1,\text{newh}}$. The measured elasticity of wages of all workers with respect to productivity increases with the contract elasticity $\phi_{1,\text{stay}}$, but there is a substantial difference. The reason for this difference is that the group of job stayers changes over time: this period's job stayers include last period's new hires. The larger is the difference between the cyclicality of the wage in ongoing matches $\phi_{1,\text{stay}}$ and at the start of a job $\phi_{1,\text{newh}}$, the larger is the gap between the measured elasticity for all workers and the contract elasticity for job stayers.[30] The implied wage contract that matches our estimates in Table 4 has an average wage growth with tenure of 2% per year, $\phi_{0,\text{stay}} = 0.02$, and an elasticity of the wage with respect to aggregate productivity of $\phi_{1,\text{stay}} = 0.25$.

D2. Present value of wages and productivity

For the simulated wage contracts, which we calibrated to be consistent with our estimates for the response of the average wage of new hires and all workers to changes in productivity, we have all the information necessary to calculate the expected net present value of wages at the start of a match.[31] Since we assumed a stochastic process for productivity, we can calculate the expected net present value of productivity as well.

In order to later compute expectation and variance of wages and productivity it is useful to

[30] If $\phi_{1,\text{newh}}$ is smaller or not much larger than $\phi_{1,\text{stay}}$ then the estimated elasticity for all workers can be smaller than the contract elasticity for job stayers because of the exogenous wage growth in wages in continuing job relationships. This pattern disappears when we set $\phi_{0,\text{stay}} = 0$.

[31] The only additional piece of information we need is a discount rate, for which we use the three-month T-bill rate or the bank prime loan rate (FRED series TB3MS or MPRIME).

have access to a moving-average type representation for productivity.

$$\gamma_1 = \frac{1}{\psi_1 - 1}(\log y_t - \log y_{t-1}) \tag{20}$$

$$\gamma_2 = \log y_t - \psi_1\gamma_1 = \frac{1}{\psi_1 - 1}\log y_t + \frac{\psi_1}{\psi_1 - 1}\log y_{t-1} \tag{21}$$

$$\log y_{t+k} = \gamma_2 + \gamma_1\psi_1^{k+1} + \frac{1}{\psi_1 - 1}\left\{\sum_{s=1}^{k}(\psi_0 + \upsilon_t)\left(1 - \psi_1^{k+1+s}\right)\right\} \tag{22}$$

Notice that we are considering an *exploding* series here. Therefore it is not possible to simply take the MA-representation of an AR(2). Based on the distributional assumptions given in (18) we can now compute conditional expectations and variances:

$$\begin{aligned} E(\log y_{t+k}|\log y_t) &= \gamma_2 + \gamma_1\psi_1^{k+1} + \frac{\psi_0}{\psi_1 - 1}\sum_{s=1}^{k}\left(1 - \psi_1^{k+1+s}\right) \\ &= \gamma_2 + \gamma_1\psi_1^{k+1} + \frac{k(1-\psi_1) - \psi_1(1-\psi_1^k)}{(1-\psi_1)^2}\psi_0 \\ &= \log y_t + \frac{\psi_1\left(1-\psi_1^k\right)}{1-\psi_1}(\log y_t - \log y_{t-1}) \\ &\quad + \frac{k(1-\psi_1) - \psi_1(1-\psi_1^k)}{(1-\psi_1)^2}\psi_0 \end{aligned} \tag{23}$$

$$\begin{aligned} V(\log y_{t+k}|\log y_t) &= \frac{\sigma_\upsilon^2}{(1-\psi_1)^2}\sum_{s=1}^{k}\left(1 - \psi_1^{k+1+s}\right)^2 \\ &= \left(\frac{1}{(-1+\psi_1)^2}k^2 + \frac{\psi_1\left(-1+\psi_1^k\right)\left(-2+\psi_1(-1+\psi_1^k)\right)}{(-1+\psi_1)^2(-1+\psi_1^2)}k\right)\sigma_\upsilon^2 \end{aligned} \tag{24}$$

Not surprisingly, the conditional variance is growing at the rate of k^2, so when computing the present values we need to hope that our discounting and separations probabilities will counter this effect. Any potential divergence is going to come from this conditional variance term.

