

IZA DP No. 2006

On-the-Job Search, Productivity Shocks and the Individual Earnings Process

Fabien Postel-Vinay
Hélène Turon

March 2006

On-the-Job Search, Productivity Shocks and the Individual Earnings Process

Fabien Postel-Vinay

*University of Bristol, CREST-INSEE
and IZA Bonn*

Hélène Turon

*University of Bristol
and IZA Bonn*

Discussion Paper No. 2006
March 2006

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
Email: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

On-the-Job Search, Productivity Shocks and the Individual Earnings Process*

Individual labor earnings observed in worker panel data have complex, highly persistent dynamics. We investigate the capacity of a structural job search model with i.i.d. productivity shocks to replicate salient properties of these dynamics, such as the covariance structure of earnings, the evolution of individual earnings mean and variance with the duration of uninterrupted employment, or the distribution of year-to-year earnings changes. Specifically, we show within an otherwise standard job search model how the combined assumptions of on-the-job search and wage renegotiation by mutual consent act as a quantitatively plausible “internal propagation mechanism” of i.i.d. productivity shocks into persistent wage shocks. The model suggests that wage dynamics should be thought of as the outcome of a specific acceptance/rejection scheme of i.i.d. productivity shocks. This offers an alternative to the conventional linear ARMA-type approach to modelling earnings dynamics. Structural estimation of our model on a 12-year panel of highly educated British workers shows that our simple framework produces a dynamic earnings structure which is remarkably consistent with the data.

JEL Classification: J41, J31

Keywords: job search, individual shocks, structural estimation, covariance structure of earnings

Corresponding author:

Hélène Turon
Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
United Kingdom
Email: helene.turon-lacarrieu@bris.ac.uk

* The authors thank Gadi Barlevy, Jeff Campbell, Morris Davis, Jean-Marc Robin and Randy Wright for their inspiring thoughts on this paper. Very useful feedback was also received from conference participants at the 2005 Essex Thanksgiving conference and seminar audiences at Bristol, Paris I, CREST, UCL, Queen Mary, Carnegie Mellon, Minnesota, the Chicago Fed, Toulouse and Oxford. The usual disclaimer applies.

1 Introduction

Individual labor earnings, as they are observed from worker panel data, follow a complex and as yet not fully understood dynamic pattern. While no claim has been made so far to discovery of the “true” earnings process, the numerous existing empirical analyses of this process seem to concur on a twofold conclusion.¹ First, earnings shocks are highly persistent over time. Second, it takes a fairly rich mix of random processes to replicate the intricate autocovariance structure of earnings. Indeed the archetypal ARMA-type decomposition of the individual earnings process features a martingale or a highly persistent Markov component, on top of a fixed-effect and a transitory (MA) shock. While the dynamic properties of individual earnings—notably persistence—are well diagnosed by this reduced-form approach, the economic mechanisms at the root of these properties are still unknown.

In this paper we aim to offer a candidate theoretical foundation of individual earnings dynamics by investigating the capacity of a structural model of job search with simple i.i.d. productivity shocks to capture the main aspects of observed earnings processes. Specifically, we show how the combined assumptions of *on-the-job search* (with search frictions) and *wage renegotiation by mutual consent* can act as a realistic “internal propagation mechanism” of i.i.d. productivity shocks. This combination of assumptions, which we shall motivate momentarily, implies that purely transitory productivity shocks are translated into persistent wage shocks with a covariance structure that we find to be consistent with the data.²

The intuitive mechanism at work is as follows. Consider firms and workers who are matched in pairs, each match facing an idiosyncratic productivity (or “match quality”) shock in every period. Also assume that, through on-the-job search, workers occasionally contact outside firms who then

¹The related literature is literally huge. A somewhat arbitrary selection includes the seminal papers by Lillard and Willis (1978), Lillard and Weiss (1979), MaCurdy (1982) and Abowd and Card (1989), the comparative analyses of recent developments by Baker (1997) or Alvarez, Browning and Ejrnæs (2001). See also Blundell and Preston (1998) for an application to U.K. data, and Meghir and Pistaferri (2004) as an example of a state-of-the-art paper in this field.

²Of course it is not our contention that i.i.d. productivity shocks is by any standards a “better” assumption than any different, more complicated ARMA process. We nonetheless adopt this assumption first for tractability reasons (see footnote 7 below), and second because we view one of our contributions as the identification and analysis of a potential endogenous explanation of income persistence, given the least possible persistent process for what we now refer to as “productivity” and will in fact cover the broader concept of “match quality” in the model.

compete over their services with their current employer. Because of search frictions, worker-firm pairings produce a positive surplus which the wage rate splits into the worker's value and the employer's profit. In this process, the maximum wage that the firm is willing to pay leaves the firm with zero profit and follows productivity shocks. The minimum wage that the worker is willing to receive yields the worker her/his outside option value, which equals the value of unemployment except in periods when the worker raises an outside offer, in which case it equals the value of this offer. As wages can only be renegotiated by mutual consent, neither party can force the other to renegotiate if the latter is not willing to do so. Three distinct situations then arise. First, if the match receives a sufficiently adverse productivity shock to make it unprofitable for the firm to keep employing the worker at her/his current wage, then the firm has a credible threat to fire the worker which it can use to renegotiate the wage downward. Second, whenever the worker receives an outside job offer paying a higher wage than her/his current wage, s/he can credibly threaten to accept it in order to force her/his employer into upward wage renegotiation. This will lead to an efficient separation if the outside offer is greater than the maximum wage the firm is able to pay. Finally, in any other event (i.e. no sufficiently adverse productivity shock and no sufficiently good outside job offer), neither party is in a position to force the other to renegotiate, and the wage remains unchanged.

The wage is only altered when one of those outside-option constraints becomes binding, in which case it is revised up or down by just enough to satisfy whichever constraint is binding. Indeed the pattern of wage dynamics implied by the model just sketched can be summarized graphically as in Figure 1 (which we adapt from MacLeod and Malcomson, 1993): because of search frictions and the rule of mutual consent, i.i.d. productivity shocks are only infrequently translated into wage shocks, hence wage shock persistence.

< **Figure 1 about here.** >

The intuitive idea that renegotiation by mutual consent likely causes some form of “price stickiness” has been around for a while (as studies surveyed by Malcomson, 1997, suggest). Yet, our paper is, inasmuch as we know, the first to formalize it in the context of a structural job search

model and to provide a quantitative analysis of the resulting individual income dynamics.

Our theoretical model can be seen as a version of the matching model of labor market equilibrium, now routinely referred to as the Diamond-Mortensen-Pissarides, or “DMP” model,³ in which employed job search is allowed. Although virtually any wage formation mechanism can be embedded into the DMP model, the typical (and by far dominant) practice is to assume a Nash-like sharing rule, whereby each party receives a given share of the match surplus at all times. Hidden underneath this constant-share feature is the assumption that wages are renegotiated at least every time the match is hit by a productivity shock. This assumption is somewhat arbitrary as in general the occurrence of a shock to match productivity does not give either of the matched partners a credible threat to force the other to renegotiate. As advocated by, e.g., Malcomson (1997, 1999), renegotiation by mutual consent is a more natural assumption, at least for its consistency with a number of legal and/or economic facts.⁴

Besides formalizing the above theoretical mechanism, we show that its quantitative implications are remarkably well borne out by the data: we estimate our structural model on a sample of high-educated British workers taken from the BHPS and provide an in-depth fit analysis of the model. In so doing, we contribute to the growing body of research carrying out structural estimation of various forms of search models, which have so far been essentially geared to the description of cross-sectional wage dispersion. As a consequence, estimation of these models tends to mostly rely on the cross-sectional dimension of the data, leaving aside the question of individual earnings

³From Diamond (1982) and Mortensen and Pissarides (1994). For a complete exposition of the DMP model and many extensions thereof, see Pissarides (2000).

⁴Mutual agreement is indeed a prerequisite to wage renegotiation under English law, which is relevant to the data we use in the latter part of this paper. In the U.S., while the employment-at-will doctrine would in principle leave scope for more responsiveness of wages to productivity shocks, the empirical evidence reviewed in Malcomson (1997, 1999) reveals that wage changes occur much less frequently than would be consistent with a strict application of the employment-at-will rule, suggesting that mutual consent, although not an explicit legal provision in the U.S., may nonetheless be common practice.

On the theory side, Mortensen and Pissarides (2003, footnote 4) recognize, without actually using it for their purposes, that the assumption of renegotiation by mutual consent “may well generate more realistic wage dynamics”. Fella (2004) does in turn implement this type of negotiation within the standard DMP model (without on-the-job search). However, ignoring on-the-job search leads to the counterfactual prediction that wage profiles unambiguously (stochastically) decline over the job spell. On the other hand, existing versions of the DMP model with on-the-job search (Pissarides, 2000 chapter 4, Shimer, 2006), mostly shut down between-employer competition by assuming that the worker’s outside option is always unemployment, even when s/he winds up with an outside job offer. This, combined with the assumption that wages are renegotiated every time a shock hits the match, implies that individual wages fluctuate along with match productivity.

dynamics.⁵ Yet search models are inherently dynamic and have strong predictions about the process followed by individual wages over time. What little attention has been paid to those predictions has led to the conclusion that, in the absence of individual-level shocks, job search models fail to accommodate the observed downward wage flexibility.⁶ By contrast, we consider individual-level shocks and exploit as much as possible the observed dynamics of individual labor income, as is allowed by the longitudinal dimension of our panel and the dynamic predictions of our model.

But most importantly, we offer a structural counterpart to the aforementioned “reduced-form” literature on individual income processes. A typical application in this literature is to use a particular permanent/transitory decomposition of incomes to test the permanent-income/life-cycle hypothesis. Now surely, as e.g. Baker (1997) notes, such tests have “an obvious dependence” on the specific decomposition of income. Other frequent fields of application include the study of wage rigidity or that of wage growth over the working life. In both cases, the relevant policy implications vary quite a lot from one possible underlying theoretical framework to another. These examples are meant to illustrate the importance of understanding the economic forces governing individual wage dynamics from a dual theoretical and quantitative standpoint. This paper suggests that the combination of on-the-job search and renegotiation by mutual agreement is a promising candidate explanation of the widely documented persistence of earnings shocks. In particular, our structural approach highlights the interplay between job mobility and earnings dynamics: the model predicts that the individual probabilities of transitions between labor market states condition the individual earnings process in a way that is consistent with the data. More generally, our theory suggests that the income process should be thought of as following a particular acceptance/rejection scheme of underlying i.i.d. productivity shocks, of which labor market transition rates are a key determinant.

The rest of the paper is organized as follows. In the next section we pose the theoretical model. In Section 3 we go on to derive the model’s solution in connection with our estimation procedure, which is presented in Section 4, together with the data. Section 5 contains estimation results and an analysis of the model’s performance at replicating some features of the earnings data. Finally,

⁵See Eckstein and Van den Berg (2005) or Postel-Vinay and Robin (2006) for overviews.

⁶See Postel-Vinay and Robin (2006) for a review of this argument.

we conclude and discuss a number of potentially interesting extensions in Section 6.

2 Theory

2.1 The environment

Basics. We consider a labor market where a unit mass of workers face a continuum of identical firms producing a multi-purpose good which they sell in a perfectly competitive market. Time is discrete and the economy is at a steady state. Workers can either be unemployed or matched with a firm. Firms operate constant-return technologies and are modeled as a collection of job slots which can either be vacant and looking for a worker, or occupied and producing.

The output flow y_t of a firm-worker match in period t is defined as:

$$y_t = p \cdot \varepsilon_t. \tag{1}$$

It is the product of a worker fixed-effect p and a transitory *match- and period-specific* shock ε_t . We should emphasize that because the transitory shock is match-specific, a realization of ε is *not* carried over from one firm to the other in case the worker changes firms.

The population distribution of (log) worker fixed effects $\ln p$ is denoted as $H(\cdot)$. Identification requires normalization of one of the components of (1). We choose to normalize the mean value of $\ln p$ at zero: $E_H(\ln p) = 0$.

When a worker and a firm meet, the idiosyncratic component of (potential) match productivity is drawn from a distribution $M(\cdot)$ with support $[\varepsilon_{\min}, \varepsilon_{\max}]$. Every ongoing firm-worker match draws a new value of ε_t at the beginning of each period t from that same distribution $M(\cdot)$. Depending on the realized value of ε_t , the match can go on under the same wage contract or under a renegotiated contract. The precise cutoff values of the transitory shock under which a contract is renegotiated are determined below. Finally, transitory shocks ε_t are uncorrelated over time or across matches.⁷

⁷These assumptions have a twofold content: first, productivity shocks occur in every period and are the single most frequent random event in this environment. Second, these shocks are serially uncorrelated. These two—admittedly disputable—assumptions greatly simplify our analytical characterization of wage dynamics (see subsection 3.3). While it would be feasible to relax the first assumption (though at a great cost in added complexity and with little benefit in terms of new qualitative insights brought for our purposes), introducing serially correlated shocks would render

Future discounting. In general we assume that workers and firms are infinitely lived, forward-looking, risk-neutral and have a common exogenous per-period discount factor of β . Yet in the main text we shall work with a simple version of the model in which $\beta = 0$, i.e. workers and firms are completely myopic and only care about current period flow revenues. The full model accounting for forward-looking behavior (i.e. the model with $\beta > 0$) is presented in detail in Appendix A.

