THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Bias in returns to tenure when firm wages and employment comove

Citation for published version:

Snell, A, Martins, P, Stüber, H & Thomas, J 2018, 'Bias in returns to tenure when firm wages and employment comove: A quantitative assessment and solution', *Journal of Labor Economics*, vol. 36, no. 1, pp. 47-74. <https://doi.org/10.1086/693867>

Digital Object Identifier (DOI):

[10.1086/693867](https://doi.org/10.1086/693867)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Journal of Labor Economics

Publisher Rights Statement:

Accepted for publication to Journal of Labor Economics on 1/12/2016
<http://www.journals.uchicago.edu/doi/abs/10.1086/693867>

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Bias in Returns to Tenure When Firm Wages and Employment Comove: A Quantitative Assessment and Solution

Andy Snell, *University of Edinburgh*

Pedro Martins, *Queen Mary University of London*

Heiko Stüber, *Institute for Employment Research (IAB) and
Friedrich-Alexander Universität Erlangen-Nürnberg (FAU)*

Jonathan P. Thomas, *University of Edinburgh*

It is well known that unless worker-firm match quality is controlled for, reduced-form estimates of returns to firm tenure will be biased. In this paper, we show that there is a further pervasive source of bias, namely, the comovement of firm employment and firm wages. We argue that firm-year fixed effects must be used to eliminate this bias. Estimates from two large-panel data sets from Germany and Portugal show that the bias is empirically important. Finally, we show that the results extend to tenure correlates used in macroeconomics, such as the minimum unemployment rate since joining the firm.

I. Introduction and Overview

There is a large empirical literature that attempts to identify and consistently estimate returns to firm tenure (RTT). The aim of this literature is

This work was partly supported by the Economic and Social Research Council (grant ES/L009633/1) and priority program 1764 of the German Research Foundation (DFG). We are grateful to Gary Solon for helpful discussions. Contact the cor-

[*Journal of Labor Economics*, 2018, vol. 36, no. 1]

© 2017 by The University of Chicago. All rights reserved. 0734-306X/2018/3601-0002\$10.00

Submitted March 14, 2016; Accepted December 1, 2016; Electronically published November 8, 2017

to obtain the pure causal effect of tenure on wages (Altonji and Shakotko 1987)—that is, the effect on the wage of one more year of tenure, holding constant years of experience and job-match quality broadly interpreted. In turn, this causal effect is implicitly or explicitly viewed as being a measure of the returns to firm-specific human capital and/or to contractual mechanisms that reward tenure for incentive reasons. The traditional approach is to use coefficient estimates of wages on deterministic tenure in a Mincer regression to obtain a measure of RTT. This reduced-form method is easy to implement and avoids having to make structural economic assumptions about worker entry and exit from the firm.

However, the existence of unobservable wage shocks that drive firm hiring and worker exit may complicate the interpretation of reduced-form estimates; their existence will make tenure endogenous. Put another way, in the presence of such shocks, reduced-form estimates cannot be interpreted as the causal effect of tenure on wages. Much of the literature has focused on worker-firm match quality as the key unobservable confounding factor for RTT. In particular, if we believe that better matches tend to last longer, tenure will be endogenous and failing to control for match quality will induce upward bias in reduced-form RTT estimates. Three canonical methods have been used to circumvent this problem: (i) the two-step estimator of Topel (1991), (ii) the instrumental variable (IV) approach of Altonji and Shakotko (1987), and (iii) the method of controlling for completed tenure of Abraham and Farber (1987). More recently, the emergence of very-large-panel data sets that record complete work histories of workers has allowed investigators to absorb unobserved match quality by adding firm-worker match fixed effects (see, e.g., Battisti 2012). The downside of doing this is that—as in Topel’s (1991) method—the estimated tenure effect will include the effect of linear experience, and this must be backed out using an auxiliary regression. The upside, however, is that it automatically controls for the impact of time-invariant worker and firm heterogeneity; employing fixed effects for this purpose avoids the concern that RTT estimates may be sensitive to the investigator’s selection of controls. A specification that controls for match quality using worker-firm interaction (match) fixed effects provides us with our fourth traditional method for eliminating the upward bias in RTT due to unobserved match quality.

In this paper, we identify a further and potentially equally pervasive source of bias to RTT: the existence of a time-varying wage component that is common to all of a firm’s workers but that comoves with its employment. We argue that even in a world where match quality is irrelevant, the failure to account for these wage components will bias estimates of returns to tenure,

responding author, Andy Snell, at a.j.snell@ed.ac.uk. Information concerning access to the data used in this paper is available as supplemental material online.

most likely in a downward direction. The mechanism generating the bias is simple: suppose firms that have a relatively high wage at time t also have high employment, high hiring, and low average firm tenure at t or that firms that have a relatively low wage at t also have low employment, low hiring, and high average firm tenure at t .¹ This induces negative feedback from equal-treatment wage shocks to tenure. In this paper, we show that traditional estimators—ones designed to eliminate the effects of unobservable worker-firm match quality—are not immune to potentially sizable biases arising from this effect.

