Centre for Structural Econometrics

Discussion paper No. 09/03

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

Fabien Postel-Vinay
University of Bristol
and
Paris School of Economics

Helene Turon
University of Bristol

October 2009

Centre for Structural Econometrics
University of Bristol Department of Economics
8 Woodland Road, Bristol BS8 1TN, UK.
http://www.efm.bris.ac.uk/cse/
econ-cse@bristol.ac.uk

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

Fabien Postel-Vinay†
University of Bristol
and
Paris School of Economics

Hélène Turon‡
University of Bristol

First version June 2005—This version July 2008

Abstract

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 productivity shocks. This offers an alternative to the conventional linear ARMA-type approach to modeling 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.

Keywords: Job Search, Individual Shocks, Structural Estimation, Covariance Structure of Earnings.
JEL codes: J41, J31.

*The authors thank Gadi Barlevy, Jeff Campbell, Morris Davis, Jean-Marc Robin, Rob Shimer and Randy Wright for their inspiring thoughts on this paper, and two anonymous referees for very constructive suggestions. Very useful feedback was also received from conference participants at the 2005 Essex Thanksgiving conference and the 2006 SED meeting (Vancouver), and seminar audiences at Bristol, Paris I, CREST, UCL, Queen Mary, Kent, Carnegie Mellon, Minnesota, the Chicago Fed, Toulouse and Oxford. Postel-Vinay gratefully acknowledges financial support from the ESRC (grant reference RES-063-27-0090). The usual disclaimer applies.

†Address: Department of Economics, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK. Tel: +44 117 928 8431. E-mail `Fabien.Postel-Vinay@bristol.ac.uk`. Internet `www.efm.bris.ac.uk/economics/staff/vinay/`. Postel-Vinay is also affiliated with CEPR and IZA.

‡Address: Department of Economics, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK. Tel: +44 117 928 8400. E-mail `helene.turon-lacarrieu@bris.ac.uk`. Internet `www.efm.bris.ac.uk/economics/staff/turon/`.

1 Introduction

In this paper we aim to offer a theoretical representation of observed individual earnings dynamics by investigating the capacity of a structural model of job search with simple i.i.d. productivity shocks to capture the covariance patterns 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.

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 which then 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. Under a mutual-consent rule for renegotiation (meaning that neither party can force the other to renegotiate against their will), three distinct situations 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. match quality shocks and outside jobs offers are only infrequently translated into wage shocks, hence wage shock persistence.

< Figure 1 about here. >

The intuitive idea that renegotiation by mutual consent 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.

We estimate our structural model on a sample of highly 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 dynamics. 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.[1] 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 large longitudinal dimension of our panel and the dynamic predictions of our model.

To our knowledge, our paper is the first full-fledged analysis of individual wage dynamics within a structural job search model, perhaps with the exception of Flinn (1986).[2] In that insightful

[1] See Eckstein and Van den Berg (2005) or Postel-Vinay and Robin (2006) for reviews of these arguments.

[2] We should also mention a related and independently written paper by Yamaguchi (2006), which we discuss in

paper, Flinn also offers a search-based interpretation of the observed patterns of wage covariances by combining job search with Jovanovic's (1979) model of learning about match quality. Although Flinn's empirical motivation is very much the same as ours, his approach is different. In his model wages are assumed to equal match productivity, which follows an exogenously specified process in any given match (essentially a match-specific fixed effect plus an i.i.d. shock). It is the workers' endogenous turnover decisions (driven by the workers' learning about the fixed component of match quality) that "filter" the exogenously specified productivity process into sequences of wages that are consistent with NLSY data. We, on the other hand, completely abstract from learning and keep the modeling of turnover decisions to a minimal level of simplicity. Our focus is on the division of match rents as a mechanism to explain wage dynamics, something that was left essentially unmodeled in Flinn's paper (see Flinn, 1986, footnote 2). As such we view our contribution as complementary to Flinn's.

Another and perhaps more general contribution of our paper is to offer a structural counterpart to the large empirical literature on individual labor income processes (as Flinn, 1986, also does), in which sophisticated stochastic processes (typically, but not exclusively, ARMA-type processes) are fitted to longitudinal wage data taken from worker or household panels.[3] Because it is primarily based on statistical models, that literature remains somewhat out of touch with economic theory. Although links to behavioral interpretations of the highlighted autocovariance patterns are sometimes informally discussed, quantitative consistency with a formal behavioral model is typically not addressed.

It seems important, however, to understand the economic forces governing individual wage dynamics from a dual theoretical and quantitative standpoint.[4] This paper suggests that the

Section 6.

[3]That 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.

[4]The following examples may illustrate that claim: A typical application of the empirical literature on wages 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.

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.

Finally, 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,[5] 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 provides neither of the matched partners with 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.[6]

The rest of the paper is organized as follows. In the next section we pose the theoretical model.

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

[6] Mutual 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.

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. Section 6 tackles the issue of persistent productivity shocks within our structural model. Finally, we conclude and discuss a number of potentially interesting extensions in Section 7.

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 sold in a perfectly competitive market. Workers and firms are infinitely lived, forward-looking, risk-neutral and have a common exogenous per-period discount factor of β. Time is discrete and the economy is at a steady state. Workers are either unemployed or matched with a firm. Firms operate constant-return technologies and are modeled as a collection of job slots which are either 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 *period- and match-specific* shock ε_t. We should emphasize that because this 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 quality 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, throughout most of this paper we assume that transitory shocks ε_t are uncorrelated over time or across matches, implying that any source of persistence in a worker's productivity either over time or across jobs is picked up by the worker fixed effect, p. We opt for this admittedly disputable assumption based on the following two arguments. First, it greatly simplifies the analytical characterization of wage dynamics. Indeed our relatively simple, closed-form characterization of wage dispersion and wage dynamics (see subsection 3.3), which we view as an attraction of our approach, becomes impossible under any realistic pattern of serially correlated shocks, with otherwise little gain in terms of new qualitative insights for our purposes.[7] Second, commonly available worker- or household-level data of the type routinely used in the analysis of individual earnings dynamics typically do not convey any direct information on productivity, which poses difficulties in the identification of the productivity process. Although a combination of the model's structure and the wage information contained in the data would in principle ensure identifiability of some measure of persistence in productivity, we will argue later in the paper that this is not the case in practice, at least with the BHPS data that we are using. We defer the discussion of all those issues until section 6, where we offer a simple way to relax the assumption of i.i.d. shocks and attempt to estimate the resulting slightly generalized version of our model.

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.

[7] Arguably the simplest way to increase persistence (while preserving serial independence) of match quality 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. Yamaguchi (2006) takes this route in a closely related (and independently conducted) research, and finds it intractable to solve the model explicitly. He thus resorts to indirect inference for estimation.

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]

We denote the worker's valuation of this match by $V(\phi, p)$, and the firm's valuation by $\Pi(\varepsilon, \phi, p)$. $V(\cdot)$ and $\Pi(\cdot)$ are 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 further assume 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), and we denote the lifetime value of unemployment by $V_0(p)$.

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

$$S(\varepsilon, p) = [V(\phi, p) - V_0(p)] + [\Pi(\varepsilon, \phi, p) - 0]. \tag{2}$$

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.[9] Moreover, total match surplus only depends on the determinants of current and future match output flows. Given the assumed jump process for transitory shocks ε, those only include the permanent component of match productivity, p, and the current value of its transitory component, ε.

Job search, match formation and match dissolution. 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.

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

[9] Intuitively, total match surplus involves the present discounted sum of expected future flow values of match surplus, which in turn are the sum of wage flows net of foregone unemployment income flows $(\phi - b \cdot p)$ plus net profit flows $(p \cdot \varepsilon - \phi)$. As the ϕ terms cancel when net income and profit flows are added, surplus flows are independent of any wage value.

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. Match breakup is thus formally disconnected from idiosyncratic shocks to ε. This assumption calls for the following comments. The formation or continuation of any firm-worker match is subjected to the minimal requirement that total match surplus be nonnegative. Thus, implicit in the specifications detailed above is the assumption $S(\varepsilon_{\min}, p) \geq 0$ for all p, which amounts to truncating the "true" underlying distribution of productivity shocks from below.[10] However an alternative, probably more general view on the 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

The wage rule. Wage contracts stipulate a constant wage and are only renegotiable by mutual consent in continuing matches. In other words, no firm or worker can force their match partner to revise the wage against the latter's interest, unless the former has a credible threat to leave the match.[11] The implications of this wage rule for wage dynamics (of which we gave an informal account in Figure 1 in the introduction) were analyzed theoretically by MacLeod and Malcomson (1993): In continuing matches, if one party has 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

[10] In fact the condition $S(\varepsilon_{\min}, p) \geq 0$ will turn out not to depend on p, as the surplus will be shown to be proportional to p—see below. Let us then define r_0 by $S(r_0, p) = 0$ and 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 r_0 yields a positive potential match surplus, while any draw falling short of r_0 entails a negative surplus and causes match dissolution. Hence the match destruction rate δ can be seen as equaling $M_0(r_0)$, while $M(\cdot)$ simply coincides with $M_0(\cdot)$ truncated below at r_0.