At this point we would like to insist that the assumption $\beta = 0$ is there merely for simplicity of exposition in the main text. It is made without important loss of substance or intuition. Indeed the form of the empirical model derived from our theory is independent of the value of β , so that the analysis from section 3 onward is applicable irrespective of that value. More specifically, we show in Appendix A that β only affects the way in which productivity shocks translate into wage shocks. As our estimation procedure is based on wage data only (and on the assumption that productivity shocks are i.i.d.), knowledge of β is unnecessary for the identification of the parameters of our wage process. It would only matter if we wanted to recover the distribution of productivity shocks implied by our model and by the earnings dynamics observed in our data.

Unemployment income. In any given period, an unemployed worker with permanent productivity component p (henceforth a “type- p worker”) receives a flow income of $b \cdot p$, $b > 0$. This contains the assumption that unemployment income depends on the permanent individual productivity parameter in the same way (i.e. multiplicatively) as productivity in a match with a firm. This assumption is inessential—although not quantitatively innocuous—and again is made because it simplifies the formal model somewhat. It also implies that, while unemployment income may differ between workers, it does not fluctuate over time.

the model analytically intractable. The relatively simple, closed-form characterization of wage dynamics presented in subsection 3.3, which we view as an attraction of our approach, would become impossible under serially correlated shocks. We should finally mention that, in an attempt to deal with the lack of persistence of match productivity, we have experimented by estimating our model using different lengths of the unit time period, and found no qualitative difference in our results. Another possible way to increase persistence (while preserving serial independence) of productivity shocks would be to model their occurrence as a jump process similar to Mortensen and Pissarides (1994). Again, this would add great complexity in the formalization of the wage process and would also raise issues of identification of the shock arrival rate in the estimation stage.

Surpluses. Consider a match between a firm and a type- p worker, with current productivity ε . Denote the current wage in this match by ϕ .⁸ Under our current simplifying assumption of infinite future discounting, the worker’s valuation of this match net of the foregone unemployment income is $\phi - b \cdot p$. Further assuming that a vacant job slot is worth 0 to the firm (as naturally results from free entry and exit of vacant jobs on the search market), the firm’s net valuation of that same match equals $p \cdot \varepsilon - \phi$. Accordingly, the total surplus accruing from the match is the sum of those two valuations and simply equals $p \cdot (\varepsilon - b)$. An important thing to note about total match surplus is that it is independent of any wage value.

Job search, match formation and match dissolution. One core assumption in this paper is that the labor market is affected by job search frictions: firms and workers are brought together in pairs through random search. Specifically, any unemployed worker has a per-period probability λ_0 of meeting a firm. We also allow employed workers to raise job offers and assume that they have a per-period probability λ_1 of meeting a potential alternative employer. Note that we only allow workers (in any employment state) to contact at most one firm per period. Moreover, we assume that contacts made in earlier periods cannot be recalled.

Not all firm-worker contacts are a priori conducive to an actual job move: for employed job seekers, the decision of whether or not to quit an ongoing match for a new one involves a comparison of match surpluses which will be carried out in full detail in the next subsection.

Finally, all matches have a common, exogenous breakup probability of δ per period. This assumption calls for the following comments. The formation or continuation of any firm-worker match implies the minimal requirement that total match surplus be *nonnegative*. Strictly abiding by the simple specifications detailed above, total match surplus equals $p \cdot (\varepsilon - b)$. Thus what we are assuming here is simply $\varepsilon_{\min} \geq b$, i.e. we are only truncating the “true” underlying distribution of productivity shocks from below at b .⁹ However an alternative, probably more general view on the

⁸We omit period subscripts t when they are not strictly necessary.

⁹More precisely, assume that there is an underlying latent sampling distribution of *potential* match qualities, say $M_0(\cdot)$, with support $(-\infty, \varepsilon_{\max})$. Then, any draw of a productivity shock above b yields a positive potential match surplus, while any draw falling short of b entails a negative potential match surplus and causes match dissolution. Hence the match destruction rate δ can be seen as equaling $M_0(b)$, while $M(\cdot)$ simply coincides with $M_0(\cdot)$ truncated

essence of the assumption of an exogenous job destruction rate is that transitory shocks to match quality—that potentially cause individual *income* fluctuations—are of a fundamentally different nature than shocks leading to a job loss. While our formal setup clearly takes it that shocks to ε are match-specific, we do not give any specific interpretation of the random event causing match destruction, which can reflect adverse shocks to any combination of the match, the individual, the market or the firm.

2.2 Wage determination

General principles. The central principle we stick to is that renegotiation only occurs by mutual consent in continuing matches. In other words, no firm or worker can force their match partner to revise the terms of the employment relationship against the latter’s interest, unless the former has a credible threat to leave the match.¹⁰

The implications for wage dynamics of renegotiation by mutual consent were analyzed theoretically by MacLeod and Malcomson (1993).¹¹ These implications, of which we gave an informal account in Figure 1 in the introduction, can be summarized as follows. If one party does have a credible threat to dissolve the match, i.e. if the value of her/his outside option exceeds the value s/he gets from the existing relationship, the other party will consent to wage renegotiation up or down to the point where this outside option is matched.¹² The existing match will only survive if the surplus it generates is greater than the sum of surpluses generated by the outside options (i.e. a vacant job and an alternative worker-firm match or an unemployed worker). In case neither party has a credible threat to leave the match, there is no mutual consent to revise the wage and the current terms of employment continue to apply.

In newly created matches, there are no pre-existing terms of the (potential) employment relationship and a start-up wage has to be determined. Here we follow the approach of Postel-Vinay

below at *b*.

¹⁰See Malcomson (1997, 1999) for a motivation of this principle.

¹¹The general framework used by MacLeod and Malcomson (1993) is similar to the one we use in this paper, even though their main focus is on the design of efficient contracts when specific investments undertaken by either contracting party are potentially subject to hold-up from the rival party. Such specific investments are not explicitly present in our model.

¹²That is, when renegotiation occurs, outside options act as bounds on the players’ payoffs. This is known in the bargaining literature as the *outside option principle*. (See e.g. Sutton, 1986, or Binmore, Shaked and Sutton, 1988).

and Robin (2002a,b) and assume that firms make take-it-or-leave-it offers to the workers, so starting wages with new employers give workers the value of their outside option.¹³ The latter equals $b \cdot p$ for a worker hired from unemployment or the maximum value the worker could extract from her/his previous employer if the new match follows a job-to-job quit. Thus in the latter case, as explained in Postel-Vinay and Robin (2002a,b), we effectively let the incumbent and the outside employer Bertrand-compete for the worker's services.

Negotiation baselines. It is useful to introduce at this stage the following convention for the description of all wages. As will shortly become clear, at any time, the wage that the worker receives, ϕ , can be expressed as a the product of her/his type p and a value of match-specific productivity r (which will not necessarily be equal to match-specific productivity in the current match, ε). More precisely, all wages will take the form:

$$\phi = \phi(r, p) = \underbrace{p \cdot b}_{\text{Unemployment value}} + \underbrace{p \cdot (r - b)}_{\text{Surplus of match with productivity } r}, \quad (2)$$

where r gives a measure of the surplus that the worker enjoys over and above the value of being unemployed, $b \cdot p$. Note that $\varepsilon \geq r \geq b$ necessarily holds, otherwise either the firm would earn negative profits and fire the worker or the worker would find it preferable to quit into unemployment. Again for reasons that will become clear shortly, we will term r the worker's *negotiation baseline*.

In the next three paragraphs, we define the negotiation baseline formally.

Starting wages. First consider an unemployed, type- p worker meeting a job-advertising firm. Given the assumptions just discussed and a current match quality of ε , the potential match yields positive surplus, and a starting wage contract must be signed. As mentioned above, we assume that, in a newly created match, the employer extracts all the match rent by offering the worker her/his reservation value. In this simple context of infinite time discounting, this entails a starting wage equal to $b \cdot p$, and a negotiation baseline of b .

¹³That is, in terms of a Nash bargaining approach, we assume that the worker has zero bargaining power in newly formed matches. Extending the model to allow for positive worker bargaining power is of potential quantitative importance (see Dey and Flinn, 2005, and Cahuc, Postel-Vinay and Robin, 2006), yet it complicates the writing of the model somewhat. We leave this extension for later work.

Outside offers. We now examine the situation which arises when an already employed worker with current match productivity $p \cdot \varepsilon$ and current negotiation baseline r meets another potential employer through on-the-job search. We denote the match-specific component of productivity in the outside firm as η .

Consistently with the assumptions listed above, we let the incumbent employer and the “poacher” Bertrand-compete for the worker’s services. The worker extracts the whole surplus from the less productive of the two potential matches, which translates into a new negotiation baseline of $\min\{\varepsilon, \eta\}$. If $\varepsilon < \eta$, the poacher profitably attracts the worker with a wage offer of $p \cdot \varepsilon$ (plus one cent)—an offer that the incumbent employer is unable to match without incurring losses. Alternatively, if $\varepsilon \geq \eta$, then the incumbent can profitably retain the worker by matching the poacher’s maximal wage offer of $p \cdot \eta$. Here as a result of the outside offer, the worker stays in her/his current job but can force wage renegotiation up to her/his new outside option, $p \cdot \eta$. In this latter case, however, renegotiation only takes place if the worker gains from it, i.e. if $\eta \geq r$ (otherwise we assume that the worker always has the option to conceal the outside offer s/he has received from the poacher).

We can summarize the possible outcomes of an outside offer received by the worker as follows:

$$\left\{ \begin{array}{ll} \eta > \varepsilon & \text{Worker quits, mobility wage } p \cdot \varepsilon, \text{ new negotiation baseline } \varepsilon, \\ \varepsilon > \eta > r & \text{Worker stays, renegotiated wage } p \cdot \eta, \text{ new negotiation baseline } \eta, \\ r > \eta & \text{Offer is discarded, nothing changes.} \end{array} \right.$$

Productivity shocks. A last potential cause of wage change is the occurrence of a productivity shock. Again consider a match with productivity $p \cdot \varepsilon$, and current negotiation baseline r , and assume a new transitory shock value of ε' is drawn. One of three situations can arise.¹⁴

A first, simple case is $\varepsilon' \geq \varepsilon$. In this case, the worker would like to capture some of the extra surplus brought by the gain in match productivity through a wage increase. But the worker’s only outside option is to resign and become unemployed, thus achieving a value of $p \cdot b$, equivalent to a negotiation baseline of b . This is never preferable to keeping the existing contract which has a

¹⁴For simplicity of exposition, we describe here the case where no outside offers are raised by the worker. The fact that productivity shocks and outside offers can occur simultaneously will naturally be taken into account in the following sections.

negotiation baseline of $r \geq b$. In other words, the worker cannot force the firm to raise the wage. The match thus goes on with an unchanged wage after such a productivity gain.

In the second case, $\varepsilon > \varepsilon' \geq r$, the match has undergone a loss of productivity and the firm's profit has decreased from $p \cdot (\varepsilon - r)$ to $p \cdot (\varepsilon' - r)$. The firm would thus want to share some of this loss with the worker by lowering the wage. But as long as $\varepsilon' \geq r$, profits remain positive at the current wage $p \cdot r$. At that point the firm's only outside option is to fire the worker, thus ending up with a vacant job worth 0, while going on with the existing contract still gives it a positive profit. Hence the firm cannot force the worker to accept a wage cut, and the match again goes on with an unchanged wage.

The third, more complicated case is when $r > \varepsilon' \geq \varepsilon_{\min}$. In this case, since $\varepsilon' \geq \varepsilon_{\min}$, and since by assumption $\varepsilon_{\min} \geq b$, the match is still viable, meaning that a mutually beneficial contract exists. However, keeping the existing wage $p \cdot r$ would imply a negative profit flow of $p \cdot (\varepsilon' - r)$. Here the firm is better off firing the worker than maintaining the match under the existing contract, and so has a credible threat which it can use to force the worker into renegotiation. Our assumed wage-setting rules then imply a wage cut down to the point where the firm enjoys in the continuing match the same value as in its outside option, here equal to zero with a vacant job. This leaves the worker with a wage value of $p \cdot \varepsilon'$ in the continuing match. Her/his negotiation baseline has thus been updated to ε' .

2.3 Within-period timing of events

The last thing that requires further specification before we can solve the model is the sequence of random events affecting firm-worker matches within each period. We simply assume that all of these random events are realized simultaneously at the beginning of each period. The list of such events is the following: match destruction shocks (with probability δ , any given match is dissolved), firm-worker contacts (any unemployed worker meets a potential employer with probability λ_0 , and any employed worker meets a potential alternative employer with probability λ_1), and draws of transitory match productivity shocks ε in incumbent and potential matches. We assume that job destruction shocks and outside offers cannot occur simultaneously, so that the probability that

neither occurs is $1 - \lambda_1 - \delta$. Then wage contracts are negotiated and signed for the period, wages are paid and production takes place.

3 Model solution and econometric inference

Our aim is to estimate the model parameters with a standard panel of individual data on income and labor force transitions (in this application a sub-sample of the BHPS). In this section, we thus derive some of the model's implications that are potentially useful for econometric inference.

3.1 Worker turnover

First looking at job-to-job mobility, we have $\Pr\{\text{job-to-job move} \mid \varepsilon\} = \lambda_1 \overline{M}(\varepsilon)$, implying that the unconditional probability of a job-to-job mobility equals $\lambda_1/2$.¹⁵ Next turning to transitions in and out of employment, the probability of observing a worker moving from employment into unemployment is δ , independently of the worker's type p or the particular value of ε in the worker's initial match. In the opposite direction, the unemployment exit rate is λ_0 for all workers. Incidentally, this implies a steady-state unemployment rate of $u = \frac{\delta}{\delta + \lambda_0}$, obtained from the flow-balance condition ensuring the constancy of the unemployment rate: $\lambda_0 u = \delta(1 - u)$. This latter condition will be used at various points below.