For wages, note that we can write:

$$\log w_{t+k}^k = \log w_t^0 + k\phi_{0,\text{stay}} + \phi_{1,\text{stay}}(\log y_{t+k} - \log y_t) + \sum_{s=1}^{k}\nu_{t+s} \tag{25}$$

$$E(\log w_{t+k}^k|\log w_t^0) = k\phi_{0,\text{stay}} + \phi_{1,\text{stay}}(E(\log y_{t+k}|\log y_t) - \log y_t) \tag{26}$$

$$V(\log w_{t+k}^a|\log w_t^0) = \phi_{1,\text{stay}}^2 V(\log y_{t+k}|\log y_t) + k\sigma^2 \tag{27}$$

The final step is the computation of expected present values. First consider productivity:

$$
\begin{aligned}
E_t PV(y) &= e^{\log y_t} + \sum_{k=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^k E_t e^{\log y_{t+k}} \\
&= e^{\log y_t} + \sum_{k=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^k e^{\{E_t(\log y_{t+k}|\log y_t)+0.5V_t(\log y_{t+k}|\log y_t)\}} \qquad (28)
\end{aligned}
$$

To compute the expected present value of wages, we can use the assumption that wages are a linear function of contemporaneous productivity.

$$
\begin{aligned}
E_t PV(w) &= w_t^0 + \sum_{k=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^k E_t e^{\log w_t^0 + k\phi_{0,\text{stay}} + \phi_{1,\text{stay}}(\log y_{t+k} - \log y_t) + \sum_{s=1}^{k} \nu_{t+s}} \\
&= w_t^0 \left(1 + \sum_{k=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^k E_t e^{k\phi_{0,\text{stay}} + \phi_{1,\text{stay}}(\log y_{t+k} - \log y_t) + \sum_{s=1}^{k} \nu_{t+s}}\right) \\
&= w_t^0 \left(1 + \sum_{k=1}^{\infty} \left(\frac{1-\delta}{1+r}\right)^k \exp\left(\begin{array}{l} k\left(\phi_{0,\text{stay}} + 0.5\sigma^2\right) \\ \quad +\phi_{1,\text{stay}}\left(E_t \log y_{t+k} - \log y_t\right) \\ \quad +0.5\phi_{1,\text{stay}}^2 V_t \log y_{t+k} \end{array}\right)\right) \qquad (29)
\end{aligned}
$$

Table 12 shows the results of these calculations for our simulated wage contracts. The third number in each cell in the table reports the response of the permanent wage with respect to permanent productivity, $d \log \bar{w}_t / d \log \bar{y}_t$, for a given set of parameters of the wage contract. By the argument in Section 4., this elasticity is a good summary statistic for the cyclicality of the wage contract that affects labor market volatility. It is clear from the table that the elasticity of the permanent wage with respect to permanent productivity is always very close to the elasticity of the wage of new hires with respect to current productivity, suggesting that the latter is a good observable proxy for the cyclicality of the wage contract. For the contract that is consistent with our estimates, we find an elasticity of the permanent wage with respect to permanent productivity of 0.8.

D3. Response of job creation to productivity

We now turn to the question how much the observed type of wage wage contracts amplifies the effect of productivity shocks on job creation. As a benchmark, first consider the case of fully flexible wages that respond one-for-one to changes in productivity. In the calibration of Mortensen and Nagypal (2007), all vacancy posting costs are per-period costs, $c(q_t) = c/q_t \Rightarrow -q_t c'(q_t)/c(q_t) = 1$, and the elasticity of the matching function with respect to unemployment equals $\mu = 0.6$. Thus, with $d \log \bar{w}_t / d \log \bar{y}_t = 1$, the job creation condition (5) predicts an elasticity of the job finding rate with respect to permanent productivity of $d \log p_t / d \log \bar{y}_t = (1-\mu)/\mu = 0.7$, see equation (7). Our estimates for the process for productivity imply that productivity is very close to a random walk, $d \log \bar{y}_t / d \log y_t = 1.04$, so that the elasticity of the job finding rate with respect to current productivity is roughly equal to the elasticity with respect to permanent productivity, $d\log p_t/d\log y_t = d\log p_t/d\log \bar{y}_t = 0.7$. In the data, a regression of the log of the job finding rate on the log of productivity gives a coefficient of 7.6 (Mortensen and Nagypal 2007). Thus, the model underpredicts the volatility of the job finding rate in response to technology shocks by a factor 10.