[11] See Malcomson (1997, 1999) for a motivation of this principle.

other party consents to wage renegotiation up or down to the point where this outside option is matched.[12] The existing match only survives 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.[13]

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 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.[14] The latter equals $V_0(p)$ 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. At any time, the wage that the worker receives, ϕ, can be expressed as a function of her/his type p and a value of match-specific productivity r (which will not in general be equal to match-specific productivity in the current match, ε), defined as follows:

$$\phi = \phi(r,p) \quad \Leftrightarrow \quad V(\phi,p) = V_0(p) + S(r,p) \quad \Leftrightarrow \quad \Pi(\varepsilon,\phi,p) = S(\varepsilon,p) - S(r,p), \tag{3}$$

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

[13] Even though we appeal to MacLeod and Malcomson (1993) as a theoretical foundation of our wage setting mechanism, alternative justifications exist. 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 deliver a similar wage rule (Harris and Holmström, 1982; Thomas and Worral, 1988). More generally, some degree of worker risk aversion (coupled with an inability to transfer wealth across time) can be appealed to to justify the focus on constant-wage contracts. Finally, Hall (2005) also considers this wage setting mechanism which he interprets as a social norm favoring wage rigidity while restricting wages to lie within the bargaining set (thus avoiding inefficiencies in the allocation of labor).

[14] In terms of a Nash bargaining approach, we thus 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. At the other extreme, Flinn (1986) assumes a worker bargaining power of 1, whereby s/he gets paid the full marginal product at all times.

where r gives a measure of the surplus that the worker enjoys over and above the value of being unemployed, $V_0(p)$. Alternatively, r can be seen as the quality of the match from which the worker was last able to extract the whole surplus. For reasons that will become clear shortly, we will term r the worker's *negotiation baseline*.

The negotiation baseline is defined formally in the next three paragraphs. Before we proceed, however, it is important to note that $\phi(r,p)$ is a strictly increasing function of r. This flows directly from the monotonicity properties of $V(\cdot)$ and $S(\cdot)$, which together with (3) imply $\frac{\partial \phi}{\partial r} = \left(\frac{\partial S}{\partial r}\right) / \left(\frac{\partial V}{\partial \phi}\right) > 0$.

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. This case is encompassed by equation (3) and entails a starting wage equal to $\phi(r_0,p)$, and a negotiation baseline of r_0 such that:

$$V(\phi(r_0,p),p) = V_0(p) \quad \Leftrightarrow \quad S(r_0,p) = 0. \tag{4}$$

A formal definition of r_0 based on (4) will be given below, where we will establish that r_0 is independent of p. For now we should note that, in any generic match with quality ε and negotiation baseline r, $\varepsilon \geq r \geq r_0$ 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.

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 $\phi(\varepsilon, p)$ (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 $\phi(\eta, p)$. 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, $\phi(\eta, p)$. 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:

$$\begin{cases} \eta > \varepsilon & \text{Worker quits, mobility wage } \phi(\varepsilon, p)\text{, new negotiation baseline } \varepsilon, \\ \varepsilon > \eta > r & \text{Worker stays, renegotiated wage } \phi(\eta, p)\text{, new negotiation baseline } \eta, \\ r > \eta & \text{Offer is discarded, nothing changes.} \end{cases}$$

Productivity shocks. A last potential cause of wage change is the occurrence of a productivity shock. 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.[15]

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 $V_0(p)$, equivalent to a negotiation baseline of r_0. This is never preferable to keeping the existing contract which has a negotiation baseline of $r \geq r_0$. 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 $S(\varepsilon, p) - S(r, p)$ to $S(\varepsilon', p) - S(r, p)$. 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 $\phi(r, p)$. At that point the firm's only outside option is to fire the worker, thus

[15] For 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.

ending up with a vacant job worth 0, while carrying 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 $\varepsilon_{\min} \geq r_0$, the match is still viable, meaning that a mutually beneficial contract exists. However, keeping the existing wage $\phi(r, p)$ would imply a negative profit of $S(\varepsilon', p) - S(r, p)$. 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 rule then implies 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 $\phi(\varepsilon', p)$ in the continuing match. Her/his negotiation baseline has thus been updated to ε'.

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 matches and in 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, wages are paid and production takes place.

2.3 Value Functions and Wages

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:

$$V_0(p) = b \cdot p + \beta \cdot V_0(p) \qquad \Leftrightarrow \qquad V_0(p) = \frac{b \cdot p}{1-\beta}. \tag{5}$$

Total match surplus $S(\varepsilon, p)$**.** As we saw earlier (footnote 9), 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 $p \cdot (\varepsilon - b)$. It does not depend on any wage value.

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, 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 above, 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:

$$S(\varepsilon, p) = p \cdot (\varepsilon - b) + \beta(1-\delta) \cdot \int_{\varepsilon_{\min}}^{\varepsilon_{\max}} S(\varepsilon', p)\, dM(\varepsilon'), \tag{6}$$

where $\overline{\varepsilon} = E_M(\varepsilon)$. $S(\cdot)$ is therefore proportional to p and independent of any wage value.

Before going any further, it is worth looking at the negotiation baseline r_0 that workers start with when they leave unemployment. By its definition (4), r_0 satisfies $S(r_0, p) = 0$ and thus:

$$r_0 = b - \frac{\beta(1-\delta)}{1-\beta(1-\delta)}(\overline{\varepsilon} - b). \tag{7}$$

Note that $r_0 < b$ provided that $\overline{\varepsilon} > b$, which is a necessary condition for any trade at all to take place in the labor market.

Worker values $V(\phi(r,p), p)$ and wages $\phi(r,p)$. 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:

$$V(\phi(r,p), p) = \phi(r,p) + \beta \cdot \left[\delta V_0(p) + (1-\delta) E\left(V\left(\phi\left(r', p\right), p\right) \mid r, \text{continuing employment}\right)\right], \tag{8}$$

The definition (3) also states that $V(\phi(r,p), p) = V_0(p) + S(r,p)$, which implies:

$$\begin{aligned}
\phi(r,p) &= V_0(p) + S(r,p) - \beta \cdot \left[V_0(p) + (1-\delta) E\left(S\left(r', p\right) \mid r, \text{continuing employment}\right)\right] \\
&= p \cdot \left[r - \beta(1-\delta)\left(E\left(r' \mid r, \text{continuing employment}\right) - \overline{\varepsilon}\right)\right] \\
&\overset{\text{def.}}{=} p \cdot \varphi(r),
\end{aligned} \tag{9}$$

where the second equality uses (5), (6) and (7) to substitute for $V_0(\cdot)$, $S(\cdot)$ and r_0, respectively. The term $E(r' \mid r, \text{continuing employment})$ appearing in (9) hinges on the period-to-period evolution of r, which will be analyzed in the next section. However, what is most important about (9) is that it establishes that all wages are multiplicatively separable in the worker's type p and a function of the negotiation baseline, $\varphi(r)$. A closed-form expression of $\varphi(\cdot)$ can be derived (see below footnote 25), however the only property of $\varphi(\cdot)$ that really matters for what follows is that it is a strictly increasing function.

3 Model Solution and Econometric Inference

Our aim is to estimate the model parameters with a standard and readily available 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

Our model's very simple structure makes it particularly easy to estimate the set of transition parameters in a first stage, independently of the rest of the model. The key property that facilitates this decomposition is that the only dimension of heterogeneity that impacts worker turnover is ε, the transitory productivity shock, which is i.i.d. across periods and matches, and as such can very easily be integrated out of the likelihood of labor market spell durations.[16] 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 $\frac{\lambda_1}{2}$.[17] 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] While analytically convenient for our estimation purposes, we should acknowledge that this property has some theoretically unappealing consequences. One is that the hazard rates of job destruction, job finding and job-to-job move are 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 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 unfortunate disconnect. Another possibility would be to allow for persistent productivity shocks. We discuss these possible extensions in Section 7.