Most importantly, we see that all the transition probabilities can be retrieved by maximization of the likelihood of observed job and unemployment spell durations.¹⁶

3.2 Wage distributions

The model laid out in section 2 implies that all wages have a multiplicative form $\phi(r, p) = p \cdot r$. It thus predicts that log-wages are additively separable into a worker fixed effect $\ln p$ and a transitory/persistent match-specific component $v = \ln r$, the dynamics of which will be characterized in the next subsection.

¹⁵Throughout this paper, a bar over a cdf will be used to denote the survivor function.

¹⁶The hazard rates of job destruction, job finding and job-to-job move are thus constant with respect to spell duration. This counterfactual prediction could be improved upon by introducing some worker heterogeneity in these transition probabilities. Another counterfactual prediction of our simple model concerning job transitions is that the instantaneous job separation rate is independent of the current wage paid in the job. Again, heterogeneity in the λ_1 's (in a form that would be correlated with productive heterogeneity in the p 's) would potentially correct this problem. We discuss these possible extensions in Section 6.

Focusing on steady-state cross-sectional distributions for now, we first notice that (thanks essentially to the proportionality of all income flows to p) the transitory component of wages v is independent of p in a cross-section of employed workers. Conditional on p , log-wages are thus distributed as v and we now seek to determine the steady-state population distribution of v , $G(v)$. Let us consider flows in and out of the stock $(1-u)G(v)$ of employed workers with a (log) negotiation baseline less than $v = \ln r$. Workers exit this pool either if their match has been dissolved (probability δ) or if their new negotiation baseline is greater than r (probability $\lambda_1 \bar{M}(r)^2$).¹⁷ Two flows of workers enter this pool: $\lambda_0 u$ previously unemployed workers, and employed workers with a previous negotiation baseline greater than r with probability $(1-\delta)M(r)$.¹⁸

We can thus now write the balance of flows in and out of the stock $(1-u)G(v)$. Denoting as $F(\cdot)$ the sampling distribution of log-productivity shocks, so that $F(\ln r) = M(r)$, we have:

$$\begin{aligned} (1-u)G(v) &= \left(1 - \delta - \lambda_1 \bar{F}(v)^2\right) \cdot (1-u)G(v) + (1-\delta)F(v) \cdot (1-u)\bar{G}(v) + \lambda_0 \cdot u \\ \iff G(v) &= \frac{1 - (1-\delta)\bar{F}(v)}{1 - (1-\delta)\bar{F}(v) + \lambda_1 \bar{F}(v)^2}. \end{aligned} \quad (3)$$

Note the existence of a mass at $v_0 = \ln b$, $G(v_0) = \frac{\delta}{\delta + \lambda_1}$ due to the unemployed workers all being hired at the minimum negotiation baseline v_0 . More precisely, we can decompose $G(\cdot)$ as:

$$G(v) = \frac{\delta}{\delta + \lambda_1} \cdot \mathbf{1}_{\{v \geq v_0\}} + \left[\frac{1 - (1-\delta)\bar{F}(v)}{1 - (1-\delta)\bar{F}(v) + \lambda_1 \bar{F}(v)^2} - \frac{\delta}{\delta + \lambda_1} \right] \cdot \mathbf{1}_{\{v \geq v_{\min}\}}, \quad (4)$$

where $v_{\min} = \ln \varepsilon_{\min}$ is the lower support of $F(\cdot)$. Also, as expected, $G(\cdot)$ is identically equal to 1 in the absence of on-the-job search (i.e. if $\lambda_1 = 0$).¹⁹

¹⁷As observing an increase in the negotiation baseline requires the worker to receive an outside offer and the productivities of both the poaching and the incumbent firms to be above the current negotiation baseline.

¹⁸As, conditional on remaining employed, workers will have a new negotiation baseline lower than r if the idiosyncratic shock of their incumbent firm is drawn lower than r , whether they raise an outside offer or not.

¹⁹Incidentally we may emphasize that, contrary to standard job search models without idiosyncratic productivity shocks, our model features a non-degenerate equilibrium wage distribution even if one assumes away any risk of unemployment (i.e. if $\delta = 0$). Absent productivity shocks, continuously employed workers would gradually climb up the wage ladder as they receive outside job offers (at a speed that depends on the particular assumptions on wage determination. See Burdett and Mortensen (1998) for the canonical wage-posting model, Postel-Vinay and Robin (2002a) for an offer-matching model closer to the one of this paper, Burdett and Coles (2003) for a model of explicit wage-tenure contracts, and finally Mortensen (2003) for an overview.) Job loss then acts as a “reset button” for this process of wage progression, as workers who have experienced a spell of unemployment essentially have to start over at the bottom of the wage ladder (e.g. in our model, they start over with a negotiation baseline of v_0). In the standard model, without this reset button, all workers would end up at the top of the wage distribution and only one wage (equal to marginal productivity) would be observed in the long-run equilibrium of a homogeneous labor market. Here we see that, by causing occasional wage cuts, productivity shocks prevent this from happening.

3.3 Wage dynamics

Wage dynamics are driven by the combination of two distinct forces: job offers and idiosyncratic shocks to match productivity. Given knowledge of the process governing the arrival of job offers (i.e. given knowledge of the arrival rate λ_1 , which we saw is identified from job spell durations and job transitions), observed individual wage dynamics thus convey information about the distribution of match productivity shocks.

Dynamics over one period. At any period t , an employed worker earns a wage ϕ_t such that $\ln \phi_t = \ln p + v_t$, and we are left to analyze the dynamics of $v_t = \ln r_t$, the log of the worker's current negotiation baseline r_t .

When period $t+1$ begins, with probability δ the match is hit by a dissolution shock. In this case the worker becomes unemployed and her/his income flow becomes equal to $p \cdot b$. With probability $1 - \delta$, the worker stays employed and her/his continuation wage depends on the new value of her/his negotiation baseline r' . We thus now examine the value of r' in the various possible cases.

If the match continues, but the worker fails to find any outside job opportunity (probability $1 - \delta - \lambda_1$), the only source of randomness is the realization of ε_{t+1} . Our various assumptions concerning the wage setting rules then imply the following. If $\varepsilon_{\min} \leq \varepsilon_{t+1} < r_t$, then the match is maintained but under a renegotiated contract and the new negotiation baseline is ε_{t+1} . Otherwise, if $\varepsilon_{t+1} \geq r_t$, then none of the parties can force the other to renegotiate and the match goes on under an unchanged contract, leaving both negotiation baseline and wage unchanged.

Next consider the situation in which the match continues and the worker manages to contact a poacher (probability λ_1). The idiosyncratic productivity component η_{t+1} of a potential match with the poacher is drawn at random from $F(\cdot)$. We scan over all possible values of the shocks ε_{t+1} and η_{t+1} and see what happens in each case.

First, if $\varepsilon_{\min} \leq \varepsilon_{t+1} < r_t$ then the worker has a choice between playing off the two firms against each other or simply discarding the poacher's offer. The former option will yield the worker a new negotiation baseline of $\min\{\varepsilon_{t+1}, \eta_{t+1}\}$, and the latter a negotiation baseline of ε_{t+1} . One thus sees

that the worker's optimal choice yields a continuation value of the negotiation baseline of ε_{t+1} .²⁰

Second, if $\varepsilon_{t+1} \geq r_t$, then playing off the two employers against each other again yields a new negotiation baseline of $\min\{\varepsilon_{t+1}, \eta_{t+1}\}$, while ignoring the poacher's offer amounts to continuing a relationship with the incumbent employer under unchanged terms, thus keeping a negotiation baseline at r_t and consequently an unchanged wage. It follows that the worker's optimal choice yields a continuation renegotiation baseline equal to $\max\langle r_t, \min\{\eta_{t+1}, \varepsilon_{t+1}\} \rangle$.²¹

Summarizing the above, the conditional distribution of the continuing (log) negotiation baseline $v_{t+1} \mid v_t$ is as follows:

$$v_{t+1} \mid v_t = \begin{cases} v_t & \text{with probability } (1 - \delta) \bar{F}(v_t) - \lambda_1 \bar{F}(v_t)^2, \\ v' < v_t & \text{with density } (1 - \delta) f(v'), \\ v' > v_t & \text{with density } 2\lambda_1 f(v') \bar{F}(v'), \end{cases} \quad (5)$$

whereas with probability δ the worker becomes unemployed and v_{t+1} is irrelevant.

Conditional on individual fixed-effects p , we thus predict that wages follow a first-order, nonlinear Markovian process based on a specific acceptance/rejection scheme of i.i.d. wage innovations.²² We also predict that the rates of transition between labor market states (δ and λ_1) are key determinants of the individual earnings process. This strong prediction of our structural model highlights the interplay between job mobility and income dynamics: job mobility reflects the intensity of labor market competition between employers (as measured by the frequency at which employed workers raise outside job offers), which in turn conditions the observed (dynamic) behavior of wages. If validated empirically, that prediction may help with the interpretation of observed wage dynamics.

The empirical properties of the process in (5)—and its differences with the conventional linear

²⁰The optimal choice is to let the firms compete whenever $\eta_{t+1} > \varepsilon_{t+1}$. The outcome of the Bertrand game thus triggered is that the worker joins the poaching firm with a negotiation baseline of ε_{t+1} .

²¹The optimal choice is to let the firms compete whenever $\eta_{t+1} > r_t$. The outcome of the Bertrand game is then that the worker joins the poaching firm if $\eta_{t+1} > \varepsilon_{t+1}$, and stays with her/his incumbent employer, with a wage raise, if $\varepsilon_{t+1} \geq \eta_{t+1} > r_t$.

²²Incidentally, our specific model is not the only one suggesting that this type of acceptance/rejection scheme is the right way to think about wage dynamics. The process in (5) is indeed formally reminiscent of predictions obtained by Harris and Holmström (1982) and Thomas and Worrall (1988) in models of self-enforcing wage contracts designed to allocate risk between a risk-neutral employer and a risk-averse employee faced with uncertainty about match productivity and/or market opportunities. (Much of the related theoretical literature on labor market contracts indeed emphasizes risk sharing as the main driving force behind individual income dynamics—see Malcomson, 1999, for a comprehensive review.)

ARMA specification—will be analyzed in a later section. For the time being, we derive the following moment which will be useful for estimation. Integration of (5) implies that, conditional on employment at two consecutive dates t and $t + 1$:

$$\begin{aligned} E(v_{t+1} | v_t, \text{employment at } t, t + 1) &= E_F(v) - \int_{v_t}^{v_{\max}} \left(\bar{F}(x) - \frac{\lambda_1}{1 - \delta} \bar{F}(x)^2 \right) dx \\ &= v_t + \frac{\lambda_1}{1 - \delta} \int_{v_t}^{v_{\max}} \bar{F}(x)^2 dx - \int_{v_{\min}}^{v_t} F(x) dx. \end{aligned} \quad (6)$$

This latter equation shows that conditional expected wage growth—i.e. $E(v_{t+1} - v_t | v_t, \text{employment at } t, t + 1)$ —is the sum of a positive term reflecting the impact of outside job offers causing wage increases, and a negative term coming from adverse productivity shocks causing downward wage renegotiation. As intuition suggests, the former dominates among workers with a relatively low current negotiation baseline v_t (which translates into a relatively low wage conditional on their type p), while the latter dominates for workers with a high current negotiation baseline (which has little chance to be exceeded by the minimum of a pair of random draws from $F(\cdot)$).²³

Dynamics over s periods. We now consider the cross-sectional distribution of v_{t+s} conditional on employment at dates $t, \dots, t + s$, i.e. the distribution of negotiation baselines conditional on at least s periods of continuous employment. Designating the cdf of this distribution by $G_s(\cdot)$, we show the following in Appendix B:

$$G_s(v) = \left[1 - \left(\bar{F}(v) - \frac{\lambda_1}{1 - \delta} \bar{F}(v)^2 \right)^s \right] \cdot G_\infty(v) + \left(\bar{F}(v) - \frac{\lambda_1}{1 - \delta} \bar{F}(v)^2 \right)^s \cdot G(v), \quad (7)$$

where

$$G_\infty(v) = \frac{F(v)}{1 - \bar{F}(v) + \frac{\lambda_1}{1 - \delta} \bar{F}(v)^2}.$$

Hence as one conditions on more periods of continuous employment, the cross-sectional distribution of negotiation baselines gradually shifts from $G(v)$ to $G_\infty(v)$. An interesting property of this shift (see Appendix B) is that it features a monotonically increasing mean, i.e. $E_{G_s}(v)$ increases with s . Hence, from a cross-section perspective, our model predicts positive returns to continuous

²³Note that, beyond means, higher-order moments of the conditional distribution of $v_{t+1} | v_t$ are functions of v_t . This leaves scope for ARCH-type effects, as were detected in US data by Meghir and Pistaferri (2004).

employment in that the mean negotiation baseline increases with the duration of continuous employment. At the point of hire from unemployment, workers all start out with a negotiation baseline of $v_0 = \ln(b)$. From there, any outside offer will lead to a wage increase and a new negotiation baseline while productivity shocks cannot lead to wage decreases as $\varepsilon_{\min} \geq b$.²⁴ As one looks at populations that have been continuously employed for longer periods, the fraction of individuals having not yet received their first outside offer since they got out of unemployment (and hence the fraction of individuals with a low negotiation baseline of v_0) gradually shrinks. This selection effect is the driving force behind the increase in the mean wage with s .