According to our estimates, the elasticity of the permanent wage with respect to permanent productivity equals 0.8. In order to assess how much this amount of wage rigidity amplifies fluctuations in job creation using equation (7), we need a value for the ratio of wages over productivity $\bar{w}_t/\bar{y}_t$. Since no direct calibration target is available, we close the model in order to calibrate this ratio. For simplicity, we solve the model in steady state, so that job creation equation (11) and Bellman equation (12) for a filled job J simplify to

$$\begin{aligned} c(q) &= J \qquad (30)\\ (1+r)J &= y - w + (1-\delta)J \qquad (31)\end{aligned}$$

In order to solve for the wage, we need to complement this labor demand side of the model with Bellman equations for an employed worker W and an unemployed worker U.

$$\begin{aligned} (1+r)W &= w + (1-\delta)W + \delta U \qquad (32)\\ (1+r)U &= z + pW + (1-p)U \qquad (33)\end{aligned}$$

We assume, without loss of generality, that in steady state workers receive a fraction φ of the surplus generate by a match, so that we get the following surplus sharing rule.

$$\frac{W-U}{\varphi} = \frac{J}{1-\varphi} \qquad (34)$$

The steady state wage can be calculated from the Bellman equation for a filled job J, using the surplus sharing rule.

$$w = y - (r+\delta)J = y - (r+\delta)(1-\varphi)S \qquad (35)$$

where total match surplus $S = W - U + J$ is given by

$$S = \frac{y-z}{r+\delta+\varphi p} \qquad (36)$$

Substituting and simplifying, we get an expression for the wage as a weighted average of productivity y and the flow value of unemployment z,

$$w = \Phi y + (1-\Phi)z \qquad (37)$$

where the weight Φ can be written in terms of direct calibration targets only

$$1 - \Phi = \frac{(r+\delta)(1-\varphi)}{r+\delta+\varphi p} \simeq \frac{1-\varphi}{\varphi}\frac{r+\delta}{p} \tag{38}$$

The approximation is valid for $p >> r + \delta$. Notice that we have left the endogenous variable p in this expression, which is why we did not use steady state job creation equation (30) and why the wage does not depend on vacancy posting costs. We do this, because the average level of p is directly observable and typically used as a calibration target.

Substituting this expression into the expression for the elasticity of the job finding rate with respect to productivity (7) in the main text, evaluated in steady state, we get,

$$\frac{d\log p_t}{d\log \bar{y}_t} \simeq -\frac{c(q)}{qc'(q)}\frac{1-\mu}{\mu}\left[1 + \frac{\varphi}{1-\varphi}\frac{p}{r+\delta}\frac{y}{y-z}\left(1 - \frac{d\log \bar{w}_t}{d\log \bar{y}_t}\right)\right] \tag{39}$$

where the approximation is valid for $\varphi p >> r + \delta$.

Since the job finding rate p and the separation rate δ are observable and their average levels are typically used as calibration targets, there is no controversy about the ratio $p/(r+\delta)$ in steady state, which equals 12 in the US data.[32] This high ratio, which corresponds to a relatively low unemployment rate, strongly amplifies the effect of small surplus $y/(y-z)$ as well as wage rigidity $1 - d\log \bar{w}_t / d\log \bar{y}_t$. Assuming per-period vacancy posting costs as in the standard model, $-q_t c'(q_t)/c(q_t) = 1$, using a value for $\mu = 0.6$ as in Mortensen and Nagypal (2007), assuming the Hosios condition is satisfied in steady state so that $\varphi = \mu$ and using our estimate for wage rigidity, $d\log \bar{w}_t / d\log \bar{y}_t = 0.8$, we find that $d\log p_t / d\log \bar{y}_t = 5$ for $z/y = 0.4$ as in Shimer (2005), $d\log p_t / d\log \bar{y}_t = 9$ for $z/y = 0.7$ as in Mortensen and Nagypal (2007) and $d\log p_t / d\log \bar{y}_t = 49$ for $z/y = 0.95$ as in Hagedorn and Manovskii (2008). Thus, given the observed response of wages to changes in productivity, the model can comfortably match the observed regression coefficient of the job finding rate on productivity of 7.6 for reasonable values of the replacement ratio.