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

3.2 Wage Distributions

The model laid out in section 2 implies that all wages have a multiplicative form $\phi(r,p) = p \cdot \varphi(r)$. It thus predicts that log-wages are additively separable into a worker fixed effect $\ln p$ and a transitory/persistent match-specific component $v = \ln \varphi(r)$. Idiosyncratic shocks to the negotiation baseline r therefore only impact wages through the monotonically increasing transformation $r \mapsto v = \ln \varphi(r)$. So knowledge of the sampling and population distributions of v, denoted $F(\cdot)$ and $G(\cdot)$ respectively, is sufficient to characterize or simulate individual labor market trajectories. Indeed it will prove more convenient to work with the *transformed* negotiation baseline and match productivity shock—$\ln \varphi(r)$ and $\ln \varphi(\varepsilon)$—than with the underlying r and ε. Specifically, as the sampling distribution of ε is $M(\varepsilon)$, the corresponding distribution for $v = \ln \varphi(\varepsilon)$ is simply $F(v) = M\left[\varphi^{-1}(e^v)\right]$.

The dynamics of $v = \ln \varphi(r)$ will be characterized in the next subsection. 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 (transformed) negotiation baseline less than $v = \ln \varphi(r)$. Workers exit this pool either if their match has been dissolved (probability δ) or if their new negotiation baseline is greater than r (probability $\lambda_1 \overline{F}(v)^2$).[18] 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 $(1-\delta) F(v)$.[19]

We can thus now write the balance of flows in and out of the stock $(1-u)G(v)$:

$$(1-u)G(v)\left(\delta + \lambda_1 \overline{F}(v)^2\right) = (1-\delta)F(v) \cdot (1-u)\overline{G}(v) + \lambda_0 \cdot u$$
$$\iff \quad G(v) = \frac{\delta + (1-\delta)F(v)}{\delta + (1-\delta)F(v) + \lambda_1 \overline{F}(v)^2}. \tag{10}$$

[18] As observing an increase in the negotiation baseline requires the worker to receive an outside offer, and productivity draws that are above her/his current negotiation baseline at both the poaching and the incumbent firm.

[19] As, conditional on remaining employed, workers will have a new negotiation baseline lower than r if their current-period draw of the idiosyncratic shock falls short of r, whether they raise an outside offer or not.

Note the existence of a mass at $v_0 = \ln \varphi(r_0)$, $G(v_0) = \frac{\delta}{\delta+\lambda_1}$ due to the unemployed workers all being hired at the minimum negotiation baseline v_0. More precisely, we can decompose $G(\cdot)$ as:

$$G(v) = \frac{\delta}{\delta+\lambda_1} \cdot \mathbf{1}_{\{v \geq v_0\}} + \left[\frac{\delta + (1-\delta) F(v)}{\delta + (1-\delta) F(v) + \lambda_1 \overline{F}(v)^2} - \frac{\delta}{\delta+\lambda_1} \right] \cdot \mathbf{1}_{\{v \geq v_{\min}\}}, \tag{11}$$

where $v_{\min} = \ln \varphi(\varepsilon_{\min})$ 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 $\lambda_1 = 0$).[20]

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 of job offers and job destruction shocks δ, which we saw are 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 $v_t = \ln \varphi(r_t)$, the worker's current (transformed) 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 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 assumptions concerning the wage setting rules then

[20] Incidentally we may emphasize that, contrary to most job search models that have a "job-ladder" design and no 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 ladder and only one wage 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.

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}.[21]

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 negotiation baseline equal to $\max\langle r_t, \min\{\eta_{t+1}, \varepsilon_{t+1}\}\rangle$.[22]

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

$$v_{t+1} \mid v_t = \begin{cases} v_t & \text{with probability } (1-\delta)\,\overline{F}(v_t) - \lambda_1 \overline{F}(v_t)^2, \\ v' < v_t & \text{with density } (1-\delta)\, f(v'), \\ v' > v_t & \text{with density } 2\lambda_1 f(v')\,\overline{F}(v'), \end{cases} \tag{12}$$

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, nonlin-

[21]The 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}.

[22]The 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 $\varepsilon_{t+1} \geq \eta_{t+1} > r_t$.

ear Markovian process based on a specific acceptance/rejection scheme of i.i.d. wage innovations.[23] 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 (12)—and its differences with the conventional linear ARMA specification—will be analyzed in section 5.2. For the time being, we derive the following moment which will be useful for estimation. Integration of (12) implies that, conditional on employment at two consecutive dates t and $t+1$:

$$E\left(v_{t+1} \mid v_t, \text{ employment at } t, t+1\right) = v_t + \frac{\lambda_1}{1-\delta}\int_{v_t}^{v_{\max}} \overline{F}(x)^2\, dx - \int_{v_{\min}}^{v_t} F(x)\, dx.$$

This expression shows that conditional expected wage growth—i.e. $E\left(v_{t+1} - v_t \mid v_t, \text{ employment at } t, t+1\right)$—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 $F(\cdot)$).[24, 25]

[23] Incidentally, our 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 (12) is indeed formally reminiscent of predictions obtained by Harris and Holmström (1982) and Thomas and Worral (1988) in models of self-enforcing wage contracts. See the discussion in footnote 13.

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

[25] Incidentally, equation (13) can be rewritten in terms of the initial negotiation baseline $r = \varphi^{-1}(e^v)$ as follows:

$$E\left(r' \mid r, \text{continuing employment}\right) = \overline{\varepsilon} - \int_r^{\varepsilon_{\max}} \left(\overline{M}(r') - \frac{\lambda_1}{1-\delta}\overline{M}(r')^2\right) dr'.$$

Going back to (9), this leads to the following closed-form expression of the function $\varphi(r)$ and any wage $\phi(r,p)$:

$$\phi(r,p) = p \cdot \varphi(r) = p \cdot \left(r + \beta(1-\delta)\int_r^{\varepsilon_{\max}} \left(\overline{M}(r') - \frac{\lambda_1}{1-\delta}\overline{M}(r')^2\right) dr'\right).$$

Dynamics over s periods. We now consider the cross-sectional distribution of v_{t+s} conditional on employment at dates $t, \cdots, t+s$, 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 the Appendix:

$$G_s(v) = \left[1 - \left(1 - F(v) - \frac{\lambda_1}{1-\delta}\overline{F}(v)^2\right)^s\right] \cdot G_\infty(v) + \left(1 - F(v) - \frac{\lambda_1}{1-\delta}\overline{F}(v)^2\right)^s \cdot G(v), \quad (13)$$

where

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

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 the Appendix) 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 employment in that the mean negotiation baseline increases with the duration of continuous employment. Intuitively, a worker's wage increases in expected terms with the number of job offers received by that worker since s/he was last unemployed. As one looks at individuals that have been continuously employed for longer periods, the average number of job offers received by these individuals since they got out of unemployment (and hence the average wage or negotiation baseline) gradually increases. This selection effect is the driving force behind the increase in the mean wage with s.

Based on the definition (13), 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 the Appendix that:[26]

$$\text{Cov}(\ln\phi_t, \ln\phi_{t+s}) = \text{Var}_H(\ln p) - \text{Cov}_G\left(v_t, \int_{v_t}^{v_{\max}} \left(1 - F(x) - \frac{\lambda_1}{1-\delta}\overline{F}(x)^2\right)^s dx\right). \quad (14)$$

[26]We 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.

(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 toward 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 (13) and $T-1$ covariances from (14) on which to base an estimation of the $F(\cdot)$ distribution and the variances of the measurement error and fixed-effect distributions.

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.[27] 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.

Our working sample is obtained from the following selection of the raw data. First, we drop the few individuals that have gaps in their records. We then 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.[28]

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

[28] Ideally, 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.

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 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 599 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 and a series of third moments of the distribution of wage changes, as derived in subsection 3.3.[29] Specifically, we match the following $3T-1$ moments (where T is

[29]Note that we shall only estimate $F(\cdot)$—the sampling distribution of *transformed* match quality shocks $v = \ln\varphi(\varepsilon)$—, not the underlying sampling distribution of match quality shocks $M(\cdot)$. As we already mentioned, knowledge of $F(\cdot)$ is sufficient to simulate wages from our model. This also explains why we will not need an estimate of the discount factor β. Indeed the wage process derived from our theory—see (12)—is independent of β, which only affects the function $\varphi(\cdot)$, i.e. the way in which match quality 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. when they occur), knowledge of β is unnecessary for the identification of the parameters of our wage process. It would only matter if

the panel length in years), for $s = 0, 12, 24, \cdots, 12(T-1)$:[30]

$$
\begin{aligned}
&E(\ln \phi_{t+s}) = E_{G_s}(v) \\
&\mathrm{Var}(\ln \phi_{t+s}) = \mathrm{Var}_H(\ln p) + \mathrm{Var}_{G_s}(v) + \sigma^2_{\mathrm{me}} \\
&\mathrm{Cov}(\ln \phi_t, \ln \phi_{t+s}) = \mathrm{Var}_H(\ln p) - \mathrm{Cov}_G\left(v_t, \int_{v_t}^{v_{\max}} \left(1 - F(x) - \frac{\lambda_1}{1-\delta}\overline{F}(x)^2\right)^s dx\right).
\end{aligned}
\tag{15}
$$

All these moments are conditional on continuous employment between times t and $t+s$—that is, we will match those theoretical predictions of $E(\ln \phi_{t+s})$, $\mathrm{Var}(\ln \phi_{t+s})$ and $\mathrm{Cov}(\ln \phi_t, \ln \phi_{t+s})$ to the corresponding empirical moments computed on the subsample of individuals observed to be continuously employed between t and $t+s$.[31] 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 and fourth lines of (15) the addition of a term σ^2_{me} to the theoretical expression of cross-sectional income variances. This accounts for the presence of classical measurement error (with variance σ^2_{me}) in hourly wages.