Based on the definition (7), we can then compute any set of model-predicted moments to use in the estimation. In practice, as we discuss in the next subsection, we shall use all first- and second-order moments of $G_s(\cdot)$.

The earnings autocovariance structure also conveys potentially useful information about earnings dynamics. We establish in Appendix B that:²⁵

$$\text{Cov}(\ln \phi_t, \ln \phi_{t+s}) = \text{Var}_H(\ln p) - \text{Cov}_G\left(v_t, \int_{v_t}^{v_{\max}} \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2\right)^s dx\right). \quad (8)$$

(Subscripts indicate the distribution with respect to which expectations are taken.) Again these autocovariances are the sum of a constant term (the population variance of the fixed-effect $\ln p$), and a term that decreases down to zero as s goes to infinity. This reflects the limited persistence of wage shocks in our model: the memory of the initial negotiation baseline v_t gradually fades out as workers are hit by productivity shocks and/or outside offers causing renegotiation.

A panel length of T (i.e. T different dates at which we observe a cross-section of individual wages) thus provides us with $3T - 1$ moment conditions (T means and T variances from (7) and $T - 1$ covariances from (8)) on which to base an estimation of the $F(\cdot)$ distribution.

²⁴The negotiation baseline resulting from the first outside offer after an unemployment spell is a draw from $1 - \bar{F}(\cdot)^2$, i.e. the distribution of the minimum of two draws from $F(\cdot)$.

²⁵We continue to work conditionally on continuous employment between dates t and $t + s$. However, to avoid a notational overload, we now keep this conditioning implicit.

4 Data and estimation procedure

Structure of the analysis sample. We use a sub-sample of the British Household Panel Survey (BHPS). The BHPS is a 13-wave (1991 to 2003) panel of household data, of which we use waves 2 to 13, thus following individuals for up to 12 years.²⁶ The BHPS provides information on individual labor market spell histories and precise spell durations (down to the month or the day when not missing), together with records of individual earnings and working hours every 12 months.

There is some attrition from- and entry into the panel, both of which we assume exogenous. We also drop the few individuals that have gaps in their records.

Our working sample is obtained from the following selection of the raw data. We use data on males and females with more than five years' potential experience and aged less than 60, thus cutting five years at both ends of the individuals' working lives. We restrict our analysis to individuals with A-level education or more, both for the sake of brevity and also because the individual-level wage-bargaining/offer-matching process described in the theoretical model is arguably more relevant in high-skill labor markets. For similar reasons, we do not consider individuals observed as self-employed or employed in the public sector in their initial year in the survey.²⁷

Based on these selection rules, we then construct two separate samples. The first one draws from records of individual labor market spell histories and will serve for the estimation of transition parameters: we take all selected individuals at their first interview date, follow them throughout the 12 waves and record all their labor market spell durations and transition types (job-to-job or job-to-unemployment).

Our second sample is an income sample gathering the yearly observations of wages and working hours: we compute hourly wages using data on (before-tax) labor income received in the last month and on worked hours. We then regress these wages on indicators of year, education, gender, ethnic background and labor market cohort. We use the residuals from this latter regression as our measure of individual earnings. We finally trim the data by dropping the top and bottom 2.5% of

²⁶We discard wave 1 (1991) because of substantial coding differences between this and the subsequent waves.

²⁷Ideally, we would have liked to work on a more homogeneous set of workers. Yet as we shall see below, given the resulting sample size, this is probably the finest stratification of the original BHPS data that we can reasonably envisage.

earnings (residuals). This trimming is useful to stabilize our empirical estimates of cross-sectional wage variances.

< **Table 1 about here.** >

Our job spell sample comprises 659 initially employed individuals and 58 initially unemployed individuals, while our income sample comprises 677 individuals with a valid initial wage observation. Table 1 gives more detailed descriptive statistics for the two samples.

Estimation procedure. Following the above developments, we carry out a two-step estimation procedure. In the first step we use the data on labor market spells to estimate the transition rates δ and λ_1 using maximum likelihood on observed job spell durations and job transitions. In the second step, we use our income data to estimate the remaining parameters—i.e. the sampling distribution $F(\cdot)$ of productivity shocks and the distribution of person fixed-effects $H(\cdot)$ —by matching a series of wage means and covariances, as derived in subsection 3.3. Specifically, we match the following $3T - 1$ moments (where T is the panel length in years), for $s = 0, 12, 24, \dots, 12(T - 1)$:²⁸

$$\begin{aligned} E(\ln \phi_{t+s}) &= E_{G_s}(v) \\ \text{Var}(\ln \phi_{t+s}) &= \text{Var}_H(\ln p) + \text{Var}_{G_s}(v) + \sigma_{\text{me}}^2 \\ \text{Cov}(\ln \phi_t, \ln \phi_{t+s}) &= \text{Var}_H(\ln p) - \text{Cov}_G\left(v_t, \int_{v_t}^{v_{\text{max}}} \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2\right)^s dx\right). \end{aligned} \quad (9)$$

All these means and covariances are conditional on continuous employment between times t and $t + s$.²⁹ Also, we set our period length to be one month, hence the series of leads being taken at multiples of 12 periods to match our yearly wage data. Finally note in the second line of (9) the addition of a term σ_{me}^2 to the theoretical expression of cross-sectional income variances. This accounts for the presence of classical measurement error (with variance σ_{me}^2) in hourly wages.

We match these moments using Optimal Minimum Distance (OMD) estimation.

²⁸The first T moments (mean wages) use the normalization $E(\ln p) = 0$.

²⁹This conditioning leads us to discard the information brought by observations for individuals who experienced a complete unemployment spell between t and $t + s$, and still have a wage record at both dates. It is possible to write down the means and autocovariances in (9) conditional on employment at t and $t + s$ only. However, the corresponding formulae are cumbersome and unemployment is a sufficiently rare event in our sample to make the loss of information entailed in the more stringent conditioning inconsequential.

For convenience, we choose to parameterize the population distribution of productivity shocks, $G(\cdot)$ rather than the sampling distribution $F(\cdot)$. The latter can then easily be retrieved from $G(\cdot)$ using equation (3). As detailed in equation (4) in the theoretical section, $G(\cdot)$ is the sum of a mass point at v_0 corresponding to entry wages for previously unemployed workers who all start off their employment spell with a negotiation baseline of v_0 , and a transformation of the sampling distribution $F(\cdot)$ with lower support v_{\min} . We thus parameterize $G(\cdot)$ as the sum of a mass point (at v_0) and a normal distribution, truncated below at v_{\min} . Note that because match surpluses are nonnegative for any realization of the productivity shock, it has to be the case that $v_{\min} \geq v_0$. In fact we assume $v_{\min} = v_0$. This assumption seems natural in that it means that jobs and job offers exist for values of the productivity shock down to a value leading to a match surplus of zero (see the discussion in footnote 9).

As to the distribution of worker fixed effects $H(\cdot)$, we see that only its variance, $\text{Var}_H(\ln p)$ appears in the series of moments we aim to match—see (9). The distribution $H(\cdot)$ can however be retrieved from our estimate of $G(\cdot)$ and the actual wage distribution by deconvolution. Yet as we shall see below, knowledge of this estimated variance will be sufficient for all of our simulation exercises.

5 Results

5.1 Parameter estimates

Transition rates. The estimated arrival rates of outside offers and job destruction shocks are reported in Table 2, in monthly values. The probability of receiving an outside offer is about 1.5 percent per month. Because the probability of job switching is equal to $\lambda_1/2$ (see subsection 3.1), this number says that about one in 130 employed workers from our high-skill sample switches jobs each month. With a job destruction rate of just under a third of one percent (implying an average waiting time of about 25 years between two spells of unemployment), the average duration of a job spell, $(\delta + \lambda_1/2)^{-1}$, is in the order of 7.9 years.

< **Table 2 about here.** >

Finally turning to the unemployment exit rate λ_0 , we estimate it at 4.11 percent monthly implying a mean unemployment duration of just over two years. Combined with the estimated value of δ , it also implies an unemployment rate of 7.2%. Considering that we have a sample of highly educated workers, this may sound like a long duration and a high rate. These numbers, however, accommodate the contents of the data.³⁰

Distributions. Table 3 contains the OMD estimates and standard errors of the various distributional parameters involved in our moment conditions (9).³¹ A graphical rendering of the resulting distributions $F(\cdot)$ and $G(\cdot)$ is provided in the bottom-right panel of Figure 2 below.

< **Table 3 about here.** >

The first thing we notice from the first row of Table 3 is the relatively poor precision of our estimate of σ_{me}^2 , the variance of the measurement error. Based on this sample we can only reject equality to zero at the 10 percent level (p -value of 0.068 against the alternative $\sigma_{\text{me}}^2 = 0.012$, our point estimate). Hence in the bottom row of Table 3 we also report parameter estimates obtained subject to the constraint $\sigma_{\text{me}}^2 = 0$. Comparison of the two rows shows that the only parameter on which these constraints have any impact at all is σ —the standard deviation of $G(\cdot)$, the population distribution of negotiation baselines, as explained in section 4—, which is estimated slightly higher in the absence of measurement error.

The minimum negotiation baseline, $v_0 = v_{\min}$ that workers start off with when first hired from unemployment is estimated at 1.78, to be compared with a value of the mean negotiation baseline (which equals the mean log wage residual under our normalization $E_H(\ln p) = 0$) over all employed workers of 2.17. Recalling that $G(v)$ is parameterized as the sum of a mass of $\frac{\delta}{\delta + \lambda_1}$ at v_0 and a normal distribution of mean μ and variance σ^2 truncated at v_{\min} , our estimated transition rates δ and λ_1 put the fraction of employed workers with a minimum negotiation baseline at 17.9%.³²

³⁰Unemployment notoriously appears to be highly persistent in the BHPS data (Stewart, 2004). Also, the mean unemployment duration implied by our estimated value of λ_0 may seem at odds with the mean duration of unemployment spells reported in Table 1. We should bear in mind that our estimated λ_0 aims to fit both spell durations and the initial (un)employment rate observed in our sample.

³¹The reported standard errors do not account for the presence of the nuisance parameters δ and λ_1 , which appear in the theoretical moments and which we fix to their estimated value from our first estimation step. Hence the standard errors reported in Table 3 are likely slightly understated.

³²The model further has it that v_0 should equal the mean log wage among workers who were hired out of unem-

The variance of the worker fixed-effect distribution, $\text{Var}_H(\ln p)$, is 0.062. Taking the variance of the measurement error to equal 0.012, this suggests that the relative contributions of the permanent component of earnings $\ln p$, the transitory component of earnings v_t , and measurement error to the total cross-sectional earnings variance (which equals 0.145) are 34% for $\ln p$, about 50% for v_t , and 8% for the measurement error (this latter number being little more than indicative, given the lack of precision with which we estimate σ_{me}^2).

Finally, as we use 12 years of data, (9) corresponds to a set of $3T - 1 = 35$ moments to match in the estimation, for only 5 parameters to estimate. Table 3 reports the test statistic and p -value of a Wald test of overidentifying restrictions. With a p -value of 0.55, the model passes this specification test.

5.2 Simulations and fit analysis

Matched moments. Figure 2 displays the moments listed in (9) that we match in the estimation, as they are predicted by our model with the estimated parameters and as they are observed in our sample. In all panels the dotted lines materialize 95% confidence bands around the empirical moments.

< **Figure 2 about here.** >

The top left panel shows the progression of the mean log-wage as one conditions on zero to $12T (= 144)$ months of continuous employment. In the data, this mean log-wage increases from 2.17 for the whole employed population to 2.24 for the subset of workers with a duration of continuous employment of 12 years or more.³³ This wage growth is well replicated by our model.

The top right panel shows the evolution of the variance of the distribution of log-earnings conditional on zero to $12T$ months of continuous employment. We showed above that the underlying conditional distribution, $G_s(\cdot)$, gradually shifts from the steady-state population distribution $G(\cdot)$ when we condition on a minimum length of continuous employment of zero periods, to $G_\infty(\cdot)$ when s

ployment within the last month. Unfortunately we can only obtain a very imprecise direct estimate of this mean, as the number of unemployment exits occurring within a month of an interview date—i.e. a date at which we have a wage observation—in our sample is very small (27 observations in total). Interestingly, however, this direct estimate is 1.87 (std. err. 0.11), which is statistically consistent with the estimate of v_0 shown in Table 3.

³³These individuals account for 34.7% of our sample.

is increased indefinitely. The predicted variance of $G_s(\cdot)$ clearly declines with the minimum duration of uninterrupted employment, s . In spite of the poor precision of the empirical counterparts of these conditional wage variances, the Figure is also suggestive of a decline of the latter with s , at least up to about $s = 96$ months (8 years). While staying well within the confidence bands, the model still seems to have a tendency to over-predict this decline at long lags. Yet at this level of precision, we may say that the empirical pattern is consistent with the model.

Finally, the bottom-left panel graphs wage autocovariances at 0 to $12T$ lags (again conditional on continuous employment over the number of lags considered). As expected, these autocovariances decline as one looks at longer lags, in a way that the model captures well. Note in particular the convergence of $\text{Cov}(\ln \phi_t, \ln \phi_{t+s})$ toward $\text{Var}(\ln p)$ as the number of lags increases.