Drivers of a firm's wage/employment comovements may include both aggregate (business cycle) shocks and firm-specific shocks. In both cases, the shocks that are the root cause of the problem are assumed to impact all workers in the firm. We call these common components of wages equal-treatment shocks following the relevant macro literature (see, e.g., Gertler and Trigari 2009; Snell and Thomas 2010; Moscarini and Postel-Vinay 2013 for macro models subject to within-firm equal treatment). Because equal-treatment shocks are the same for each worker in a firm in a particular year, we propose that they be controlled for via the addition of firm-year interaction fixed effects to panel wage regressions while at the same time also controlling for the more traditional match-quality problem.²

In an empirical application, we use two large samples drawn from matched-panel data sets from Germany and Portugal to show that the four traditional methods produce RTT estimates that are substantially lower than that obtained using our proposed correction. If we take the average RTT estimate from the four traditional methods as a benchmark, then adding firm-year fixed effects to wage equations (while controlling for worker-firm match quality) increases estimated RTT in Germany by about 2.5% of wages and in Portugal by about 3.5% of wages at 10 years of tenure. This amounts to about 20% and 40%, respectively, of the bias-corrected RTT level itself. Although investigators may have been aware of this problem (see, e.g., a discussion on high wage/employment growth firms in Topel 1991), to the best of our knowledge, we are the first to quantify its importance and propose a (simple) solution.

One interesting supplementary result from our estimation method is that the fitted firm-year fixed effects appear to follow a unit root; like unobserved match quality, the equal-treatment shocks also appear to have a permanent impact on a worker's wages within a firm. Given that entry into and exit out of a firm are likely driven by permanent (rather than transient) wage

¹ There is also a steady-state cross-sectional effect: high-paying firms tend to have low labor turnover and hence longer tenure. However, this type of time-invariant cross-sectional effect is usually removed via the addition of firm fixed effects.

² Aggregate business cycle shocks can be controlled for by including time-fixed effects; in fact, we find that shocks below the aggregate level account for almost all of the bias.

shocks, this is consistent with our finding that equal-treatment shocks cause bias in RTT estimates. It suggests that if one wishes to obtain the causal effects of tenure on wages, one must control for all permanent wage shocks whether they arise from equal-treatment shocks or from match quality.

A further implication of our results is that using regressors that interact macroeconomic variables, such as unemployment, with deterministic tenure will also result in biased inference. Canonical examples of such variables are Beaudry and DiNardo's (1991) minimum unemployment rate during a worker's tenure (*minu*) and a new-hire dummy interacted with unemployment to measure the incremental cyclicalities of new-hire wages. The empirical importance of these variates found in the literature adds a further twist because their omission from Mincer equations will be yet another source of bias to RTT estimates. Another way of saying this is that if wage growth within the firm contains both the effects of human capital and implicit contracts, then to consistently estimate these separate effects requires inclusion of the relevant contract variate (e.g., *minu*) and firm-year fixed effects. We examine some of these issues in Section III below.³

The key result in this paper, that there is yet another source of pervasive bias to RTT estimates obtained via reduced-form estimation, may lead the investigator to conclude that a safer way to proceed is via a fully specified structural model of wage shocks and worker mobility (for a recent example of such a model, see Buchinsky et al. 2010). However, one key finding of our work is that it is firm-specific (heterogeneous) comovement that drives the biases we find and not macro (aggregate) effects. A structural model with heterogeneous firm hiring (and firing)—as opposed to cyclically related hiring—may be hard to specify and identify empirically. Additionally, estimates gleaned from structural models may be only as good as the veracity of their underlying assumptions. As far as reduced-form modeling goes, our paper has a clear message: to avoid substantial RTT bias, one must control for not only worker-firm match quality but also equal-treatment shocks.

The paper is laid out as follows. Section II revisits the traditional econometric model of RTT and the implications for wages. We outline the four traditional estimation methods of Abraham and Farber (1987), Altonji and Shakotko (1987), Topel (1991), and the addition of match fixed effects. We estimate RTT for these methods using Portuguese and German panel data and plot the corresponding RTT profiles together with that obtained using our proposed correction. We then offer an anatomy of the bias from a theoretical and empirical viewpoint. Importantly, here we show that the bias is driven by heterogeneous (across firms) employment/wage comovements. Section III looks at the implications of our analysis when contractual vari-

³ There may, of course, be other sources of wage growth within the firm arising from wage contracts, such as back-loading, for which observable controls are not available. We discuss these issues in Sec. III below.

ables play a role, in particular, tenure-related macro variables associated with wage contracts. Section IV offers concluding comments.

II. Estimates of RTT: A Comparison of Traditional Methods and the Corrected Method

In this section, we estimate RTT using the four traditional methods outlined above and compare the implied RTT profiles obtained with that obtained using our proposed correction for firm employment/wage comovements. To begin, we revisit the bias caused by unobserved match quality and outline the methods that were designed to deal with it. We call these four methods T (Topel), AS (Altonji and Shakotko), AF (Abraham and Farber), and MFE (match fixed effects). To outline these methods, we use a somewhat simplified archetypal model of RTT. We assume that log wages w_{ijt} for worker i in firm j at time t are given by

$$w_{ijt} = \alpha + \beta\tau_{ijt} + \gamma E_{it} + \varepsilon_{ijt}, \quad (1)$$

with

$$\varepsilon_{ijt} = \theta_{ij} + \varpi_{jt} + u_{ijt}, \quad (2)$$

where τ_{ijt} is the worker's tenure and E_{it} is her lifetime work experience. The error consists of job-match quality θ_{ij} (which also may include a worker and firm fixed effect), an idiosyncratic error u_{ijt} that is assumed to be uncorrelated with the regressors (especially tenure), and an equal-treatment wage component ϖ_{jt} —the innovation in our study. The coefficient β is the per-year RTT within the firm.⁴ The traditional problem (dealt with in T, AS, etc.) arises when the job-match quality θ_{ij} is correlated with the tenure of worker i . When the match is good (high θ_{ij}), the worker's separation hazard may fall (see in particular Bowlus 1995) and expected tenure will rise. This makes tenure endogenous and biases β upward. The aspiration of the traditional RTT estimation methods is to estimate the causal effect of tenure on wages in the presence of unobserved match quality θ_{ij} . The point of this paper is to show that the existence of equal-treatment wage elements (ϖ_{jt}), in addition to match quality, undermines this aspiration.