We match these moments using Equally Weighted Minimum Distance (EWMD) estimation.[32]

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 (10). As detailed in equation (11) 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

we wanted to recover the distribution of match quality shocks implied by our model and by the earnings dynamics observed in our data.

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

[31] This 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 (15) conditional on employment at t and $t+s$ only. However, the corresponding formulæ 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.

[32] Optimal minimum distance (OMD) estimation results are available upon request. They are qualitatively very similar to those reported here. As suggested by a referee, we report EWMD results in the main text as OMD estimates are fraught with a well-known small-sample bias problem (Altonji and Segal, 1996).

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 10).

As to the distribution of worker fixed effects $H(\cdot)$, we see that only its variance, $\mathrm{Var}_H(\ln p)$ appears in the series of moments we aim to match. The distribution $H(\cdot)$ can however be retrieved from our estimate of $G(\cdot)$ and the actual wage distribution by deconvolution. Yet knowledge of $\mathrm{Var}_H(\ln p)$ 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 unemployment spell durations observed in our BHPS sample.[33]

[33] Unemployment 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.

Distributions. Table 3 contains the EWMD estimates and standard errors of the various distributional parameters involved in our moment conditions (15).[34]

< 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^2_{me}, the variance of the measurement error. Equality to zero is only borderline rejected (a Wald test produces a p-value of 0.023 against the alternative $\sigma^2_{\text{me}} = 0.016$, our point estimate), although the point estimate of 0.015 — or 10.8% of the cross-sectional wage variance — is of reasonable magnitude. For comparison, in the bottom row of Table 3, we also report parameter estimates obtained subject to the constraint $\sigma^2_{\text{me}} = 0$. Comparison of the two rows shows that the parameters affected by this constraint is σ, which is estimated higher in the absence of measurement error. Removing the possibility of a measurement error leads us to estimate a variance of the negotiation baseline v to account for 56% of the wage variance, as opposed to 44% in the base estimation.

The minimum negotiation baseline, $v_0 = v_{\min}$ that workers start off with when first hired from unemployment is estimated at 1.80, 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 2.18. 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%.[35]

The variance of the worker fixed-effect distribution, $\text{Var}_H(\ln p)$, is 0.063. Taking the variance of the measurement error to equal 0.015, 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 45% for $\ln p$, about 44% for v_t, and

[34] The reported standard errors are based on 500 bootstrap replications of *both stages* of our estimation procedure on 500 resamples with replacement. As such they *do account* for the presence of nuisance parameters δ and λ_1, which appear in the theoretical moments and which we fix to their first-stage estimated value in our second estimation step.

[35] The model further has it that v_0 should equal the mean log wage among workers who were hired out of unemployment 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 very close with the estimate of v_0 shown in Table 3.

11% for the measurement error.

Finally, as we use 12 years of data, (15) defines 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 0.78, the model passes this specification test.

5.2 Simulations and Fit Analysis

Matched moments. The first three panels of Figure 2 displays the moments listed in (15) 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 $12T$ $(= 132)$ 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 11 years or more.[36] 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 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.

The bottom-left panel graphs wage autocovariances at 0 to $12T$ lags (again conditional on

[36] These individuals account for 37.9% of our sample.

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. The bottom-right panel graphs the third moment of the distribution of wage changes over an increasing horizon, i.e. $(\ln \phi_{t+s} - \ln \phi_t)^3$, with $s = 12, 24, .., 12T$. Our knowledge of the distribution G_s (see equation (13)) indeed allows us to predict any moment of this distribution. We show in the Appendix the detail of the derivation of the third moment of wage changes. The graph shows that, although not matched in the estimation, this series of moments is well predicted by our model.

That having been said, our structural model permits a much finer analysis of predicted 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 $\sigma^2_{\text{me}} = 0$ and $\sigma^2_{\text{me}} = 0.015$ for the variance of the measurement error.

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). The solid line is the empirical cdf of log wage changes. The dashed line is the model-predicted counterpart of that cdf $\sigma^2_{\text{me}} = 0.015$. Finally, the dash-dot line with a mass at zero is the cdf of simulated wage changes if one removes measurement error from the model.

First comparing the "observed" and "predicted, with measurement error" distributions of yearly wage changes, 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. 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.

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 60% 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 14%, whereas the corresponding figure for wage cuts is about a 27%.

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 $1 - \overline{F}(\cdot)^2$. 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:[37]

$$E\left(\ln\phi_t \mid \text{job-to-job transition at } t\right) = E_{1-\overline{F}^2}(v)$$

$$\text{Var}\left(\ln\phi_t \mid \text{job-to-job transition at } t\right) = \text{Var}_H(\ln p) + \text{Var}_{1-\overline{F}^2}(v) + \sigma^2_{\text{me}}$$

$$\text{Cov}\left(\ln\phi_t, \ln\phi_{t+s} \mid \text{job-to-job transition between } t+1 \text{ and } t+s\right) = \text{Var}_H(\ln p). \qquad (16)$$

Computation of the first two moments in (16) 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.[38] 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 2.62 and 0.11. 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 (16) 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 $s = 0, 12, 24, \cdots, 12(T-1)$ months) among workers who have

[37] Of 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 (10). (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.

[38] The 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,[39] 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.[40] 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 (2008) 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.

[39] These numbers range from 21 to 60, depending on s.

[40] A straightforward proof of this claim relies on flow-balance equations similar to (10). Details are available on request.

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

$$\begin{cases} \ln \phi_{i,t} = z_i + a^P_{i,t} + a^T_{i,t}, & \\ a^P_{i,t} = a^P_{i,t-1} + \zeta_{i,t}, & \text{with } \zeta_{i,t} \text{ i.i.d.,} \\ a^T_{i,t} = \sum_{\ell=0}^{q} \theta_\ell \, \xi_{i,t-\ell}, & \text{with } \xi_{i,t} \text{ i.i.d. and } \theta_0 = 1. \end{cases} \tag{17}$$

Here log-earnings are modelled as the sum of an individual fixed-effect z_i, a permanent earnings shock $a^P_{i,t}$ following a martingale process, and a transitory earnings shock $a^T_{i,t}$ 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^2_ζ and σ^2_ξ.

We fit model (17) separately to our income sample from the BHPS and to a sample of simulated data based on the parameter estimates obtained above.[41] The order of the MA process is determined by looking at the sequence of autocovariances of first-differenced log wages, $\text{Cov}\left(\Delta \ln \phi_{i,t}, \Delta \ln \phi_{i,t+s}\right)$, 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 was expected given the good fit to these covariances 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 (17) as the sum of a random walk component and a serially uncorrelated—or MA(0)—component. However, for completeness we estimate model (17) under both the MA(0) and MA(1) specifications.[42]

< Table 4 about here. >

[41] We 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 (17)—see Meghir and Pistaferri (2004).

[42] Results 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 EWMD matching of the autocovariance structure of yearly wage growth. Details are available on request.

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.[43] 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 (according to our theoretical model, earnings have a steady-state distribution characterized in (10), which has a finite variance). 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 (17) 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:

$$\begin{aligned} \operatorname{Var}(\ln \phi_{i,t+s}) &= \operatorname{Var}\left(z_i + a_{i,t}^P\right) + s\sigma_\zeta^2 + \left(1+\theta_1^2\right)\sigma_\xi^2, \\ \operatorname{Cov}(\ln \phi_{i,t}, \ln \phi_{i,t+s}) &= \operatorname{Var}\left(z_i + a_{i,t}^P\right) + \theta_1\sigma_\xi^2 \times \mathbf{1}_{\{s=1\}}. \end{aligned} \tag{18}$$

Figure 5 shows the fit of (18) 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

[43] Wald tests further indicate that both specifications are accepted based on either sample at the five percent level (last column of Table 4). Moreover, neither sample allows rejection of the MA(0) against the MA(1) specification at conventional levels.