Wage changes. Using the parameter estimates obtained above, we can now create a sample of simulated data which can then be compared to the real data along any dimension we like. This will allow us to assess the model's capacity to replicate features of the data that were not directly matched in the estimation. In the reported simulations we alternatively use the values of $\sigma_{\text{me}}^2 = 0$ and $\sigma_{\text{me}}^2 = 0.012$ for the variance of the measurement error.³⁴

An intuitive way of looking at earnings dynamics beyond first- and second-order moments (which, as we just saw, the model replicates well enough) is to consider the distribution of wage changes from one year to another. In the language of our theory, this is the distribution of $v_{t+s} - v_t$, with s some chosen horizon.

< **Figure 3 about here.** >

Figure 3 plots three such distributions of wage changes over one year ($s = 12$ months, left panel) and ten years ($s = 120$ months, right panel). On both panels, the solid line is the empirical cdf of log wage changes.³⁵ The dashed line is the model-predicted counterpart of that cdf $\sigma_{\text{me}}^2 = 0.012$. Finally, the dash-dot line with a mass at zero is the cdf of simulated wage changes if one removes

³⁴The construction of a sample of simulated data obviously requires additional distributional assumptions about the person fixed-effect ($\ln p$) and the measurement error, of which we only know the variance. We assume a normal distribution for both (even though the distribution of $\ln p$ plays no part in the exercises to come).

³⁵In the one-year-horizon case, all waves of our 1992-2003 analysis panel pooled. The ten-year-horizon graph obviously has to rely on changes between the first two and last two waves only.

measurement error from the model.

First comparing the “observed” and “predicted, with measurement error” distributions of yearly wage changes (left panel), we observe a near perfect replication of the top half of the distribution (i.e. the part of the distribution corresponding to wage increases) and a slightly larger discrepancy in the bottom half, with the observed distribution tending to dominate the simulated one at the first order. Next turning to the ten-year horizon graph (right panel), we get a somewhat more distinct impression of first-order stochastic dominance, with the model overstating the median earnings change over ten years by some 4.5 percentage points. While the overall look of the graph is sensitive to the particular calibration of the measurement error, the robust conclusion is that the model tends to slightly over-predict wage cuts, particularly large ones.

The model also predicts that many of the observed wage changes are in fact artificial and only reflect measurement error: once measurement error is removed from the model, the predicted distribution of year-to-year wage changes has a mass of about two thirds at zero. Interestingly, the distribution without measurement error also seems more skewed than the observed one, as the model-predicted share of “genuine” wage increases is roughly 12%, whereas the corresponding figure for wage cuts is about a quarter. Finally note from the right panel of Figure 3 that the predicted relative importance of measurement error in explaining the distribution of reported wage changes becomes negligible as one looks at longer horizons.

Conditioning on job-to-job mobility. The model has very strong implications about earnings dynamics around a job-to-job transition. As mentioned above, the date- t negotiation baseline of a worker changing jobs at date t is equal to the date- t idiosyncratic match productivity shock at the incumbent firm, which has to be less than productivity at the poaching firm if the worker has switched from the former to the latter. In other words, conditional on observing a job-to-job transition, the negotiation baseline is the minimum of two independent draws from $F(\cdot)$, hence a draw from $1 - \bar{F}(\cdot)^2$. Most importantly, this new negotiation baseline v_t is independent of any previous negotiation baseline v_{t-s} .

Are these predictions borne out by the data? A partial answer can be sought in the following

additional moment conditions, which are implied by the above considerations:³⁶

$$\begin{aligned}
E(\ln \phi_t \mid \text{job-to-job transition at } t) &= E_{1-\bar{F}^2}(v) \\
\text{Var}(\ln \phi_t \mid \text{job-to-job transition at } t) &= \text{Var}_H(\ln p) + \text{Var}_{1-\bar{F}^2}(v) + \sigma_{\text{me}}^2 \\
\text{Cov}(\ln \phi_t, \ln \phi_{t+s} \mid \text{job-to-job transition between } t+1 \text{ and } t+s) &= \text{Var}_H(\ln p). \quad (10)
\end{aligned}$$

Computation of the first two moments in (10) has to rely on independent wage observations that coincide with the occurrence of a job-to-job mobility, i.e. on the subset of workers moving from job to job in a month preceding an interview date. Because we only have one interview—therefore at best one wage observation—each year, and because job-to-job transitions are infrequent events, we only have very few (indeed 25) such independent coincidences in our data set.³⁷ This does not allow for a very precise estimation of the mean and variance of wages conditional on mobility: the empirical mean is 2.25 (standard error of 0.09), and the empirical variance equals 0.18 (standard error of 0.05). Yet, the corresponding model predictions are 2.63 and 0.10. Hence, while our prediction of the conditional variance is still acceptable (albeit on the low side), the model seems to overstate the mean wage of job-to-job movers somewhat.

Computation of the series of conditional covariances in (10) can rely on a slightly wider set of observations, as they only involve pairs of wage observations that are anywhere on either side of a job-to-job mobility (as opposed to wage observations that exactly coincide with a job-to-job mobility). Exploiting this, Figure 4 depicts the empirical covariance of wages at dates $t = 1992$ (the initial year) and $t + s$ (for $s = 0, 12, 24, \dots, 12(T - 1)$ months) among workers who have

³⁶Of course these considerations imply much more than these moment restrictions. Indeed they even offer a potential source of nonparametric identification of our model (up to the measurement error). Consider a worker experiencing a job-to-job transition, for whom we have two wage observations, one on each side of the transition. Let $\ln \phi_b = \ln p + v_b$ denote the wage observed before the transition and $\ln \phi_a = \ln p + v_a$ the wage after the transition. From the above we know that v_a is independent of v_b . In principle we can thus retrieve $H(\cdot)$ and $G(\cdot)$ by nonparametric deconvolution—and hence $F(\cdot)$ as well, using (3). (This is an application of a general identification theorem by Kotlarski, 1967.) Unfortunately, implementation of this method requires a sufficient number of independent observations of a worker with a job-to-job mobility occurring between two valid wage observations. We only have in the order of 55 such observations in our sample.

³⁷The measurement of these “mobility” wages also runs into another problem, which is that the date at which a job transition exactly occurs is certainly measured with error. For instance, some workers are known to report as having started a new job spell, when they have really only accepted an offer for a job that is effectively to start at some (near) future date. In these cases, it is unclear whether the reported wage pertains to the new or to the old job. Because it likely adds some “pre-mobility” wages into our sample of mobility wages, this type of measurement error tends to bias our empirical estimate of the conditional mean $E(\ln \phi_t \mid \text{job-to-job transition at } t)$ downward, and that of the corresponding variance upward.

experienced at least one job-to-job transition between $t + 1$ and $t + s$. For comparison with the model’s prediction, a horizontal line at $\text{Var}_H(\ln p)$ is also drawn. Finally, at $s = 0$ the Figure reports the observed and predicted conditional variance $\text{Var}(\phi_t \mid \text{job-to-job transition at } t)$. Dotted lines represent confidence bands around the empirical moments.

< **Figure 4 about here.** >

Again given the scarce numbers of observations upon which we have to base our computations,³⁸ this Figure only paints an indicative picture of the covariance profile of individual income around a job-to-job transition. Yet it still suggests that this profile is both lower and markedly “flatter” than the corresponding unconditional covariance profile plotted in the bottom-left panel of Figure 2. Both properties are in accordance with-, and indeed quantitatively well captured by the model.

A related property of the wage process generated by our model is also worth pointing out. It can be shown that the wage earned by an employed worker is independent of the number of outside offers received by the worker since the beginning of his/her employment spell, provided that this number is at least equal to one.³⁹ This, together with the fact that the transitory component v_{t+1} of the new wage obtained by a job mover is independent of the past transitory component v_t in the previous job, has the particular implication that wage gains for job-to-job changers are independent of the number of offers raised in the past, or of the number of past job changes. Barlevy (2003) finds empirical support for such independence using NLSY data in a different (search) context.

Linear ARMA model. As we mentioned in the introduction, the literature has a long tradition of fitting ARMA-type models to individual wage trajectories. In this paragraph we take another look at the covariance structure of (observed and simulated) wages under this alternative angle.

³⁸These numbers range from 21 to 60, depending on s .

³⁹A straightforward proof of this claim relies on flow-balance equations similar to (3). Details are available on request.

The following describes a canonical ARMA specification found in the literature:

$$\left\{ \begin{array}{l} \ln \phi_{i,t} = z_i + a_{i,t}^P + a_{i,t}^T, \\ a_{i,t}^P = a_{i,t-1}^P + \zeta_{i,t}, \\ a_{i,t}^T = \sum_{\ell=0}^q \theta_{\ell} \xi_{i,t-\ell}, \end{array} \right. \quad \begin{array}{l} \text{with } \zeta_{i,t} \text{ i.i.d.}, \\ \text{with } \xi_{i,t} \text{ i.i.d. and } \theta_0 = 1. \end{array} \quad (11)$$

Here log-earnings are modelled as the sum of an individual fixed-effect z_i , a permanent earnings shock $a_{i,t}^P$ following a martingale process, and a transitory earnings shock $a_{i,t}^T$ following an MA(q) process. The order q of the latter MA process is to be determined empirically, along with the parameters of this process (the θ_{ℓ} 's) and the innovation variances, σ_{ζ}^2 and σ_{ξ}^2 .

We fit model (11) separately to our income sample from the BHPS and to a sample of simulated data based on the parameter estimates obtained above.⁴⁰ The order of the MA process is determined by looking at the sequence of autocovariances of first-differenced log wages, $\text{Cov}(\Delta \ln \phi_{i,t}, \Delta \ln \phi_{i,t+s})$, which should equal zero for any $s \geq q + 2$. The top panel of Table 4 reports these autocovariances at lags of up to 4 years for both the BHPS and simulated sample. We first notice that the pattern of wage autocovariances is very similar in the real and in the simulated data samples, which is expected given the good fit obtained with our model with respect to these covariances (as illustrated in the bottom left panel of Figure 2). Second, the magnitude of the covariance point estimates drops tenfold in both samples between the first and the second lag and remains very small thereafter. Both patterns square in well with earnings following an MA(1) process in growth rates, thus implying that earnings levels can be described along the lines of model (11) as the sum of a random walk component and a serially uncorrelated—or MA(0)—component. However, for completeness we estimate model (11) under both the MA(0) and MA(1) specifications.⁴¹

⁴⁰We make the simulated sample of equal size to the observed one. An important detail to keep in mind is that, for this paragraph, the time unit is taken to be one year (as opposed to one month, as was the case thus far), in accordance with the yearly frequency of wage observations in the BHPS data, so our simulated sample is made up of yearly observations taken from a monthly simulated dataset. Finally, we ignore the presence of measurement error in this comparative exercise. The measurement error variance cannot be identified within the specification (11)—see Meghir and Pistaferri (2004).

⁴¹Results from the literature conclude to the presence of either an MA(0) or an MA(1) component in the earnings process, but not usually to higher order MA components. For the estimation, we proceed by OMD matching of the autocovariance structure of yearly wage growth. Details are available on request.

< Table 4 about here. >

Results for both samples and both specifications are displayed in the bottom panel of Table 4. Looking at the variances of innovations for the permanent and transitory components of the earnings process, σ_{ζ}^2 and σ_{ξ}^2 , we observe again much similarity between those obtained with the simulated sample and with the BHPS data.⁴² Most intriguingly, as commonly found in the literature, we obtain a significant variance for the innovation of the permanent earnings shock, thus concluding to the presence of a random walk component of the individual earnings process, both in the real and in the simulated data. Yet we know that, at least for the simulated data, the fitted ARMA process is misspecified and the true DGP is stationary.⁴³ This illustrates the difficulty of numerically distinguishing between a process truly exhibiting a unit root and other forms of persistent, possibly nonlinear processes.

The following exercise provides further illustration of this latter point. The ARMA specification (11) was estimated by fitting moments of the wage growth rates (i.e. by taking first differences of the wage data). The results shown in table Table 4 indicate that our structural model correctly captures these latter moments, even though it was fitted to the data in levels. One also may ask the converse question of how good the ARMA model is at replicating the autocovariance structure of wage levels that was used in the estimation of the structural model. Under the simple assumption of homoskedastic innovations, this structure is characterized as follows:

$$\begin{aligned}\text{Var}(\ln \phi_{i,t+s}) &= \text{Var}(z_i + a_{i,t}^P) + s\sigma_{\zeta}^2 + (1 + \theta_1^2)\sigma_{\xi}^2, \\ \text{Cov}(\ln \phi_{i,t}, \ln \phi_{i,t+s}) &= \text{Var}(z_i + a_{i,t}^P) + \theta_1\sigma_{\xi}^2 \times \mathbf{1}_{\{s=1\}}.\end{aligned}\tag{12}$$

Figure 5 shows the fit of (12) to the data in the same way as Figure 2 did for the corresponding predictions based on the structural model.⁴⁴ The very poor fit seems to point to an inconsistency

⁴²Wald tests further indicate that both specifications are accepted based on either sample at the 5% level (last column of the Table). Moreover, neither sample allows rejection of the MA(0) against the MA(1) specification at the 5% level. It should be recognized, however, that acceptance of either specification based on simulated data is fragile and varies from one simulated sample to another from marginal acceptance to marginal rejection.