Topel's (1991) method is to first-difference incumbents' wages to remove the (presumed time-invariant) match quality. In the absence of ϖ_{jt} , regressing these incumbent wage changes on an intercept would, in this model at least, produce a consistent estimate of $\beta + \gamma$, $\widehat{\beta + \gamma}$ say. To separately identify β and γ , Topel (1991) proposed estimating a second-stage regression of $w_{ijt} - (\widehat{\beta + \gamma})\tau_{ijt}$ on the worker's initial experience on entry to the

⁴ In more general contexts where RTT is heterogeneous across workers and/or firms, β could be interpreted as the average RTT, or average treatment effect in the words of the experimental literature.

firm. Provided the latter is not correlated with job-match quality, an admittedly strong assumption,⁵ this produces a consistent estimate of γ . Subtracting the latter estimate from $\beta + \gamma$ gives a consistent estimate of β . Altonji and Shakotko (1987) proposed an IV method whereby tenure is instrumented by the deviation of tenure from its spell mean $\widetilde{\tau}_{ijt}$. By construction, this variable is orthogonal to (constant within spell) match quality. Again, in the absence of ϖ_{jt} , this would offer consistent estimates of β .⁶ Abraham and Farber (1987) propose adding duration—the final ex post tenure of the worker at the firm—as a regressor. If workers with better matches have longer completed tenure—as the traditional bias story goes—then controlling for completed tenure directly should eliminate the bias in β .⁷ Finally, the MFE method adds match fixed effects to the estimation process. This focuses on within-match variation in tenure. As was the case with Topel, within-match de-meaned tenure and de-meaned within-match experience are the same variable, and the latter must be dropped from the estimation. The result is that the coefficient on tenure becomes an estimate of $\beta + \gamma$. To estimate γ —and hence β —one would use Topel's second stage (above).

All of these methods ignore the existence of the equal-treatment wage components ϖ_{jt} . If these components positively comove with firm employment, then they will be negatively correlated with firm average tenure, and this will induce downward bias in estimates of β . We propose to augment the MFE estimator with firm-year interaction fixed effects (FYFE). The FYFE will absorb the equal-treatment wage components and eliminate the bias arising from wage/employment comovements. As with MFE and Topel, we use a second-stage regression of $w_{ijt} - (\beta + \gamma)\tau_{ijt}$ on the worker's experience at entry to the firm to obtain an estimate of γ . If it is true that more experienced workers do find better matches (and this effect does have significant traction), then the estimated γ will be upward biased and RTT will

⁵ Topel (1991) argued that more experienced workers are likely to form better matches, in line with job-shopping models of search. If true, returns to experience will be overestimated in the second stage and tenure underestimated—his RTT estimates are a lower bound. He considers in detail two further sources of bias: first, frequent job changers may be less productive, in which case more able workers' initial experience will tend to be lower, leading to γ being underestimated. Second, jobs offering low wage growth may survive with a lower probability than higher wage growth jobs. This could lead to an overestimate of $\beta + \gamma$. Topel (1991) gives evidence to suggest that these biases are not likely to be significant; we discuss the issue further in Sec. V.

⁶ As with Topel (1991), this requires that experience is not correlated with job-match quality. If it is positively correlated, again presumably because of job shopping, then the estimate of γ will be biased upward and that of β downward biased, although they argue that this effect is relatively small (see Altonji and Shakotko 1987, 450–3).

⁷ This is under their assumption that initial experience is correlated with match quality only through total job duration.

be downward biased. At worst, therefore, the RTT profile produced by FYFE will be a lower bound for the true RTT.⁸

A. Data

We draw our data from the German Beschäftigten-Historik (BeH) and the Portuguese Quadros de Pessoal (QP). Before discussing our subsamples, we give a brief overview of these two well-known data sources. We then describe the samples and the cleaning operations we perform on them.

The BeH data set is organized by worker spells. A spell is a portion of a year spent at a single firm. For the BeH, if a worker stays with one firm throughout the year, the average daily wage for that spell forms a single data point. If the worker moves to a second firm within the year, there will be two spells that year; the average wage at each firm would form a separate data point for that year. By contrast, the QP is an annual survey that records data on each worker at only one point in the year (census date in March up to 1993 and in October from 1994 onward). For the QP then, there is only one worker spell per year.

The BeH draws data from the totality of gainfully employed members of the German population who are covered by the social security system. Those not covered are self-employed, family workers assisting in the operation of a family business, civil servants (*Beamte*), and regular students. The BeH covers roughly 80% of the German workforce. We focus solely on workers employed in states of the former West Germany. Plausibility checks performed by the social security institutions and the existence of legal sanctions for misreporting guarantee that the earnings data are very reliable—in contrast with interview-based wage data, such as that in the Panel Study for Income Dynamics (PSID; for the United States) or the Socio-Economic Panel (for Germany).

Unfortunately, the BeH documents only total spell earnings and not hours worked in that spell. We therefore consider only full-time workers. Nearly all full-time workers in Germany work a standard number of hours per week, so the average daily wage should be very closely related to the hourly wage. To calculate the daily real wage (in 2005 prices), we use Germany's consumer price index (CPI). Another problem is that wages are censored at a maximum level equal to the contribution assessment ceiling of the compulsory pension insurance scheme.⁹ Earnings spells with wages above or close to (within 1% of) the truncation point are dropped. We drop all

⁸ In the presence of the mechanism identified in this paper, an additional likely upward bias exists in the estimation of γ (and hence an additional downward bias in RTT). See n. 16 below.