[44] We are only reporting second-order moments to save on space. For means, the simple ARMA decomposition (17) postulates that $E(\ln \phi_{t+s})$ 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 (17). Also we should mention that the intercept term in (18), $\operatorname{Var}\left(z_i + a_{i,t}^P\right)$, was set to fit the cross-sectional wage 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.

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 Productivity Shock Persistence

The above results show that the covariance pattern of wages—most notably the degree of wage persistence—can be accounted for by the mechanisms embedded in our simple model. Those results were obtained under the maintained assumption of i.i.d. productivity shocks, i.e. from a model where wage persistence arises solely from the mutual consent rule in wage renegotiation and the (in)frequency of outside job offers.

Clearly a fuller empirical exercise would consist of decomposing the observed persistence of wages into components of their dynamics arising from shocks to match productivity and job mobility shocks (outside offers). As argued before, however, this idea is fraught with two main problems. One is the general lack of analytical tractability of our model under non-i.i.d. productivity shocks (see the discussion below). The other and more substantive one is that the typical earnings data available to researchers (such as the BHPS data used in this paper) conveys no direct information on productivity. It may then seem more natural, given this lack of direct information, to assume away persistence in productivity rather than estimating it through its manifestation in wage persistence.

It is nonetheless true that our model remains well-defined—if considerably more cumbersome—without the i.i.d. assumption on productivity shocks. Moreover, as we argue below, the model has a tight enough structure to theoretically allow identification of at least some simple measure of productivity persistence from observed wage processes. We briefly pursue these ideas in this

[45] This 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 (17).

[46] For instance Meghir and Pistaferri (2005) show that the variance of innovations in a decomposition similar to (17) follow a relatively complex time pattern.

section.

To that end, we must go some way toward relaxing our assumption of i.i.d. productivity shocks. Instead of assuming that these shocks occur every period, we assume that the match productivity is re-drawn from distribution $F(\cdot)$ with probability τ and stays at its previous value with probability $1-\tau$. Pre- and post-shock values of match productivity are still independent.

To keep the model analytically tractable, we further assume that the arrival of an outside job offer is always accompanied by a shock to idiosyncratic match quality. This is clearly an *ad hoc* assumption, and we certainly do not defend it as particularly theoretically appealing. It has, however, the advantage of keeping the model tractable while at the same time introducing persistent productivity shocks in a simple, if somewhat constrained, way. We view this simple first pass at tackling the issue of productivity persistence as informative at least in the following way: if identification of the (scalar) measure of productivity persistence, i.e. $1-\tau$, fails in this simple case, then any hope of identifying more complex or theoretically appealing patterns of persistence from the kind of worker-level data that is typically used in the analysis of individual earnings dynamics should probably be abandoned.

The crucial theoretical simplification that comes with our assumption that outside offers are always accompanied by a shock to ε is that current match productivity does *not* become an additional state variable in the worker's choice problem (as it would under any different assumption, making the model analytically intractable).[47] Here, every time the worker is in a position to compare surpluses between her/his incumbent employer and a potential poacher, the idiosyncratic quality of these two matches are two independent draws from the distribution $M(\cdot)$. As a consequence the worker's value of employment continues to be independent of the current level of match productivity (as was obviously the case under i.i.d. productivity shocks), implying that there is no option value of being employed in a match with a higher productivity.

[47] Analytically tractable models in Postel-Vinay and Robin (2002), Dey and Flinn (2005) and Cahuc et al. (2006) have fixed match- or firm-specific heterogeneity and no transitory productivity shocks. However those models predict implausibly (downward-)rigid wages within a job spell. Match productivity in Flinn (1986) is modeled as the sum of a match fixed effect plus an i.i.d. shock, yet the model assumes that wages equal productivity at all dates so that the wage process is exogenously specified in that paper. Finally, Yamaguchi (2006) and Lise, Meghir and Robin (2008) introduce various forms of persistent productivity processes in models very similar to ours, but have to resort to numerical solutions and simulation-based techniques to estimate their model.

Given this extended model, it is easy to see that data on job and unemployment spell durations now only convey information on λ_0, δ, and $\tau \times \lambda_1$: the unconditional probability of a job-to-job transition is now $\tau\lambda_1/2$, and the rest of the analysis of subsection 3.1 stays unaffected. We now argue, however, that separate identification of λ_1 and τ is theoretically possible from a combination of the model's structure and the use of some higher-order moments of the wage data.

In our modified setting the one-period dynamics of the negotiation baseline are the following:

$$v_{t+1} \mid v_t = \begin{cases} v_t & \text{with probability } (1-\delta)\left[1-\tau F(v_t)\right] - \tau\lambda_1 \overline{F}(v_t)^2, \\ v' < v_t & \text{with density } \tau(1-\delta) f(v'), \\ v' > v_t & \text{with density } 2\tau\lambda_1 f(v') \overline{F}(v'), \end{cases} \tag{19}$$

The appendix then extends the above characterization of one-period wage dynamics to derive expressions for the $3T-1$ moments listed in (15) and matched in the estimation procedure described in Section 4.

In addition to the set of means, variances and covariances described in (15), we now consider the third moment of the distribution of wage changes in the hope of disentangling the arrival rate of productivity shocks τ from the job mobility transition rate $\tau\lambda_1/2$. Indeed, looking at the wage process (19), it appears that τ alone drives wage cuts while $\tau\lambda_1$ drives wage increases. Thus the relative size of these arrival rates will impact the skewness of the distribution of wage changes.[48] Specifically, we use the following additional moments in the estimation:[49]

$$E\left[(\ln\phi_{t+s} - \ln\phi_t)^3\right] = E\left[(v_{t+s} - v_t)^3\right] + 6\sigma_{\text{me}}^2 \cdot E\left[v_{t+s} - v_t\right]. \tag{20}$$

Details of the analytical derivation of that moment are reported in the appendix.[50]

We then proceed to match the basic set of moments listed in (15) (duly amended to account for persistence in productivity shocks — see the appendix), supplemented by the $T-1$ moments in (20)

[48] The discussion in subsection 5.2 gives us additional a priori reason to be hopeful that the (asymmetry of the) distribution of wage changes has some practical informative content, as we then observed that our simple model with i.i.d. shocks has trouble replicating the asymmetry of that distribution. Specifically, it tends to over-predict wage cuts, which may be caused by an excessively high frequency of productivity shocks.

[49] Once again, this moment is meant conditional on continuous employment between times t and $t+s$, although we keep the conditioning implicit.

[50] Note that we assume that the distribution of the measurement error is symmetric, i.e. has zero skewness. Also, looking at wage changes means that the worker fixed effect p does not appear in this expression and that we do not require any assumptions regarding the skewness of the distribution $H(p)$.

computed for $s = 0, 12, 24, \cdots, 12(T-1)$ using EWMD. Results of this estimation are reported in the third column of Table 5. The first column in that table merely repeats parameter estimates obtained from the basic model and reported earlier in Table 3 for comparison, while the second column reports estimates obtained by matching the full set of moments, including the third-order moment (20), yet constraining productivity shocks to be i.i.d. by fixing τ at 1.

< Table 5 about here. >

The main conclusion that we draw from those estimates is that the arrival rate of productivity shocks cannot be precisely estimated with our data. Looking at the third column of Table 5, the point estimate of τ is 0.282 with a (bootstrap) standard error of 0.386. While that point estimate may appear markedly lower than unity, a Wald test of the hypothesis $\tau = 1$ produces a p-value of 0.045, indicating borderline rejection (or acceptance) of the restricted model against the unrestricted one at the five percent level.[51] It should further be noted that estimates of all other parameters remain very stable across the columns of Table 5. All these findings concur to suggest that the particular value assumed by τ makes little difference to our model's ability to fit the set of moments described above.[52] In other words, given the model's structure, there does not seem to be enough information in our data to practically identify the degree of persistence of productivity shocks. We conjecture that this negative result would also obtain given a more sophisticated description of the productivity process or without the *ad hoc* simplifying assumption about the forced joint occurrence of job offers and productivity shocks made here for convenience, and that firm-level data conveying direct information on productivity are needed to credibly identify parameters of the productivity process.[53]

[51] Wald tests of overidentifying restrictions reported at the bottom of Table 5 otherwise lead to accepting either models against an arbitrary alternative.

[52] We also experimented with larger sets of moments, going up to the sixth moment of the distribution of wage changes. All specifications yielded equally imprecise estimates of τ.

[53] Of course an alternative route would be to gain identification power by further tightening the model's structure and assuming more specific and less flexible functional forms for some elements of the model. This is the route followed by Yamaguchi (2006), to which we refer the reader.