Equation (39) sheds light on the various solutions that have been proposed for the unemployment volatility puzzle. We find that on the one hand there is evidence for very little wage rigidity in the data, but on the other hand very little wage rigidity is needed to match the volatility of job creation. The intuition for this conclusion is that the wage as a fraction of productivity $\bar{w}/\bar{y}$ is very close to one so that even a small amount of wage rigidity generates a large amount of amplification, see equation (7). Mortensen and Nagypal (2007) make a similar argument, although they did not have any direct evidence on the amount of wage rigidity in the data. The observed amount of wage rigidity is consistent with a modest degree of wage stickiness e.g. as in Hall and Milgrom (2008), but can also be replicated by models with flexible wage setting, for example by reducing workers' bargaining power as in Hagedorn and Manovskii (2008). By assuming less countercyclicality in vacancy posting costs, as Pissarides (2009) does, it is even possible to match

[32] There may be disagreement about the average levels of p and δ, which depend on the time period used and the aggregation method to go from monthly data to other frequencies, but not about their ratio.

the volatility of job creation without any wage rigidity, i.e. with $d \log \bar{w}_t / d \log \bar{y}_t = 1$. In this case, $-q_t c'(q_t)/c(q_t) = k/(qK + k)$ so by making the per-period component of vacancy posting costs k arbitrarily small relative to the fixed component K, one can amplify the volatility of job creation to arbitrarily high levels. The contribution of this paper is to provide an estimate of the response of the expected net present value of wages to changes in productivity, which can be used as a calibration target and rules out models with very sticky wage setting.

E Additional tables and figures

Table 10: Robustness to alternative estimators

	Wage per hour		Earnings per person	
WLS	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.25	0.79	0.36	0.86
Std. error	0.14	0.40	0.17	0.50
Median	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.13	0.89	0.15	0.56
Std. error	0.20	0.45	0.24	0.70
Median, WLS	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.11	0.89	-0.05	0.57
Std. error	0.24	0.49	0.22	0.72

Elasticities are estimated using the two-step method described in the text. WLS weights the second step regression by the inverse of the variance of the first step estimates. Median uses the median wages instead of mean wages by quarter.

Table 11: Robustness to alternative sample selection criteria

	Wage per hour		Earnings per person	
Including supervisory workers	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.10	0.57	0.39	0.70
Std. error	0.13	0.40	0.18	0.49
Including public sector	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.06	0.70	0.33	0.57
Std. error	0.12	0.48	0.15	0.54
New hires out of unemployment	All workers	New hires	All workers	New hires
Elasticity wrt productivity	0.24	0.77	0.37	0.69
Std. error	0.14	0.55	0.17	0.70

Elasticities are estimated using the two-step method described in the text. The table compares the results for different compositions of the sample from which the CPS wages are constructed.

Table 12: Simulated long-term wage contracts

$\phi_{1,\text{newh}}$	$\phi_{1,\text{stay}}$				
	0.00	**0.25**	0.50	0.75	1.00
	0.10	0.10	0.10	0.10	0.10
0.10	-0.04	0.18	0.40	0.63	0.85
	0.10	0.11	0.12	0.13	0.14
	0.30	0.30	0.30	0.30	0.30
0.30	-0.02	0.20	0.42	0.64	0.87
	0.29	0.30	0.31	0.32	0.33
	0.50	0.50	0.50	0.50	0.50
0.50	0.00	0.22	0.44	0.66	0.88
	0.48	0.49	0.50	0.51	0.52
	0.80	**0.80**	0.80	0.80	0.80
0.80	0.03	**0.25**	0.47	0.69	0.91
	0.77	**0.78**	0.79	0.80	0.81
	1.00	1.00	1.00	1.00	1.00
1.00	0.05	0.27	0.49	0.71	0.93
	0.97	0.97	0.98	0.99	1.00

The table reports three elasticities from simulated data for individual wages, assuming long-term wage contracts with parameters $\phi_{1,\text{newh}}$ and $\phi_{1,\text{stay}}$ as described in Section **??**. The first two numbers in each cell are the elasticities for the wages of new hires and all workers, estimated from the simulated data using specification (3). Since we repeated the simulations many times and averaged the results, the standard errors of these estimates are negligible. The third number is the elasticity of the expected net present value of wages with respect to the expected net present value of productivity, calculated consistent with the stochastic processes we used for the simulations. Across rows and columns of the table we vary the parameters of the wage contracts. Different rows show results for different values for the cyclicality of the wage at the start of a contract. Different columns correspond to different values for the cyclicality of the wage in an ongoing job relationship.