⁴³Indeed, according to our theoretical model, earnings have a steady-state distribution, $G(\cdot)$, with finite variance.

⁴⁴We are only reporting second-order moments to save on space. For means, the simple ARMA decomposition (11) postulates that $E(\ln \phi_{i,t+s})$ is independent of s , which is a poor fit to the empirical pattern observed in the top-left panel of Figure 2. Figure 5 shows predictions based on both assumptions (0 or 1) about the order of the MA process in (11). Also we should mention that the intercept term in (12), $\text{Var}(z_i + a_{i,t}^P)$, was set to fit the cross-sectional wage

between the ARMA representation of wage growth and what would be the ARMA representation of wage levels.⁴⁵ The point of all this is not to say that a more sophisticated ARMA decomposition would not perform better,⁴⁶ but rather to emphasize the difficulty of choosing a specific statistical model of income dynamics without theoretical guidelines.

< **Figure 5 about here.** >

6 Conclusion

Our concern in this paper has been the ability of a simple structural model to replicate the main features of the dynamics of individual labor earnings observed in the data. Our proposed model belongs to the family of search models *à la* Diamond-Mortensen-Pissarides (DMP). Our specific assumptions are that we allow on-the-job search and assume that wages can only be renegotiated with mutual consent by the firm and the worker. We investigate whether such a model can produce a quantitatively plausible “internal propagation mechanism” of i.i.d. productivity shocks into persistent wage shocks using a 10-year panel of highly educated British workers. Our key contributions are the following.

First, we formalize the assumption of renegotiation by mutual consent in the context of a job search model with idiosyncratic productivity shocks to match quality. We then scrutinize the model outcomes in terms of individual earnings dynamics, whereas the existing job search literature usually focuses on cross-sectional wage dispersion. Because the mutual consent rule will only allow the wage level to be altered when one party has a credible threat to leave the match, wage dynamics generated by our model are more persistent than under the constant-share Nash surplus sharing rule often used for wage determination.

Second, we show that our model, when estimated with twelve waves of BHPS data, produces

variance at the initial date t . While this is not a very efficient way to estimate this parameter, its only impact is to shift the (co)variance/time profile up or down, thus changing very little to the overall fit.

⁴⁵This inconsistency was already noticed in US data by Baker (1997). Interestingly, Baker advocates an alternative specification of the wage process—which he calls the “profile-heterogeneity model”—where wages are linear functions of experience with individual-specific intercepts and slopes. Even though our structural model is nonlinear and only features endogenous (and non-systematic) heterogeneity in wage/experience profiles, it is formally closer to Baker’s preferred profile-heterogeneity model than to the linear ARMA model (11).

⁴⁶For instance Meghir and Pistaferri (2005) show that the variance of innovations in a decomposition similar to (11) follow a relatively complex time pattern.

a dynamic earnings structure which is remarkably consistent with the data. The main features of individual earnings dynamics, such as the covariance structure of earnings, the evolution of individual earnings mean and variance with the duration of uninterrupted employment, or the distribution of year-to-year earnings changes are very well matched by our model predictions.

Third, we offer a structural counterpart to the “reduced-form” literature on individual earnings processes and we establish that the combination of on-the-job search and renegotiation by mutual agreement is a promising candidate explanation of the widely documented persistence of earnings shocks. Our theory suggests that wage dynamics should be thought of as the outcome of a specific acceptance/rejection scheme of i.i.d. wage shocks, thus offering an alternative to the conventional linear ARMA-type approach. Moreover, it highlights the link between labor market competition (as measured by the probability of raising outside job offers when employed), worker mobility across jobs, and individual earnings dynamics.

There are several avenues for further work building on this paper. We now briefly discuss some of these. A first, relatively straightforward extension would be to close our theoretical model in the manner of the DMP framework to include the firm’s job creation decision. We could also allow for the job destruction to become endogenous by allowing the match surplus to be negative over some of the support of the distribution of productivity shocks. Although such a closed model would be difficult to estimate (it would be fraught with the well-known difficulty of estimating a matching function), a calibrated version of it could still be used to analyze the impact of various labor market policies. One could also simulate wage dynamics in response to business cycle shocks and assess whether our model can reproduce the dampened cyclical fluctuations of wages (as compared to fluctuations in productivity) observed in the data. Finally, this would allow us to relate our model to well known empirical results of the contract literature, such as Beaudry’s and DiNardo’s (1991) striking observation that individuals’ current wages are more strongly affected by the lowest unemployment rate since the start of their job than by the current unemployment rate or the unemployment rate at the start of the job.⁴⁷

⁴⁷Even though these authors, along with much of the related contract literature, emphasize risk-sharing considerations as the main driving force behind individual earnings dynamics, our model still bears a close formal relationship

A second possible extension would be to incorporate the impact of human capital accumulation into our model. While we have ignored it completely in the present paper, it obviously lines up as a potential cause of the positive returns to continuous employment. A simple way of introducing human capital into our framework would be to allow the individual-specific component of match productivity to grow at a constant rate over time.⁴⁸ Although this will make the derivations of the predicted moments of interest much more cumbersome, it is likely to improve the fit of the model in terms of the frequency and size of wage cuts, which the present model tends to overestimate.

A third avenue of potential improvement would be the addition of worker heterogeneity in the transition rates, i.e. the hazards of job destruction, job finding and outside offers. This would help with our model's currently counterfactual implication that these three hazards are constant with spell duration or independent of current wages.

Finally, our model could be enriched by a more careful modelling of the employer side and a more precise description of what we referred to as “match productivity” or “match quality” shocks in terms of firm-specific and truly match-specific shocks. On the theoretical side, the introduction of permanent firm heterogeneity is a far-from-trivial extension of our model, as it complicates the derivation and comparison of the workers' valuation of employment at different firms by an order of magnitude. On the empirical side, direct evidence on productivity shocks can in principle be found in matched employer-employee data of the type used in Abowd, Kramarz and Margolis (1999), Cahuc, Postel-Vinay and Robin (2006) or Guiso, Pistaferri and Schivardi (2005). In this latter reference, the authors investigate in a reduced-form setting the pattern of correlation between firm productivity shocks and individual wage shocks. Whether a job search model of the type we describe in this paper can help to explain this empirical pattern remains an open question.

References

- [1] Abowd, J. M. and D. Card (1989), “On the Covariance Structure of Earnings and Hours Changes”, *Econometrica*, 57(2), 411-45.

to theirs. Malcomson (1997) indeed notes that the assumption of renegotiation by mutual consent squares in well with the evidence documented in Beaudry and DiNardo (1991).

⁴⁸Such derivations are available upon request.

- [2] Abowd, J. M., F. Kramarz and D. N. Margolis (1999) “High Wage Workers and High Wage Firms”, *Econometrica*, 67, 251-333.
- [3] Alvarez, J., M. Browning, and M. Ejrnæs (2001), “Modelling Income Processes with Lots of Heterogeneity”, CAM discussion paper 2002-01.
- [4] Baker, M. (1997), “Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings”, *Journal of Labor Economics*, 15 (2), 338-75.
- [5] Barlevy, G. (2003), “Estimating Models of On-the-Job Search using Record Statistics”, Federal Reserve Bank of Chicago Working Paper #2003-18.
- [6] Beaudry, P. and J. DiNardo (1991), “The Effect of Implicit Contracts on the Movement of Wages over the Business Cycle: Evidence from Micro Data”, *Journal of Political Economy*, 99(4), 665-88.
- [7] Binmore, K., A. Shaked and J. Sutton (1989), “An Outside Option Experiment”, *Quarterly Journal of Economics*, 104(4), 753-70.
- [8] Blundell, R. and I. Preston (1998), “Consumption Inequality and Income Uncertainty”, *Quarterly Journal of Economics*, 113, 603-40.
- [9] Burdett, K. and M. G. Coles (2003), “Wage/Tenure Contracts and Equilibrium Search”, *Econometrica*, 71(5), 1377-404.
- [10] Burdett, K. and D. T. Mortensen (1998), “Wage Differentials, Employer Size and Unemployment”, *International Economic Review*, 39, 257-73.
- [11] Cahuc, P., F. Postel-Vinay and J.-M. Robin (2006), “Wage Bargaining with On-the-job Search: Theory and Evidence”, *Econometrica* (forthcoming).
- [12] Dey, M. S. and C. J. Flinn (2005), “An Equilibrium Model of Health Insurance Provision and Wage Determination”, *Econometrica*, 73(2), 571-627..
- [13] Diamond, P. A. (1982), “Wage Determination and Efficiency in Search Equilibrium”, *Review of Economic Studies*, 49, 761-82.
- [14] Eckstein, Z. and G. J. Van den Berg (2005), “Empirical Labor Search: A Survey”, forthcoming *Journal of Econometrics*.
- [15] Fella, G. (2004), “Optimal Severance Pay in a Matching Model”, Queen Mary University of London mimeo.
- [16] Guiso, L., L. Pistaferri and F. Schivardi (2005), “Insurance Within the Firm”, *Journal of Political Economy*, 113(5), 1054-87.
- [17] Harris, M. and B. Holmström (1982), “A Theory of Wage Dynamics”, *Review of Economic Studies*, 49, 315-33.

- [18] Kotlarski, I. (1967), “On Characterizing the Gamma and Normal Distributions”, *Pacific Journal of Mathematics*, 20, 69-76.
- [19] Lillard, L. A. and Y. Weiss (1979), “Components of Variation in Panel Earnings Data: American Scientists 1960-1970”, *Econometrica*, 47(2), 437-54.
- [20] Lillard, L. A. and R. J. Willis (1978), “Dynamic Aspects of Earnings Mobility”, *Econometrica*, 46(5), 437-54.
- [21] MacLeod, W. B. and J. M. Malcomson (1993), “Investments, Holdup, and the Form of Market Contracts”, *American Economic Review*, 83(4), 811-37.
- [22] MaCurdy, T. E. (1982), “The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis”, *Journal of Econometrics*, 18(1), 83-114.
- [23] Malcomson, J. M. (1997), “Contracts, Hold-Up, and Labor Markets”, *Journal of Economic Literature*, XXXV, 1916-57.
- [24] Malcomson, J. M. (1999), “Individual Employment Contracts”, in O. Ashenfelter and D. Card (ed.), *Handbook of Labor Economics*, chapter 35, vol. 3B, 2291-372, Elsevier Science.
- [25] Meghir, C. and L. Pistaferri (2004), “Income Variance Dynamics and Heterogeneity”, *Econometrica*, 72(1), 1-32.
- [26] Mortensen, D. (2003), *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge, MA: MIT Press.
- [27] 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.
- [28] Mortensen, D. T. and C. A. Pissarides (2003), “Taxes, Subsidies and Equilibrium Labor Market Outcomes”, in E. S. Phelps, editor, *Designing Inclusion: Tools to Raise Low-end Pay and Employment in Private Enterprise*, Cambridge: Cambridge University Press.
- [29] Pissarides, C. A. (2000), *Equilibrium Unemployment Theory*, 2nd edition. Cambridge, MA: MIT Press.
- [30] Postel-Vinay, F. and J.-M. Robin (2002a), “The Distribution of Earnings in an Equilibrium Search Model with State-Dependent Offers and Counter-Offerers”, *International Economic Review*, 43(4), 1-26.
- [31] Postel-Vinay, F. and J.-M. Robin (2002b), “Wage Dispersion with Worker and Employer Heterogeneity”, *Econometrica*, 70(6), 2295-350.
- [32] Postel-Vinay, F. and J.-M. Robin (2006), “Microeconomic Search-Matching Models and Matched Employer-Employee Data”, in R. Blundell, W. Newey and T. Persson, editors, *Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress* (forthcoming).

- [33] Shimer, R. (2006), “On-the-Job Search and Strategic Bargaining”, *European Economic Review* (forthcoming).
- [34] Stewart, M. (2004), “The Inter-related Dynamics of Unemployment and Low-Wage Employment”, mimeo.
- [35] Sutton, J. (1986), “Non-Cooperative Bargaining Theory: An Introduction”, *Review of Economic Studies*, 53, 709-24.
- [36] Thomas, J. and T. Worrall (1988), “Self-Enforcing Wage Contracts”, *Review of Economic Studies*, 55, 541-54.

APPENDIX

A The full theoretical model with forward-looking agents

In this appendix we show how the simple theory laid out in the main text can be amended in order to properly account for the agents' forward-looking behavior. We thus now assume that workers and firms are infinitely lived, forward-looking, risk-neutral and have a common exogenous per-period discount factor of $\beta \in (0, 1)$.

More notation. The first thing that needs to be amended when considering $\beta > 0$ is notation. Given a match between a firm and a type- p worker, with current productivity ε and current wage ϕ , we now denote the worker's valuation of this match will be denoted by $V(\phi, p)$, and the firm's valuation by $\Pi(\varepsilon, \phi, p)$. $V(\cdot)$ and $\Pi(\cdot)$ are now present discounted sums of future expected income or profit flows. We assume from the outset that $V(\cdot)$ is increasing in ϕ and is independent of ε , while $\Pi(\cdot)$ is increasing in ε and is decreasing in ϕ . The consistency of these assumptions will be verified later on. We still assume that a vacant job slot is worth 0 to the firm, and we now denote the lifetime value of unemployment by $V_0(p)$.