7 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 12-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 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. We pursue those ideas in more recent work (Postel-Vinay and Turon, 2008), where we include firing costs and minimum wage regulations to the present setting. Our mutual consent assumption then gives rise to endogenous agreements on firm-worker separations with severance packages. We use this setup to analyze (inter alia) the impact of firing restrictions on wage inequality. Finally, closing the model would allow us to relate our framework 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.[54]

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. Although this will make the derivations of the predicted moments of interest much more cumbersome,[55] it is likely to improve the fit of the model

[54]Even 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 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).

[55]See also Yamaguchi (2006) and Bagger, Fontaine, Postel-Vinay and Robin (2006) whose focus is on the decomposition of young workers' wage growth between human capital accumulation and job search.

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 can be enriched by a more careful depiction 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, Section 6 has demonstrated the need for direct evidence on productivity shocks to ensure practical identification of the model. Such evidence 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 currently ongoing research, Lise, Meghir and Robin (2008) extend the model we have offered in this paper to allow for more sophisticated, firm-specific productivity processes and plan to estimate their model on matched employer-employee data.

References

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

[2] Abowd, J. M., F. Kramarz and D. N. Margolis (1999) "High Wage Workers and High Wage Firms", *Econometrica*, 67, 251-333.

[3] Altonji, J. and L. Segal (1996), "Small-sample Bias in GMM Estimation of Covariance Structures", *Journal of Business and Economic Statistics*, 14(3), 353-366.

[4] Alvarez, J., M. Browning, and M. Ejrnæs (2001), "Modelling Income Processes with Lots of Heterogeneity", CAM discussion paper 2002-01.

[5] Bagger, J., F. Fontaine, F. Postel-Vinay and J.-M. Robin (2006), "A Feasible Equilibrium Search Model of Individual Wage Dynamics with Experience Accumulation", University of Bristol mimeo.

[6] Baker, M. (1997), "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings", *Journal of Labor Economics*, 15 (2), 338-75.

[7] Barlevy, G. (2008), "Identification of Search Models using Record Statistics", *Review of Economic Studies*, 75(1), 29-64.

[8] 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.

[9] Binmore, K., A. Shaked and J. Sutton (1989), "An Outside Option Experiment", *Quarterly Journal of Economics*, 104(4), 753-70.

[10] Blundell, R. and I. Preston (1998), "Consumption Inequality and Income Uncertainty", *Quarterly Journal of Economics*, 113, 603-40.

[11] Burdett, K. and M. G. Coles (2003), "Wage/Tenure Contracts and Equilibrium Search", *Econometrica*, 71(5), 1377-404.

[12] Burdett, K. and D. T. Mortensen (1998), "Wage Differentials, Employer Size and Unemployment", *International Economic Review*, 39, 257-73.

[13] Cahuc, P., F. Postel-Vinay and J.-M. Robin (2006), "Wage Bargaining with On-the-job Search: Theory and Evidence", *Econometrica*, 74(2), 323-64.

[14] Dey, M. S. and C. J. Flinn (2005), "An Equilibrium Model of Health Insurance Provision and Wage Determination", *Econometrica*, 73(2), 571-627.

[15] Diamond, P. A. (1982), "Wage Determination and Efficiency in Search Equilibrium", *Review of Economic Studies*, 49, 761-82.

[16] Eckstein, Z. and G. J. Van den Berg (2005), "Empirical Labor Search: A Survey", forthcoming *Journal of Econometrics*.

[17] Fella, G. (2004), "Optimal Severance Pay in a Matching Model", Queen Mary University of London mimeo.

[18] Flinn, C. (1986), "Wages and Job Mobility of Young Workers", *Journal of Political Economy*, 94(3), S88-S110.

[19] Guiso, L., L. Pistaferri and F. Schivardi (2005), "Insurance Within the Firm", *Journal of Political Economy*, 113(5), 1054-87.

[20] Hall, R. E. (2005), "Employment Fluctuations with Equilbrium Wage Stickiness", *American Economic Review*, 95(1), 50-65.

[21] Harris, M. and B. Holmström (1982), "A Theory of Wage Dynamics", *Review of Economic Studies*, 49, 315-33.

[22] Kotlarski, I. (1967), "On Characterizing the Gamma and Normal Distributions", *Pacific Journal of Mathematics*, 20, 69-76.

[23] Lillard, L. A. and Y. Weiss (1979), "Components of Variation in Panel Earnings Data: American Scientists 1960-1970", *Econometrica*, 47(2), 437-54.

[24] Lillard, L. A. and R. J. Willis (1978), "Dynamic Aspects of Earnings Mobility", *Econometrica*, 46(5), 437-54.

[25] Lise, J., C. Meghir and J.-M. Robin (2008, "Matching, Sorting and Wages", University College London *mimeo*.

[26] MacLeod, W. B. and J. M. Malcomson (1993), "Investments, Holdup, and the Form of Market Contracts", *American Economic Review*, 83(4), 811-37.

[27] 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.

[28] Malcomson, J. M. (1997), "Contracts, Hold-Up, and Labor Markets", *Journal of Economic Literature*, XXXV, 1916-57.

[29] 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.

[30] Meghir, C. and L. Pistaferri (2004), "Income Variance Dynamics and Heterogeneity", *Econometrica*, 72(1), 1-32.

[31] Mortensen, D. (2003), *Wage Dispersion: Why Are Similar Workers Paid Differently?* Cambridge, MA: MIT Press.

[32] 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.

[33] 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.

[34] Pissarides, C. A. (2000), *Equilibrium Unemployment Theory*, 2nd edition. Cambridge, MA: MIT Press.

[35] Postel-Vinay, F. and J.-M. Robin (2002a), "The Distribution of Earnings in an Equilibrium Search Model with State-Dependent Offers and Counter-Offers", *International Economic Review*, 43(4), 1-26.

[36] Postel-Vinay, F. and J.-M. Robin (2002b), "Wage Dispersion with Worker and Employer Heterogeneity", *Econometrica*, 70(6), 2295-350.

[37] Postel-Vinay, F. and J.-M. Robin (2006), "Microeconometric 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).

[38] Postel-Vinay, F. and H. Turon (2008), "Severance Packages", University of Bristol *mimeo.*

[39] Shimer, R. (2006), "On-the-Job Search and Strategic Bargaining", *European Economic Review*, 50(4), 811-30.

[40] Stewart, M. (2004), "The Inter-related Dynamics of Unemployment and Low-Wage Employment", mimeo.

[41] Sutton, J. (1986), "Non-Cooperative Bargaining Theory: An Introduction", *Review of Economic Studies*, 53, 709-24.

[42] Thomas, J. and T. Worrall (1988), "Self-Enforcing Wage Contracts", *Review of Economic Studies*, 55, 541-54.

[43] Yamaguchi, S. (2006), "Job Search, Bargaining, and Wage Dynamics", University of Wisconsin mimeo.

APPENDIX: Derivation of Theoretical Moments

All derivations in this appendix are made under the general specification of Section 6, which allows for non-i.i.d. productivity shocks. This encompasses the main baseline case of Sections 2-5, which is obtained by setting $\tau = 1$ in what follows.

Derivation of $G_s(\cdot)$. We construct $G_s(\cdot)$ by induction. Considering the conditional distribution of $v_{t+1} \mid v_t$ in (19), one has for any $s \geq 1$:

$$g_{s+1}(v) = \left[1 - \tau F(v) - \frac{\tau\lambda_1}{1-\delta}\overline{F}(v)^2\right] \cdot g_s(v) - \left[\tau f(v) - \frac{2\tau\lambda_1}{1-\delta} f(v)\overline{F}(v)\right] \cdot G_s(v) + \tau f(v), \quad (21)$$

which integrates as:

$$G_{s+1}(v) = \left[1 - \tau F(v) - \frac{\tau\lambda_1}{1-\delta}\overline{F}(v)^2\right] \cdot G_s(v) + \tau F(v). \quad (22)$$

Given the initial condition $G_0(\cdot) \equiv G(\cdot)$, this difference equation solves as (13) in the main text. With the notation $b(v) = 1 - \tau F(v) - \frac{\tau\lambda_1}{1-\delta}\overline{F}(v)^2$, we can rewrite this equation as:

$$G_s(v) = (1 - b(v)^s) \cdot G_\infty(v) + b(v)^s \cdot G(v), \quad (23)$$

It is straightforward to check that $\forall v$, $\overline{G}_\infty(v) \geq \overline{G}(v)$, i.e. $G_\infty(\cdot)$ first-order stochastically dominates $G(\cdot)$. Integration by parts further shows that:

$$E_{G_s}(v) = v_0 + \int_{v_0}^{v_{\max}} \overline{G}_\infty(x)\,dx - \int_{v_0}^{v_{\max}} b(v)^s \cdot \left[\overline{G}_\infty(x) - \overline{G}(x)\right] dx, \quad (24)$$