⁹ In a sensitivity analysis in Snell et al. (2016), we found that our core result—the downward bias when FYFE is not controlled for—was robust with respect to artificially censoring the highest wages in our already censored sample. This suggests that the original censorship is not impacting our results.

spells that have missing tenure. This means that a worker enters the data only when he joins a firm after January 1, 1975.¹⁰ For this reason and to match the data period used for Portugal, we drop the first 12 years and use worker spells dated 1986 and beyond.

The QP covers all workers except the self-employed and those employed in the public sector; of course, the unemployed and the inactive are also not included.¹¹ There are several wage variables, all of them expressed in monthly values (the most common type of pay in Portugal), including base wages, tenure-related payments, overtime pay, subsidies, and other payments (this latter category includes bonuses and profit- or performance-related pay). All QP wages have been deflated using Portugal's CPI and are expressed in 2005 euros. There is also information about normal hours and overtime hours per month. The benchmark measure of pay adopted in this study is based on the sum of all five types of pay divided by the sum of the two types of hours worked, resulting in a measure of total hourly pay. Tenure—in both data sets—is measured (in rounded years) as the current year minus the reported start year.

From the QP, we sample all workers from the 127 largest firms that existed throughout the entire period 1986–2009. From the BeH, we sample all full-time, full-year worker spells from the 100 largest (West German) firms that existed throughout the same time period.¹² The motivation for using large firms is to enable us to get good estimates of the ϖ_{ji} for subsequent analysis. An additional reason to focus on a relatively small number of firms is to allow a subsequent computation of diagnostic regressions (below) that involve more than $2n$ regressors with two-dimensional fixed effects. Of course, estimated RTT of workers in large firms may not be representative of RTT that exist in the economy at large. But in an earlier version of this paper (Snell et al. 2016), we showed that our main result—that there is substantial downward RTT bias if you fail to control for equal-treatment wage components—is robust with respect to changing the sample to (a) one consisting of the 1,000 largest firms and (b) a randomly drawn sample of 10,000 (mostly small) firms.

¹⁰ For this analysis we use only the years 1986–2009, but for the identification of firm entrants and the calculation of firm tenure we use BeH data from 1975 onward. However, we exclude all spells starting January 1, 1975, because the tenure could be left censored.

¹¹ See Martins et al. (2012) for further details and a recent example of work using the QP.

¹² The BeH reports establishment-level data, and the QP reports firm-level data. In the paper, we refer to both as firms. Focusing on full-year spells eliminates anomalies such as supposed full-time workers working for two firms at the same time and workers who have short-tenured jobs. It also gives a cleaner approach to estimating within-firm wage growth—especially when we use first differences (Topel) in a regression. Finally, having a maximum of one observation per worker per year makes the German sample more comparable to the Portuguese sample.

Despite the small number of firms, the sample still yields around 3.3 million data points for Portugal and 12.8 million data points for Germany, around 5% and 3% of the total available data from the QP and the BeH, respectively, over this period. We also observe a substantial proportion of workers in more than one firm—just under 5% of workers in the Portuguese data and just under 10% in the German data. These are higher proportions than one would expect if workers joined firms randomly. This suggests that the labor markets within which these large firms operate have a high degree of segmentation from the rest of the labor market.

Table 1 offers some summary statistics from the two samples. It shows some stark differences in the two samples and in the two labor markets. Aside from average wages being very much lower in our Portuguese sample (as we would expect), wages are more than twice as variable therein. Average tenure, however, is substantial in both samples. Average separation rates (which can be backed out from average tenure) are around 10% per year, considerably lower than the around 30% level in the United States (see, e.g., Hobbijn and Şahin 2007). The sixth and last rows give the average firm sizes in our core sample of 100 or 127 large firms and in the wider economy (as recorded in the BeH and QP, respectively). Firms in general are smaller in Portugal than in Germany. Our samples also indicate that variation in size may be higher in Portugal than in Germany. The stark differences in the labor markets is reassuring for our analysis; if we find similar results from both countries, then those results will have greater external validity than those based on a single data set.

B. Implementation of the Methods and Estimates

We generalize the tenure function in equation (1) by allowing RTT to follow a quartic function.¹³ Experience is modeled via a quadratic function, and we control for business cycles and common trends using year fixed effects in all methods.¹⁴ We control for time-invariant worker heterogeneity using first differences in Topel, match fixed effects in MFE and FYFE, and worker fixed effects in the other specifications.¹⁵ We now give specific implementation details method by method.

¹³ We also tried adding a tenure-zero dummy to the quartic to capture any additional wage effect of being a new hire that the quartic specification cannot easily capture; while we find that there is a significant pay increase in the first year, in line with previous work (e.g., Altonji and Shakotko 1987, table 1), the impact on RTT is small in both data sets; likewise, the bias we find is virtually unchanged.

¹⁴ In Snell et al. (2016) we found that the bias in MFE (RTT from FYFE minus that from MFE) was virtually unchanged when we generalized the experience function to a quartic.

¹⁵ As well as avoiding the use of subjectively selected controls that may otherwise drive differences in our RTT estimates, our fixed effects also imply that estimates of γ (linear experience) are driven by the same source of variation in each of the five

Table 1
Sample Summary Statistics

	Portugal	Germany
Average log monthly wage (2005 euros)	7.01	7.91
Standard deviation of log monthly wage	.637	.260
Standard deviation of log average annual firm employment	.882	.760
Average tenure (years)	12.9	9.26
Standard deviation of tenure	10.2	7.28
Number of worker spells per firm per year	1,172	5,278
Number of years available (1986–2009) ^a	22	24
Number of tenure categories available	51	36
Average firm size in QP and BeH in 1997	9.5	14.6

NOTE.—Worker-level data for 1990 and 2001 are not available from the QP. For comparison purposes, we present average establishment size in both cases. Average firm size in Portugal is marginally higher at 11.0. BeH = Beschäftigten-Historik; QP = Quadros de Pessoal.