We define total match surplus as the value of the match relative to the combined value of a vacant job and an unemployed worker:

$$S(\varepsilon, p) = [V(\phi, p) - V_0(p)] + [\Pi(\varepsilon, \phi, p) - 0]. \quad (\text{A1})$$

We shall start working under the provisional assumption that $S(\cdot)$ is independent of any wage value. This will be shown later to be a consistent assumption given risk-neutrality of workers and firms and given our (privately efficient) surplus-sharing mechanism.⁴⁹ Moreover, total match surplus only depends on the determinants of current and future match output flows. Given the assumed memoryless process for transitory shocks ε , those only include the permanent component of match productivity, p , and the current value of its transitory component, ε .

Wages and negotiation baselines. Given the rules of wage determination exposed in the main text and given the definitions of the various relevant value functions in the previous paragraph, it is still the case that all wages can be expressed as functions of the worker's type p and a negotiation baseline r as follows:

$$\phi = \phi(r, p) \quad \Leftrightarrow \quad V(\phi, p) = V_0(p) + S(r, p). \quad (\text{A2})$$

Again, r measures the quality of the match from which the worker was last able to extract the whole surplus. In functional terms, r is the inverse of the wage function $r \mapsto \phi(r, p)$.

⁴⁹Intuitively, total match surplus involves the present discounted sum of expected future flow values of match surplus, which we saw in the main text are independent of any wage value.

Finally notice that equation (A2) encompasses the special case of starting wages offered to unemployed workers:

$$\phi_0(p) = \phi(r_0, p) \Leftrightarrow V[\phi_0(p), p] = V_0(p). \quad (\text{A3})$$

Unemployment value $V_0(p)$. In any given period, an unemployed worker receives a flow income of $b \cdot p$. In the following period, that same worker can either fail to meet a firm (an event of probability $1 - \lambda_0$), in which case s/he stays unemployed and gets a continuation value of $V_0(p)$, or s/he can meet a firm and be hired (probability λ_0). In this latter case, our assumption that firms are able to extract the entire surplus from newly formed matches implies that the worker's continuation value from finding a job is again equal to $V_0(p)$. Hence, given the worker's discount factor of β , the value of unemployment is simply defined by $V_0(p) = b \cdot p + \beta \cdot V_0(p)$.

Total match surplus $S(\varepsilon, p)$. As we saw in the main text, in any given period, the flow surplus from a match between a firm and a type- p worker with current match-specific parameter ε is $p \cdot (\varepsilon - b)$.

If the match is dissolved in the following period, then the worker becomes unemployed and receives a continuation value of $V_0(p)$, while the employer is left with the option of opening a vacant job slot, which is worth zero. Hence the continuation surplus of the firm-worker pair is zero in this case.

If the match is not dissolved, then given the new value of the transitory component of productivity ε' the continuation surplus associated with the incumbent match is $S(\varepsilon', p)$. However, with probability λ_1 , the worker meets a potential alternative employer with match quality η' (drawn from $M(\cdot)$) and associated potential match surplus $S(\eta', p)$. As described in the main text, two configurations can arise. Either $\varepsilon' > \eta'$, in which case the worker stays with the incumbent firm (possibly under a renegotiated contract) with a continuation surplus equal to $S(\varepsilon', p)$, or $\eta' > \varepsilon'$ and the worker joins the poaching firm. In this latter case, the incumbent firm is left with a value of 0 while Bertrand competition between the two employers implies that the worker extracts all the surplus from the incumbent match, i.e. $S(\varepsilon', p)$. All this implies that the sum of the worker's continuation value and the incumbent firm's continuation profits is equal to $S(\varepsilon', p)$, whether an outside offer was raised by the worker or not.

Summing up, given a common discount factor of β for the worker and the employer, total match surplus $S(\varepsilon, p)$ is defined recursively by:

$$S(\varepsilon, p) = p \cdot (\varepsilon - b) + \beta(1 - \delta) \cdot \int_{\varepsilon_{\min}}^{\varepsilon_{\max}} S(\varepsilon', p) dM(\varepsilon') = p \cdot \left(\frac{\bar{\varepsilon} - b}{1 - \beta(1 - \delta)} + \varepsilon - \bar{\varepsilon} \right), \quad (\text{A4})$$

where $\bar{\varepsilon} = E_M(\varepsilon)$.

Before going any further, it is worth looking at the negotiation baseline r_0 that workers start with when they leave unemployment. Given (A3) and (A4), r_0 satisfies $S(r_0, p) = 0$ and thus $r_0 = \frac{b - \beta(1 - \delta)\bar{\varepsilon}}{1 - \beta(1 - \delta)}$. Starting wages $\phi_0(p)$ can then be expressed as $\phi(r_0, p)$. Also note that $r_0 < b$ provided that $\bar{\varepsilon} > b$, which is a necessary condition for any trade at all to take place in the labor market.

Worker value $V(\phi(r, p), p)$. The current-period flow earnings of an employed worker is simply her/his current wage $\phi(r, p)$. When the next period begins, with probability δ the match is hit by a dissolution shock. In this case the worker becomes unemployed and therefore receives a continuation value of $V_0(p)$. With probability $1 - \delta$, the worker stays employed and her/his continuation value depends on the new value of her/his wage and hence on the new value of her/his negotiation baseline r' . Formally:

$$\begin{aligned} V[\phi(r, p), p] &= \phi(r, p) + \beta \cdot [\delta V_0(p) + (1 - \delta) E(V[\phi(r', p), p] | r, \text{continuing employment})] \\ &= \phi(r, p) + \beta \cdot [V_0(p) + (1 - \delta) E(S(r', p) - S(r_0, p) | r, \text{continuing employment})] \\ &= \phi(r, p) + \beta \cdot [V_0(p) + (1 - \delta) p \cdot E(r' - r_0 | r, \text{continuing employment})], \end{aligned} \quad (\text{A5})$$

where the second equality uses $S(r_0, p) = 0$ and the third one uses the formal expression of the match surplus in (A4). The term $E(r' | r, \text{continuing employment})$ appearing in (A5) hinges on the period-to-period evolution of r , which was analyzed in the main text (see equation (5) and subsection 3.3). In particular, we showed in equation (6) that:

$$E(r' | r, \text{continuing employment}) = \bar{\varepsilon} - \int_r^{r_{\max}} \left(\bar{M}(r') - \frac{\lambda_1}{1 - \delta} \bar{M}(r')^2 \right) dr', \quad (\text{A6})$$

Substituting into (A5), we can express the worker's value function as:

$$V(\phi(r, p), p) - V_0(p) = \phi(r, p) - p \cdot \left(r_0 + \beta(1 - \delta) \int_r^{r_{\max}} \left(\bar{M}(r') - \frac{\lambda_1}{1 - \delta} \bar{M}(r')^2 \right) dr' \right). \quad (\text{A7})$$

Wage values. Combining equations (A2), (A4) and (A7), we now see that all wages have the following form:

$$\phi(r, p) = p \cdot \left(r + \beta(1 - \delta) \int_r^{r_{\max}} \left(\bar{M}(r') - \frac{\lambda_1}{1 - \delta} \bar{M}(r')^2 \right) dr' \right), \quad (\text{A8})$$

with in particular entry wages being defined by $\phi_0(p) = \phi(r_0, p)$. The model with forward-looking agents thus still predicts that log-wages are additively separable into a worker fixed effect $\ln p$ and a transitory/persistent match-specific component:

$$v = v(r) = \ln \left(r + \beta(1 - \delta) \int_r^{r_{\max}} \left(\bar{M}(r') - \frac{\lambda_1}{1 - \delta} \bar{M}(r')^2 \right) dr' \right). \quad (\text{A9})$$

This transformation of r into $v(r)$ substitutes the simpler $r \mapsto v = \ln r$ that arose in the main text from the $\beta = 0$ case. Equation (A9) still defines a monotonically increasing transformation of r , and as such implies no substantial change for the analysis in sections 3 and above. In particular, idiosyncratic shocks to the negotiation baseline r thus only impact wages through $v = v(r)$. So knowledge of the sampling and population distributions of v , denoted $F(\cdot)$ and $G(\cdot)$ respectively, is sufficient to simulate individual labor market trajectories with our model. As we saw in the main text, it is more convenient to work with the transformed negotiation baseline and match productivity shock— $v(r)$ and $v(\varepsilon)$ —than with the underlying r and ε . Specifically, as the sampling distribution of ε is $M(\cdot)$, the corresponding distribution for v is simply $F(\cdot) = M \circ v^{-1}(\cdot)$.

B Derivation of theoretical moments

Derivation of $G_s(\cdot)$. We construct $G_s(\cdot)$ by induction. Considering the conditional distribution of $v_{t+1} | v_t$ in (5), one has for any $s \geq 1$:

$$g_{s+1}(v) = \left[\bar{F}(v) - \frac{\lambda_1}{1-\delta} \bar{F}(v)^2 \right] \cdot g_s(v) - \left[f(v) - \frac{2\lambda_1}{1-\delta} f(v) \bar{F}(v) \right] \cdot G_s(v) + f(v), \quad (\text{B1})$$

which integrates as:

$$G_{s+1}(v) = \left[\bar{F}(v) - \frac{\lambda_1}{1-\delta} \bar{F}(v)^2 \right] \cdot G_s(v) + F(v). \quad (\text{B2})$$

Given the initial condition $G_0(\cdot) \equiv G(\cdot)$, this difference equation solves as (7) in the main text.

Taking up the expression in (7), it is straightforward to check that $\forall v$, $\bar{G}_\infty(v) \geq \bar{G}(v)$, i.e. $G_\infty(\cdot)$ first-order stochastically dominates $G(\cdot)$. Integration by parts further shows that:

$$E_{G_s}(v) = v_0 + \int_{v_0}^{v_{\max}} \bar{G}_\infty(x) dx - \int_{v_0}^{v_{\max}} \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2 \right)^s \cdot [\bar{G}_\infty(x) - \bar{G}(x)] dx, \quad (\text{B3})$$

which establishes that $E_{G_s}(v)$ monotonically increases with s .

Derivation of $\text{Cov}(\ln \phi_t, \ln \phi_{t+s} | \text{employment at } t, \dots, t+s)$. We begin by the derivation of the conditional expectation $E(v_{t+s} | v_t, \text{employment at } t, \dots, t+s)$. Even though this is not among the set of moments we are eventually going to directly match in the estimation, its derivation is a useful intermediate step.

It is straightforward to show using (5) that that for any differentiable function $\varphi(\cdot)$:

$$E[\varphi(v_{t+1}) | v_t, \text{employment at } t, t+1] = E_F[\varphi(v)] - \int_{v_t}^{v_{\max}} \varphi'(x) \cdot \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2 \right) dx. \quad (\text{B4})$$

Next defining $T_s(v_t) = E(v_{t+s} | v_t, \text{employment at } t, \dots, t+s)$, the conditional prediction of the negotiation baseline s periods ahead given v_t and given continuous employment between dates t and $t+s$, one notices that for any $s \geq 2$:

$$\begin{aligned} T_{s+1}(v_t) &\equiv E(v_{t+s+1} | v_t, \text{employment at } t, \dots, t+s+1) \\ &= E[E(v_{t+s+1} | v_{t+1}, \text{employment at } t+1, \dots, t+s+1) | v_t, \text{employment at } t, t+1] \\ &= E[T_s(v_{t+1}) | v_t, \text{employment at } t, t+1]. \end{aligned} \quad (\text{B5})$$

Reasoning by induction and differentiating (B4) shows that $T'_s(v_t) = \left(\bar{F}(v_t) - \frac{\lambda_1}{1-\delta} \bar{F}(v_t)^2 \right)^s$. Substituting back into (B4) and (B5) establishes the following for $T_s(v_t) = E(v_{t+s} | v_t, \text{employment at } t, \dots, t+s)$:

$$T_s(v_t) = E_F[T_{s-1}(v)] - \int_{v_t}^{v_{\max}} \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2 \right)^s dx. \quad (\text{B6})$$

Note that $E_F[T_{s-1}(v)]$ is a deterministic term depending on the ‘‘prediction horizon’’ s only. Then from

(B6) we directly obtain:⁵⁰

$$\begin{aligned}
\text{Cov}(\ln \phi_t, \ln \phi_{t+s}) &= \text{Var}(\ln p) + \text{Cov}(v_t, v_{t+s}) \\
&= \text{Var}(\ln p) + \text{Cov}(v_t, E(v_{t+s} | v_t)) \\
&= \text{Var}(\ln p) - \text{Cov}\left(v_t, \int_{v_t}^{v_{\max}} \left(\bar{F}(x) - \frac{\lambda_1}{1-\delta} \bar{F}(x)^2\right)^s dx\right). \tag{B7}
\end{aligned}$$

Note that the distributions with respect to which expectations should be taken in all these moments are distributions in the population of employed workers, meaning $G(\cdot)$ for the moments involving v_t and $H(\cdot)$ for those involving p .

⁵⁰We continue to work conditionally on continuous employment between dates t and $t + s$. However, to avoid a notational overload, we now keep this conditioning implicit.