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

Derivation of Cov $(\ln\phi_t, \ln\phi_{t+s} \mid$ employment at $t, \cdots, t+s)$. We begin by the derivation of the conditional expectation $E(v_{t+s} \mid v_t, \text{employment at } t, \cdots, t+s)$. 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 (12) that that for any differentiable function $\varphi(\cdot)$:

$$E\left[\varphi(v_{t+1}) \mid v_t,\ \text{employment at } t, t+1\right] = (1-\tau)\varphi(v_{\max}) + \tau E_F\left[\varphi(v)\right] - \int_{v_t}^{v_{\max}} \varphi'(x) \cdot b(x)\,dx. \quad (25)$$

Next defining $T_s(v_t) = E(v_{t+s} \mid v_t,\ \text{employment at } t, \cdots, t+s)$, 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 $s \geq 2$:

$$\begin{aligned} T_{s+1}(v_t) &\equiv E(v_{t+s+1} \mid v_t,\ \text{employment at } t, \cdots, t+s+1) \\ &= E\left[E(v_{t+s+1} \mid v_{t+1},\ \text{employment at } t+1, \cdots, t+s+1) \mid v_t,\ \text{employment at } t, t+1\right] \\ &= E\left[T_s(v_{t+1}) \mid v_t,\ \text{employment at } t, t+1\right]. \end{aligned} \quad (26)$$

Reasoning by induction and differentiating (25) shows that $T_s'(v_t) = b(v)^s$. Substituting back into (25) and (26) establishes the following for $T_s(v_t)$:

$$T_s(v_t) = (1-\tau) T_{s-1}(v_{\max}) + \tau E_F[T_{s-1}(v)] - \int_{v_t}^{v_{\max}} b(x)^s \, dx. \tag{27}$$

Note that $E_F[T_{s-1}(v)]$ is a deterministic term depending on the "prediction horizon" s only. Then from (27) we directly obtain:[56]

$$\begin{aligned} \text{Cov}(\ln \phi_t, \ln \phi_{t+s}) &= \text{Var}(\ln p) + \text{Cov}(v_t, v_{t+s}) \\ &= \text{Var}(\ln p) + \text{Cov}(v_t, E(v_{t+s} \mid v_t)) \\ &= \text{Var}(\ln p) - \text{Cov}\left(v_t, \int_{v_t}^{v_{\max}} b(x)^s \, dx\right). \end{aligned} \tag{28}$$

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.

Derivation of $E\left[(\ln \phi_{t+s} - \ln \phi_t)^3\right]$**.** We can rewrite (20) as:

$$\begin{aligned} E\left[(\ln \phi_{t+s} - \ln \phi_t)^3\right] &= E_G\left[E\left[(v_{t+s} - v_t)^3 + 6\sigma_{\text{me}}^2 (v_{t+s} - v_t) \mid v_t, \text{employment at } t, \cdots, t+s\right]\right] \\ &= E_G\left[R_{s,3}(v_t) - 3v_t \cdot R_{s,2}(v_t) + 3v_t^2 \cdot R_{s,1}(v_t) - v_t^3\right] + 6\sigma_{\text{me}}^2 E_G\left[R_{s,1}(v_t) - v_t\right] \end{aligned} \tag{29}$$

where $R_{s,n}(v_t) = E\left(v_{t+s}^n \mid v_t\right)$ for $n = 1, 2, 3$. In the first equality, we integrate on the distribution G as conditioning on continuous employment between the dates t and $t+s$ is independent of v_t–employment is continuous provided the match has not been hit by a destruction shock δ, which occurs independently of the individual wage.

As in equations (26) to (27), we obtain the following recurrent definition for $R_{s,n}$:

$$R_{s,n}(v_t) = (1-\tau) R_{s-1,n}(v_{\max}) + \tau E_F(R_{s-1,n}(v)) - \int_{v_t}^{v_{\max}} R_{s-1,n}'(x) \cdot b(x) \, dx \tag{30}$$

and $R_{s,n}'(x) = n x^{n-1} b(x)^s$.

Predictions of the model with respect to fourth and fifth moments of the distribution of wage changes are derived in a similar manner when computing the fit in terms of 62 moments. Details of the derivation are available upon request.

[56] We 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

$\delta\,(\times 100)$	$1/\delta$	$\lambda_1\,(\times 100)$	$\frac{1}{\lambda_1}$	$\lambda_0\,(\times 100)$	$1/\lambda_0$
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	$\mathrm{Var}(\ln p)$	σ^2_{me}	Test Statistic (df: p-value)	
1.801 (0.078)	2.219 (0.122)	0.200 (0.099)	0.063 (0.008)	0.016 (0.015)	22.23 (30: 0.784)	
						Vs. $\sigma^2_{\mathrm{me}} \neq 0$:
1.806 (0.074)	2.188 (0.176)	0.279 (0.068)	0.062 (0.007)	0.000 (const.)	24.35 (31: 0.725)	5.18 (1: 0.023)

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 $v_{\min}$. Bootstrap standard errors obtained with 500 joint estimations of the transition parameters and the ditribution parameters.

Table 3: Distributional parameters

	Cov $(\Delta \ln \phi_{i,t}, \Delta \ln \phi_{i,t+s})$, $s = \ldots$				
	0	1	2	3	4
BHPS	0.032 (0.0007)	−0.010 (−0.001)	0.000 (0.001)	−0.000 (−0.001)	0.000 (0.001)
Simulated	0.040 (0.001)	−0.012 (−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)	0.071 (0.054)	55.21 (63: 0.747)	
					Vs. $\theta_1 \neq 0$:
	0.012 (0.001)	0.009 (0.001)	0.000 (const.)	56.44 (64: 0.738)	1.23 (1: 0.267)
Simulated	0.012 (0.001)	0.014 (0.001)	0.044 (0.042)	70.79 (63: 0.234)	
					Vs. $\theta_1 \neq 0$:
	0.012 (0.001)	0.013 (0.001)	0.000 (const.)	71.76 (64: 0.236)	0.97 (1: 0.325)

Table 4: Fitting an ARMA process

Moments matched	35	46	46
Constraint	$\tau = 1$	$\tau = 1$	-
$v_{\min}$	1.801 (0.078)	1.806 (0.083)	1.848 (0.075)
μ	2.219 (0.122)	2.217 (0.121)	2.166 (0.151)
σ	0.200 (0.099)	0.197 (0.101)	0.265 (0.076)
Var $(\ln p)$	0.063 (0.008)	0.064 (0.009)	0.054 (0.016)
σ_{me}^2	0.016 (0.015)	0.017 (0.014)	0.019 (0.012)
τ	1	1	0.282 (0.386)
Test of OI restictions:			
Test Statistic	22.23	79.98	73.93
(df: p-value)	(30 : 0.784)	(41 : 0145)	(40 : 0.262)

Note: Bootstrap standard errors obtained from 500 replications of the entire two-stage estimation procedure on 500 resamples with replacement.