^a Data are missing in the QP for 1991 and 2002.

Topel. For the first stage in Topel, we estimate the following regression using data only on incumbents:

$$\Delta w_{ijt} = \delta + \beta_2 \Delta \tau_{ijt}^2 + \beta_3 \Delta \tau_{ijt}^3 + \beta_4 \Delta \tau_{ijt}^4 + \gamma_2 \Delta E_{ijt}^2 + u_{ijt}.$$

The estimate of δ , $\hat{\delta}$ say, gives an estimate of $\beta_1 + \gamma_1 + \mu_1$, where β_1 and γ_1 are the linear terms of the quartics in tenure and experience, respectively, and μ_1 is the linear trend. We then regress the levels residual $w_{ijt} - \hat{\delta} \tau_{ijt} - \hat{\beta}_2 \tau_{ijt}^2 - \hat{\beta}_3 \tau_{ijt}^3 - \hat{\beta}_4 \tau_{ijt}^4 - \hat{\gamma}_2 E_{ijt}^2$ on E_{0ijt} and $(\tau_{ijt} - tr_t)$, where E_0 is initial experience on joining the firm and tr is a time index. The coefficients from this second regression ($\hat{\gamma}_1$ and $\hat{\mu}_1$) are consistent estimates under our assumptions of γ_1 and μ_1 , respectively. The estimate of β_1 is then obtained as $\hat{\beta}_1 = \hat{\delta} - \hat{\gamma}_1 - \hat{\mu}_1$.¹⁶

MFE and FYFE. For the first stage in MFE, we estimate

$$w_{ijt} = \beta_1^* \tau_{ijt} + \beta_2 \tau_{ijt}^2 + \beta_3 \tau_{ijt}^3 + \beta_4 \tau_{ijt}^4 + \gamma_2 E_{ijt}^2 + \theta_{ij} + u_{ijt} \quad (3)$$

using match fixed effects to control for match quality θ_{ij} . To estimate FYFE, we add firm-year interaction fixed effects to equation (3). Due to the addition of match fixed effects, the estimated linear tenure coefficient $\hat{\beta}_1^*$ in MFE and FYFE is an estimate of $\beta_1 + \gamma_1$. Unlike Topel, where the linear tenure coefficient also includes the effect of trend, here the deterministic trend is

methods, namely, the wage variation of those workers who switch firms in our sample.

¹⁶ Because δ , by the reasoning of the paper, is downward biased, for this to be a consistent estimate of β_1 requires an additional assumption that initial experience is uncorrelated with duration as well as match quality (if they are positively correlated, then γ_1 will be upward biased and RTT will be downward biased). This applies equally to the second stage of MFE below. But if experience is correlated with match quality only via its correlation with duration, as assumed in AF, then this is not an additional assumption.

identified separately and absorbed in the year and firm-year fixed effects, respectively.¹⁷ To obtain a consistent estimate of γ_1 , we regress the levels residual $w_{ijt} - \hat{\beta}_1^* \tau_{ijt} - \hat{\beta}_2 \tau_{ijt}^2 - \hat{\beta}_3 \tau_{ijt}^3 - \hat{\beta}_4 \tau_{ijt}^4 - \hat{\gamma}_2 E_{ijt}^2$ on initial experience. This coefficient is subtracted from $\hat{\beta}_1^*$ to give our estimate of β_1 .

AF. In AF, we simply add the within-match variate τ_{ijt}^c —completed tenure—to the main regression. For workers whose tenure is incomplete—ongoing workers in 2009—we may use either imputed values or the actual value of tenure in 2009. We experimented with two imputations: (i) we assumed constant exit hazards after 2009, and (ii) we used the sample of workers with completed tenures to compute the expected additional tenure of workers with τ years of tenure in 2009. Both imputation methods produced profiles virtually identical to that obtained from using the value of final tenure itself, and so we simply present the profile from the latter.

AS. For AS, we adjust the tenure terms by subtracting their respective within-match means. For example, τ_{ijt}^3 becomes $\tau_{ijt}^3 - \bar{\tau}_{ijt}^3$ where the overbar denotes within-match mean. These variates are used as instruments for the tenure terms in a two-stage least-squares (2SLS) IV regression.

C. Results

The estimates and standard errors of the quartic tenure parameters β_i , $i = 1, \dots, 4$ for the four traditional methods (T, AF, AS, and MFE) together with the corrected method (FYFE) are presented in table A1 of the appendix, available online. The coefficient estimates are quite hard to map into RTT itself—which is the object of interest here. More informative is the RTT tenure profiles implied by these estimates. We plot these profiles for values of tenure from 0 to 20 years in figure 1 (Portugal) and figure 2 (Germany).

The graphs show that AS, AF, and MFE offer similar RTT estimates.¹⁸ The methods themselves are in fact quite close. Both AS and MFE measure the tenure regressors in the same way—that is, as deviations from match mean. In fact, in the absence of other regressors, the 2SLS tenure estimates

¹⁷ Topel argues that the estimate of deterministic trends from levels are upward biased because of the secular tendency for worker quality to improve. For this reason, he uses an extraneous trend estimate. However, this critique does not apply to MFE and FYFE because in those specifications the match quality of every worker is controlled for via match fixed effects.