Job spell sample: first spell					
Initial state	N. Obs.	Mean duration (months)	% censored	% job-job transitions	% job-un. transitions
Employed	659	61.6	42.4	33.8	23.8
Unemployed	58	15.2	25.9	—	—
Job spell sample: numbers of transitions					
		0	1	2	≥ 3
Percent of sample		40.8	23.9	15.0	20.3
Income sample (1992 cross-section)					
	N. Obs.	Mean		Std. Dev.	
	677	1.90 (£7.70 per hour)		0.51 (£4.96 per hour)	

Table 1: Sample descriptive statistics

$\delta (\times 100)$	$1/\delta$	$\lambda_1 (\times 100)$	$1/\lambda_1$	$\lambda_0 (\times 100)$	$1/\lambda_0$
0.323 (0.021)	309.7 (20.22)	1.481 (0.055)	67.52 (2.484)	4.108 (0.716)	24.34 (4.241)

Table 2: Transitions and mean spell durations (months)

Parameters of $G(\cdot)$					Test of OI restrictions:
$v_0 = v_{\min}$	μ	σ	$\text{Var}(\ln p)$	σ_{me}^2	Test Statistic (df: p -value)
1.779 (0.059)	2.260 (0.019)	0.202 (0.033)	0.062 (0.008)	0.012 (0.006)	28.45 (30: 0.547)
1.793 (0.060)	2.236 (0.035)	0.262 (0.025)	0.060 (0.008)	0.000 (const.)	31.78 (31: 0.428)
					$Vs. \sigma_{\text{me}}^2 \neq 0:$ 3.324 (1: 0.068)

Note: $G(v)$ is parameterized as the sum of a mass of $\frac{\delta}{\delta + \lambda_1}$ at v_0 and a normal distribution of mean μ and standard deviation σ truncated at v_{\min} .

Table 3: Distributional parameters

	Cov($\Delta \ln \phi_{i,t}, \Delta \ln \phi_{i,t+s}$), $s = \dots$				
	0	1	2	3	4
BHPS	0.033 (0.001)	-0.010 (-0.001)	0.000 (0.001)	-0.000 (-0.001)	0.000 (0.001)
Simulated	0.037 (0.001)	-0.010 (-0.001)	-0.001 (-0.001)	-0.000 (-0.001)	-0.001 (-0.001)
	ARMA parameters:			Test of OI rest.:	
	σ_ζ^2	σ_ξ^2	θ_1	Test Statistic (df: p -value)	
BHPS	0.011 (0.001)	0.010 (0.001)	0.069 (0.048)	70.68 (63: 0.237)	
	0.012 (0.001)	0.011 (0.001)	0.000 (const.)	67.73 (64: 0.351)	$Vs. \theta_1 \neq 0:$ 2.95 (1: 0.086)
Simulated	0.012 (0.001)	0.012 (0.001)	0.060 (0.037)	78.44 (63: 0.091)	
	0.013 (0.001)	0.011 (0.000)	0.000 (const.)	80.69 (64: 0.078)	$Vs. \theta_1 \neq 0:$ 2.25 (1: 0.134)

Table 4: Fitting an ARMA process

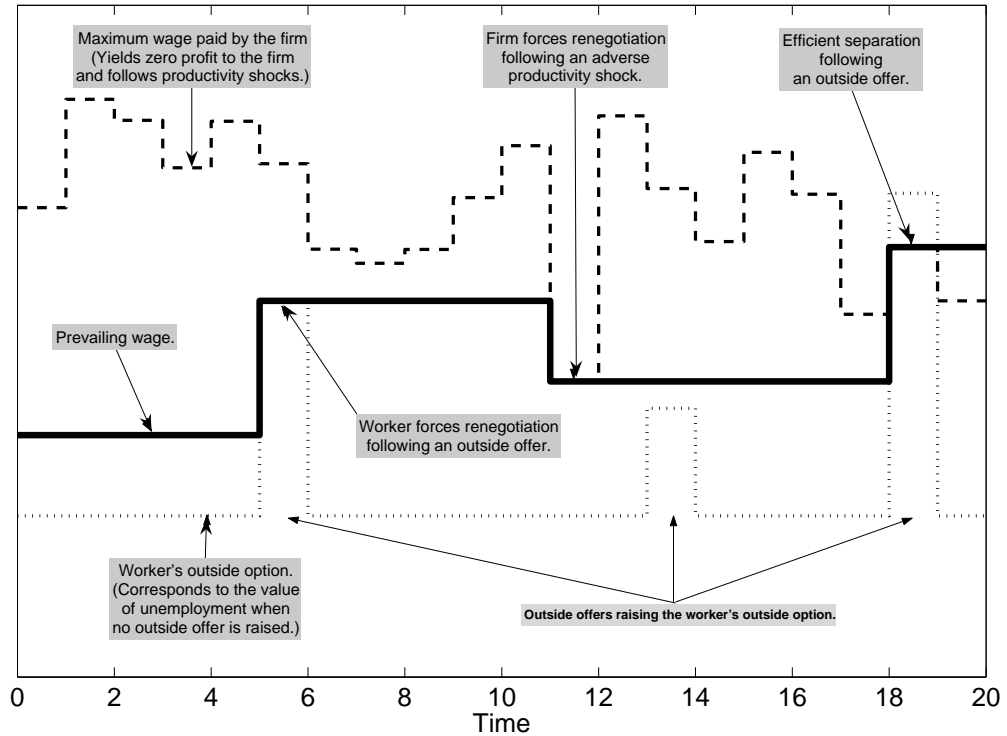


Figure 1: The wage process

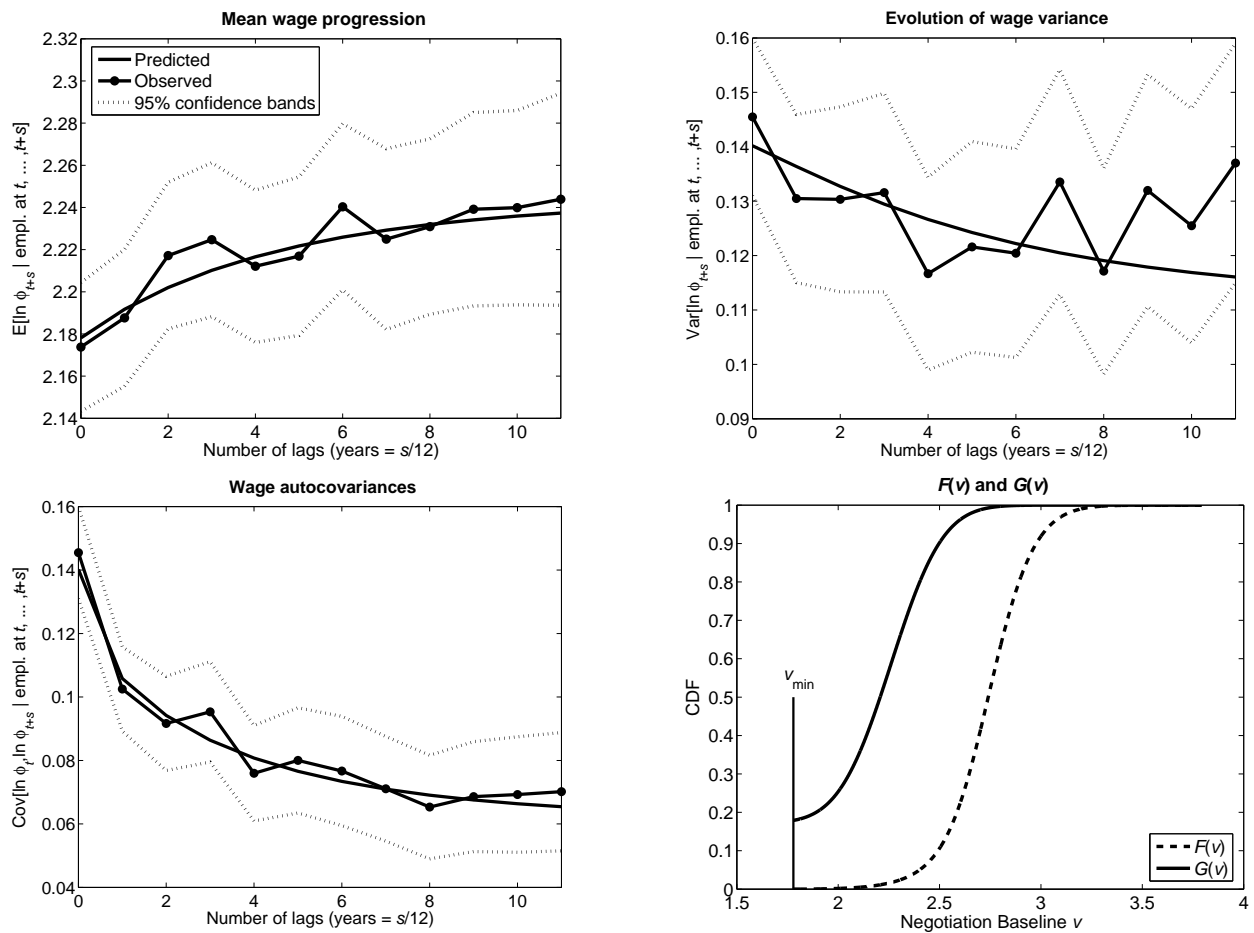


Figure 2: First- and second-order moments of the wage process

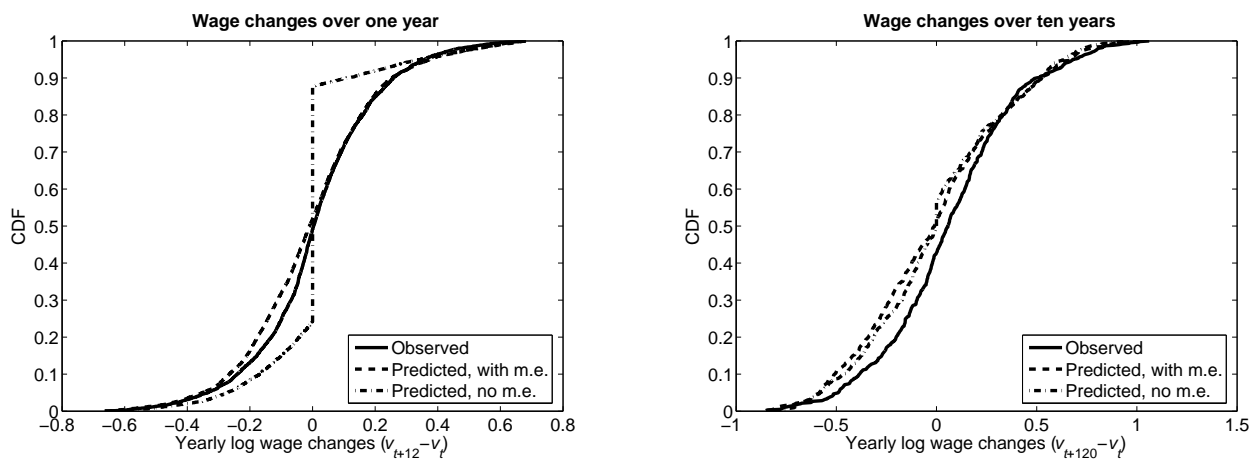


Figure 3: Distribution of year-to-year wage changes

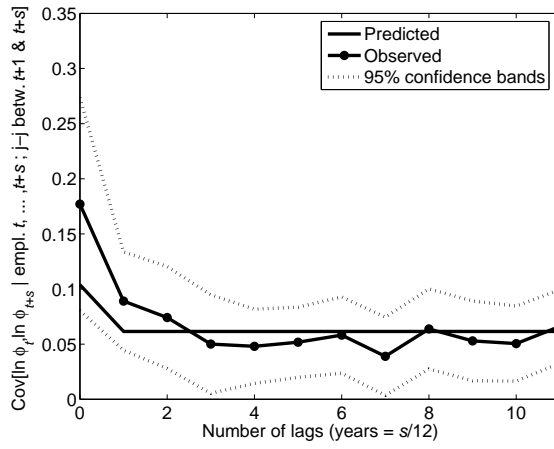


Figure 4: Wage autocovariances conditional on job-to-job mobility

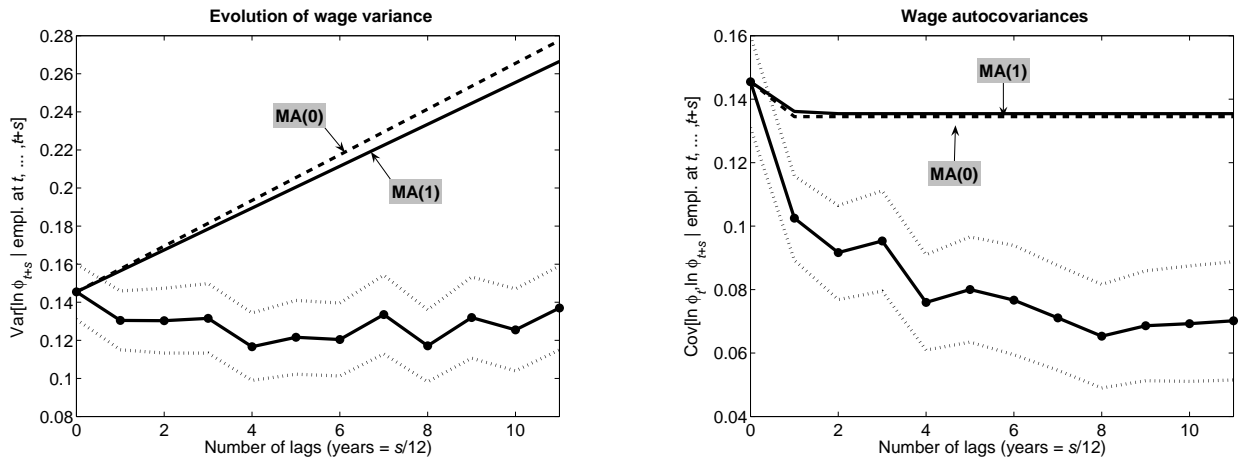


Figure 5: The wage process in levels and the ARMA decomposition