Table 5: Estimating the frequency of productivity shocks

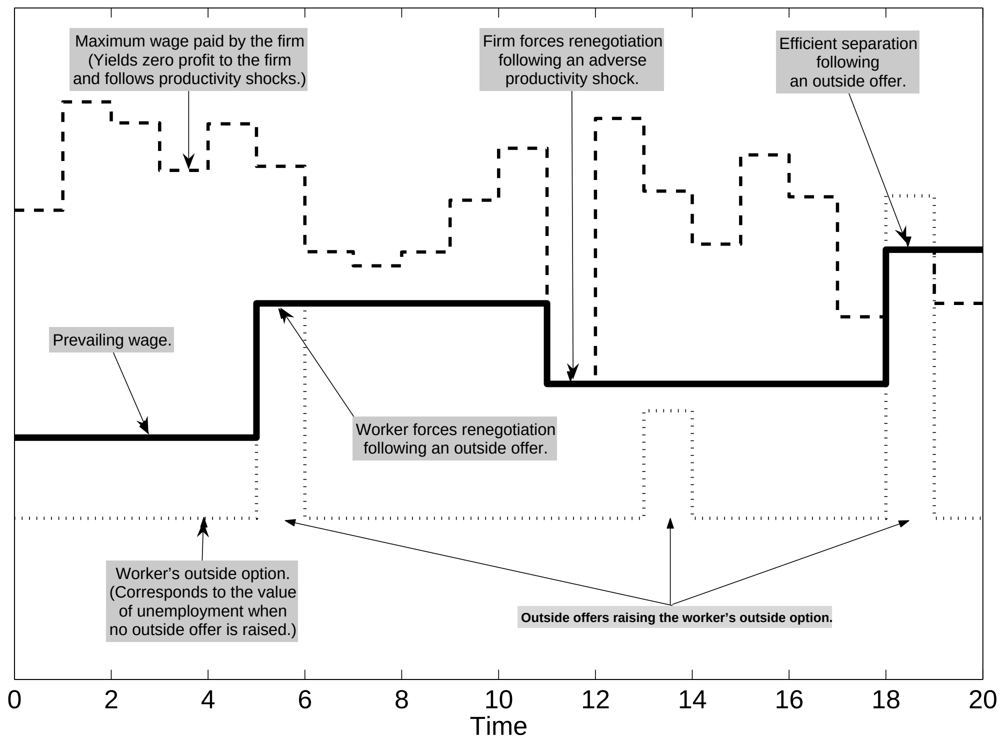

Figure 1: The wage process

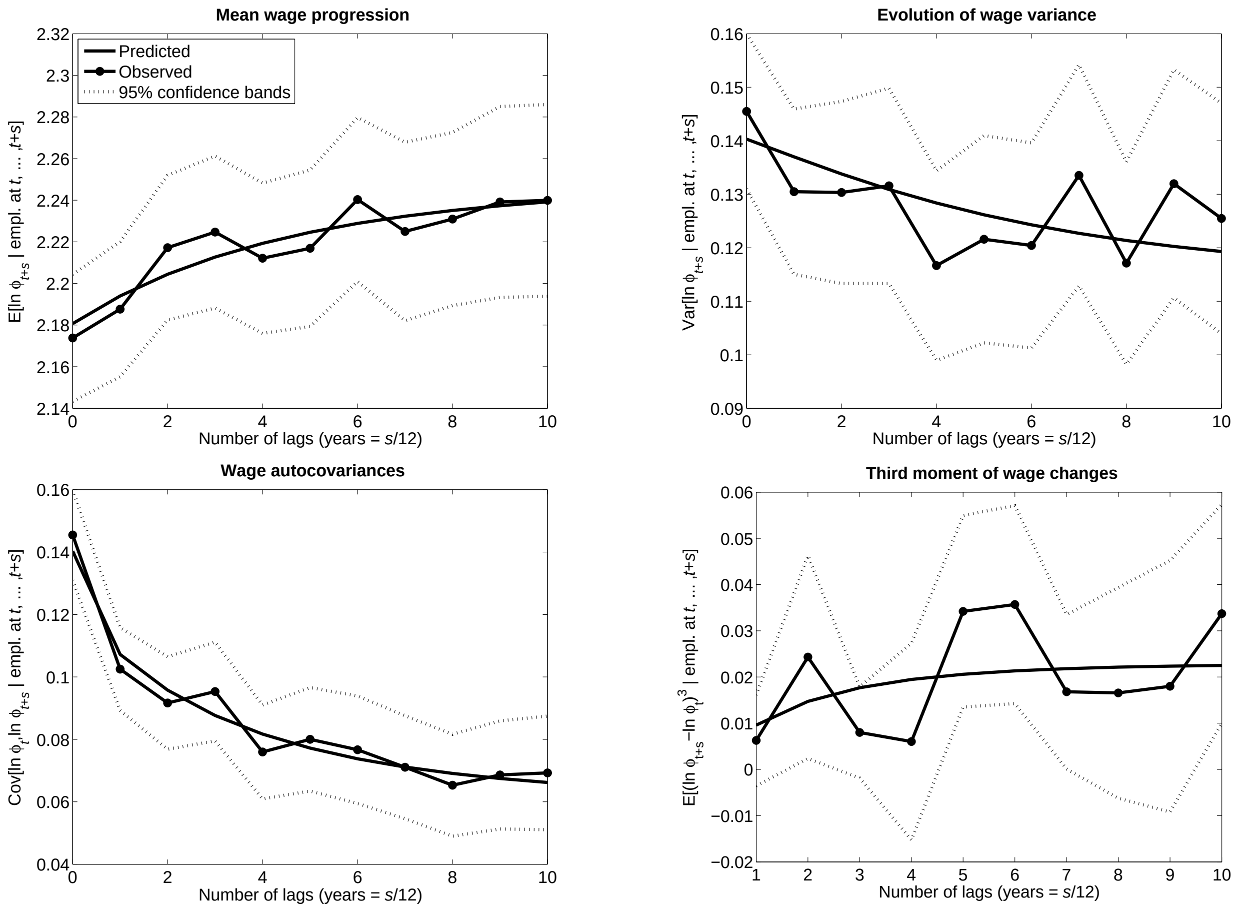

Figure 2: First-, second- and third-order moments of the wage process

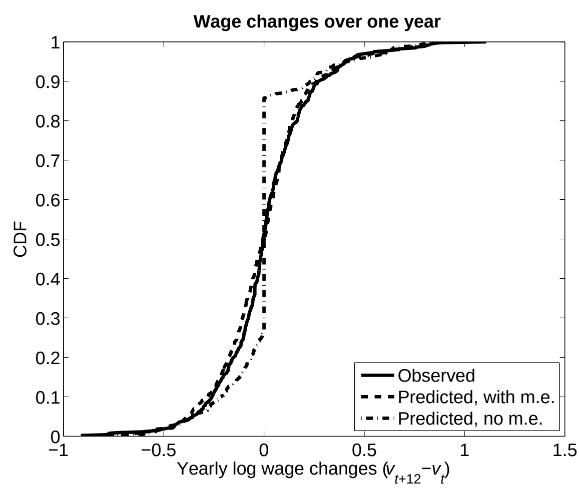

Figure 3: Distribution of year-to-year wage changes

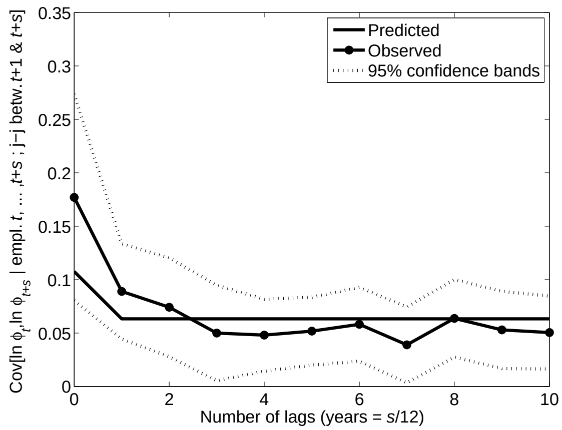

Figure 4: Wage autocovariances conditional on job-to-job mobility

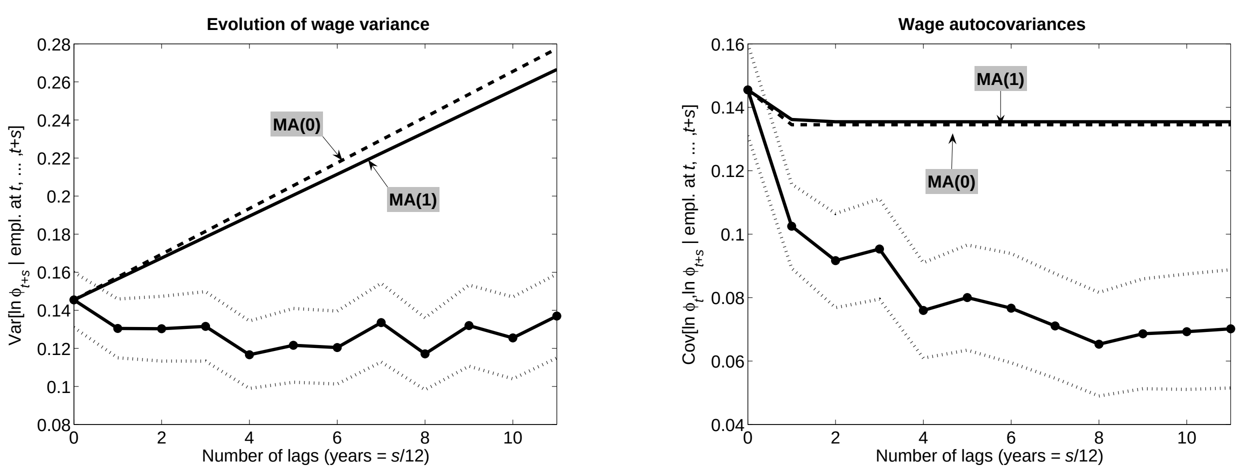

Figure 5: The wage process in levels and the ARMA decomposition