¹⁸ The 10-year tenure effects for the traditional estimators are broadly in line with what Altonji and Williams (2005) find for the United States in their reappraisal of earlier work. That tenure profiles are falling at higher tenures is not unusual in the literature. For example, in Altonji and Williams's replication exercise, when, as here, the time trend is controlled for using time dummies and a quartic in tenure is included, the IV1 estimates (AS here) have a falling tenure profile above 5 or 10 years depending on the specification, while Topel's method yields the same at 10 and 20 years in one of the two specifications reported (table 3 and n. 12 in Topel 1991). When they adjust the relative timing of wage and tenure measures, they find that the tenure effect is negative above 10 years for both AS and T (table 5 in Topel 1991).

FIG. 1.—Returns to firm tenure profiles from the five methods (Portugal). AF = Abraham and Faber; AS = Altonji and Shakotko; FYFE = firm-year interaction fixed effects; MFE = match fixed effects; T = Topel. A color version of this figure is available online.

of AS would be identical to those of MFE. As far as AF goes, if completed tenure is a good proxy for match quality, then AF will also effect an approximate within-spell de-meaning of the regressors.

By contrast with the other three methods, Topel's method produces RTT estimates that are quite low, and the dynamic pattern is also somewhat different. Of course, Topel uses a first-differences specification in contrast to the levels of MFE, AF, and AS—an important difference that sets Topel apart from the other methods. At the same time, we should also point out that the standard errors of Topel's estimates are quite high for Germany, suggesting that the corresponding RTT schedule is not as precisely estimated as the others. For Portugal, Topel's estimates are better determined, but once again the corresponding RTT lie substantially below that of the other three methods, and this is quite hard to rationalize.¹⁹

¹⁹ Reversing the original findings in the literature, Altonji and Williams (2005) argue that the reason for Topel's (1991) finding of a much higher RTT than previ-

FIG. 2.—Returns to firm tenure profiles from the five methods (Germany). AF = Abraham and Faber; AS = Altonji and Shakotko; FYFE = firm-year interaction fixed effects; MFE = match fixed effects; T = Topel. A color version of this figure is available online.

The key point, however, is that in both data sets the corrected RTT profile lies substantially above that of the other four methods.²⁰ The purest measure of the impact of adding firm-year fixed effects on RTT can be seen by looking at the vertical gap between FYFE and MFE, because the two methods differ only by the application of our proposed correction. This shows a

ously estimated is to a substantial extent accounted for by his use of a secular wage trend from an alternative data source (using a Current Population Survey–based wage index rather than from the PSID he uses in his estimations, as was the case in Abraham and Farber [1987], Altonji and Shakotko [1987], and this study) and his use of lagged wages with current tenure. They also argue, in contrast to Topel, that individual heterogeneity biases the return to experience downward in T (but not AS) and so RTT upward, as discussed in n. 6.

²⁰ Confidence bands are not displayed to avoid cluttering the graph; however, a 95% confidence interval around Portugal's FYFE curve excludes all of the other curves when tenure is below 18 years. The FYFE profile for Germany is less well defined, and its confidence interval is wider; nonetheless, it still excludes all of the other curves when tenure exceeds 12 years.

substantial bias in the case of the Portuguese data, with FYFE lying 2.4% of wages above MFE at 10 years of tenure—around 30% of the RTT level itself, although the gap falls somewhat as tenure grows toward 20 years. For Germany, it is the other way around. MFE's RTT lies only 1.3% of wages below FYFE at 10 years of tenure, but the gap grows to around 2.5% of wages as tenure increases to 20 years. If we repeat these calculations using the average of the four traditional methods as a baseline, then the bias is of course considerably larger.

We have shown that our corrected RTT profile (FYFE) lies substantially above that of the other four methods. We now try to expose and understand better the source of these differences.

D. An Analysis of the Source of the Bias

We argued above that positive comovement of firm employment and firm wages is a new (i.e., uninvestigated) source of bias in RTT estimates; when a firm's employment and wages rise (fall) together, its average tenure falls (rises), and tenure becomes endogenous. To get a better analytical handle on how this mechanism works, we use a simple model that offers a sketch of the mechanism at work. The model has only a single regressor—tenure—with a regression error consisting only of equal-treatment effects. Explicitly, we consider the panel data regression of (log) wages on individual tenure

$$w_{ijt} = \alpha + \beta \tau_{ijt} + \varpi_{jt}, \quad (4)$$

where symbols are as previously defined. We ignore match fixed effects and the usual idiosyncratic regression error here because we wish to focus on the object of interest—bias caused by the existence of ϖ_{jt} and its comovement with firm employment at time t , L_{jt} say. We assume that the data comes from all workers in n large firms that exist over T years with total number of observations $N = (\sum_{t=1}^T \sum_{j=1}^n L_{jt})$. We discuss the interpretation of ϖ_{jt} below, but for now and for illustrative purposes we take ϖ_{jt} to be (proportional to) a mean-zero shock to firm profits.

Standard textbook theory tells us that ordinary least squares (OLS) bias in the estimate of β will arise if tenure has a nonzero covariance with the error. The sample covariance of tenure with the regression error in equation (4) is

$$\text{scov} = \frac{1}{N} \sum_{t=1}^T \sum_{j=1}^n \varpi_{jt} \sum_{i=1}^{L_{jt}} (\tau_{ijt} - \bar{\tau}), \quad (5)$$

where $\bar{\tau}$ is the sample mean tenure.²¹ We can rewrite the term in braces to get

²¹ In this illustration, we abuse notation by indexing each firm's workers in the same way; i.e., worker $i = 1, \dots, L_{jt}$, when in fact the identity of workers at time t in firm j will vary from year to year.