Figure 2: Characteristics of newly hired workers over time

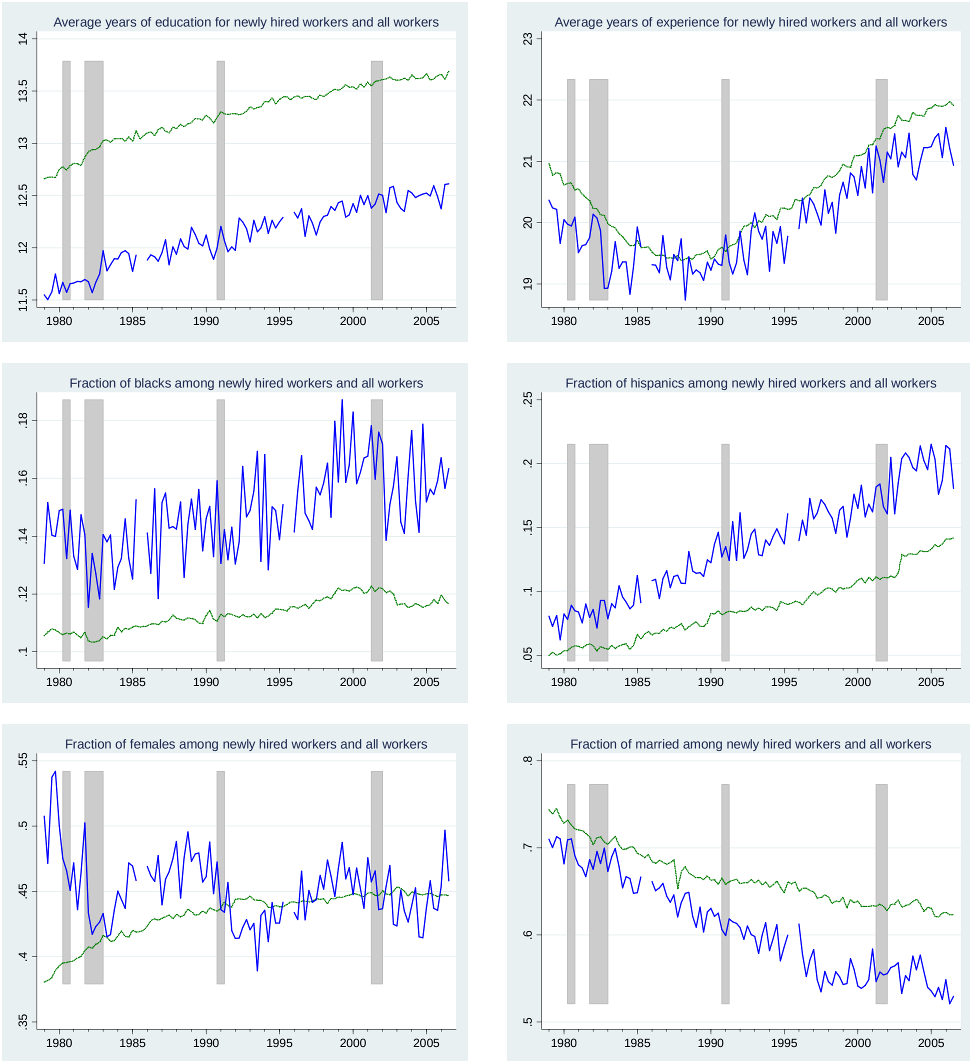

The green dotted line is the average for all workers and the blue solid line for new hires. Education coding changes in 1992. In order not to loose that observation, we regressed the average education level in the sample on a third order polynomial in time and a post 1992 dummy and took the residuals, adding back up the polynomial but not the dummy to correct the resulting level shift. The sample includes all individuals in the CPS who are employed in the private non-farm business sector and are between 25 and 60 years of age (men and women), excluding supervisory workers. New hires are workers that were non-employed at least once within the previous 3 months. The gaps in the graph are quarters when it is not possible to identify newly hired workers, see Online Appendix A. The grey areas indicate NBER recessions.