$$\text{scov} = \frac{1}{N} \sum_{t=1}^T \sum_{j=1}^n \varpi_{jt} \sum_{i=1}^{L_{jt}} \{(\tau_{ijt} - \bar{\tau}_{jt}) + (\bar{\tau}_{jt} - \bar{\tau})\}, \quad (6)$$

where $\bar{\tau}_{jt}$ is the average tenure of firm j and $\bar{\tau}$ is the long-run average tenure for all firms. The first braced summation term is by definition 0, so we can simplify to get

$$\text{scov} = \frac{1}{N} \sum_{t=1}^T \sum_{j=1}^n \varpi_{jt} \sum_{i=1}^{L_{jt}} (\bar{\tau}_{jt} - \bar{\tau}) \quad (7)$$

$$= \frac{1}{N} \sum_{t=1}^T \sum_{j=1}^n \varpi_{jt} L_{jt} (\bar{\tau}_{jt} - \bar{\tau}). \quad (8)$$

Now suppose in year t that there is positive comovement between the hiring done by firm j and its profit shock ϖ_{jt} . Effectively, this means that firms currently experiencing above-average profits (i.e., $\varpi_{jt} > 0$) will have above-average employment, above-average hiring, and below-average tenure. Hence, $(\bar{\tau}_{jt} - \bar{\tau}_t)$ will be negative and $\varpi_{jt} L_{jt}$ positive for such firms (vice versa for firms experiencing a negative profits shock). The net effect is to make scov—the OLS bias—negative.

Note that the above logic would apply to a random sample (rather than a complete sample) of workers from these firms: a randomly chosen worker who has a higher than average wage is more likely to have come from a firm that has high levels of current employment (and high levels of hiring) than from one with low levels, and that firm is more likely to have average tenure below the average for the economy as a whole. The bias argument goes through unchanged. One aspect that does change when we have only a random sample of workers is our ability to reliably estimate and control for equal-treatment effects ϖ_{jt} . In fact, in random worker samples we are likely to see very few workers working at the same firm, making identification of and controlling for ϖ_{jt} practically impossible. This is one reason we chose a sample of firms rather than a sample of workers.

The preceding arguments were an analytical sketch of the main mechanism. In section A2 of the appendix, we develop a more formal model of the bias. Our benchmark model is equation (1) with equation (2). The key additional assumptions are that there are complete data on all workers in a small number of long-lived large firms (offering a large number of data points in each year in each firm), that there is an exogenous worker exit/quit rate that we allow to be different in each firm, and that a worker's initial experience on entering the firm is exogenous. The model also admits a completely general set of fixed effects. We find that the RTT bias is a weighted average (across firms) of the comovements between a firm's wage and its current and lagged employment levels. A key special case occurs when comovements between current wages and current employment are positive while

those between current wages and lagged employment are 0. In this case, the biases from each of the four methods are negative. When we generalize to allow current wage to comove with lagged employment as well, this turns out to be relatively unimportant in determining the bias. It is the contemporaneous wage/employment comovement that matters most.

E. The Economic Mechanism Behind the Bias

Our primary claim is that the existence of equal-treatment wage components that drive employment is an important source of RTT bias. We argue in Section II.F that it is movements in these components below the macro level that matter. We now try and justify that claim. First, we discuss some models that are consistent with our approach. Then, we look at the nature of the shocks we have identified and argue that they are consistent with our contention.

Consider a standard search-matching framework adapted to large firms (Elsby and Michaels 2013), with continuous bargaining.²² Positive comovement of wages with employment requires, for example, that the higher wages after a positive firm shock be associated with more matches being made or fewer separations; for Elsby and Michaels (2013), a positive or negative shock to a firm's productivity of sufficient size will lead to the firm increasing or decreasing its vacancies and hence hiring or laying off workers. A similar story could be told in a rent-sharing or union model where positive shocks to a firm's price or productivity lead to higher employment, profits, and wages of all workers.

A number of wage-posting models with on-the-job search exhibit positive tenure effects even in the absence of specific capital accumulation. For example, Burdett and Coles (2003) show that with risk-averse workers, wage-tenure contracts can arise, in which wages increase with tenure. The function of this back-loading is to prevent turnover—firms cannot respond to outside offers, but higher pay for higher-tenured workers makes better outside offers less likely.²³ In equilibrium, different firms start new workers at different points on the same tenure ladder. This leads to wage shopping

²² To be consistent with the basic model outlined in eqq. (1) and (2), we could incorporate accumulation of general and job-specific capital and random match quality, with all three translating multiplicatively into efficiency units of labor and hence wages. The latter would, however, depend on the bargaining protocol: that a worker loses specific capital and idiosyncratic job-match quality on leaving the firm would affect the outside option, such that the bargained wage may not be identical for each efficiency unit. (Elsby and Michaels [2013] use the Stole-Zwiebel bargaining solution.) If shocks to firm productivity also affect individual productivity multiplicatively, they will affect log wages of all workers, including new hires, approximately equally.

²³ Note that retention operates in the same direction: in these models, a decrease in firm wages following a negative shock, e.g., will lead to workers with shorter tenure disproportionately quitting (as they are more sensitive to outside offers), thus

and hence experience effects. Bagger et al. (2014) look at a related model in which firms can, however, match outside offers. Wages rise with tenure but stochastically, because the firm responds to outside offers (there is no point in back-loading). These models for tractability typically do not have firm-specific shocks of the type we have in mind (firms are identical in Burdett and Coles [2003] but have different but fixed productivity in Bagger et al. [2014]). Nevertheless, in this general class of models it would be expected that a positive firm productivity shock would increase the incentive for the firm to hire (by raising the wage profile and hence the utility of a contract offered to new hires) and to want to increase incumbent pay (to reduce turnover); thus, one would see highly correlated wage shocks across workers in the firm associated with an increase in employment (and vice versa for negative shocks).²⁴

Models in which equal treatment (in the form of equal pay per efficiency unit) is imposed or derived more straightforwardly lead to the empirical relationship hypothesized here (see, e.g., Gertler and Trigari 2009; Snell and Thomas 2010) when combined with a monopsonistic setting so that higher wages are needed to increase employment (or a competitive setting but with segmented labor markets so that positive industry shocks to productivity lead to higher industry wages and employment). In a model of on- and off-the-job search with equal treatment (so a firm cannot respond to outside offers as it pays all workers the same within a period), Moscarini and Postel-Vinay (2013) analyze the effect of aggregate shocks on wage contracts. When positive shocks occur, for example, larger firms expand more rapidly than smaller ones and contract more rapidly in downswings (however, they cannot consider idiosyncratic shocks, as the equilibrium of the model can be characterized only when firm size ranking is preserved).

If mechanisms of the type described above are generating positive wage/employment comovements sufficient to underlie the bias, we would expect to find evidence that the wage model we estimate generates a sufficient change in the present value of wages to attract new and retain older workers when positive shocks occur and vice versa with negative shocks.²⁵ It is highly unlikely that transient shocks to wages will have any effect at all on attracting labor. By contrast, permanent or highly persistent movements in a firm's wage will very likely affect its hiring, worker entry, and worker exit.

Because our sample contains large numbers of workers per firm per year, we can estimate each firm-year equal-treatment component and examine its

lengthening tenure. Likewise, if the firm is laying off workers in response to a negative shock in a last-in, first-out model, tenure will lengthen.

²⁴ A model related to that of Bagger et al. (2014) that also has investment in specific and general training is Lentz and Roys (2015).

²⁵ This is particularly true of the wage-posting models, which rely on wages to attract new workers and retain existing workers.

persistence and transience. Treating the estimated firm-year fixed effects from the FYFE specification as data, we computed the first-order autocorrelation coefficients (ρ , say) in a balanced panel regression. The ρ -values for Portugal and Germany were 0.92 and 0.99, respectively—suggesting unit root or near-unit root behavior.²⁶ Additionally, we found that the residual from the FYFE, after eliminating the equal-treatment shocks and match effects, was close to white noise (ρ -values of -0.015 and 0.115 , respectively).²⁷ It seems, then, that—in our data at least—the two sources of permanent movements in a worker's wage within the firm appear to be the match effect (permanent by definition) and the equal-treatment shocks. If idiosyncratic shocks to wages were permanent, then they too may well drive labor reallocations and hence be correlated with tenure. In that case, any attempts to estimate RTT via a reduced-form (Mincer) method would be confounded at the outset because—almost by definition—we cannot control for these shocks in standard regression analysis. Our finding that idiosyncratic shocks appear to be white noise is important, therefore; it is consistent with the idea that it is equal-treatment shocks that drive the unit root behavior in wages instead.²⁸

It would be interesting to see whether the equal-treatment shocks we have estimated correlate with firm productivity and/or its product price—a topic for future research. If this turned out to be true, our results would be consistent with the hypothesis that these shocks drive hiring and that firms share rents with their workers; following a positive (permanent) shock to its product price or its productivity, a profit-maximizing firm would hire more workers and, under equal treatment or bargaining, the newly hired and the incumbents would get a share of the improved profits. In the next section, we examine what our data have to say about the role of firm-specific

²⁶ Unit root behavior of a worker's wage within a firm is a stylized fact in labor markets. See, e.g., Buhai et al. (2014).

²⁷ Under the null hypothesis that the idiosyncratic wage components follow a unit root and that workers quit the firm when the value of this process falls below some value c^* , we can show that the empirical autocorrelation coefficient still tends to unity, so our results strongly suggest that this can be rejected. However, if the components are stationary (and again workers quit at some low threshold value), simulations suggest that the autocorrelation coefficient will underestimate the true degree of persistence.

²⁸ Topel (1991, 160–2) discusses the issue in some detail and finds no evidence that individual wage growth differences are related to contemporaneous mobility. Theoretically, if the process driving a unit root in the wage reflects general human capital, then this should not affect mobility, as outside options will move in tandem with inside returns. Likewise, timing is important: if it takes time to locate a new job after a negative wage shock, so that mobility is affected only after the period of the shock, there is no bias. Our point is that even if persistent wage shocks do affect contemporaneous mobility, so long as the persistence arises only through the equal treatment component and this is controlled for as we are proposing, there will be no bias.

wage/employment comovements in generating the differences we see in figures 1 and 2.

F. The Role of Firm-Specific Wage/Employment Comovements

If the assumptions and arguments we make in this paper are correct, then the FYFE method identifies the causal effect of tenure on wages. If initial experience is positively correlated with match quality, then the FYFE method at the very least offers a lower bound on the causal effect of tenure on wages. In this section and henceforth, we refer to differences between an estimated RTT profile from the FYFE method and another method as “bias.” Use of this term is for convenience, and it comes with the obvious caveat that it is correct terminology only if the assumptions and arguments we have made in the paper are true.

The contention of the paper is that comovement between firm wages and employment leads to a bias in the estimation of β in equation (1) using traditional methods, and we have found this bias to be negative and significant, implying that the comovement is positive. In principle, the positive comovement between firm wages and employment originating from the business cycle could be one source of bias in line with this logic. However, in our estimates we had controlled for the business cycle via the addition of general year effects. In Snell et al. (2016), we found very little effect on the bias of not controlling for the business cycle in this way. The implication is that it is firm-, locality-, or industry-specific and not systemic firm wage/employment comovements that are the primary cause of the problem.

Nevertheless, to control for firm-specific wage/employment comovements the inclusion of current firm employment in the Mincer equation is a possible alternative approach to adding FYFE.²⁹ However, it is easy to show that this will remove the bias only if the elasticities of the wage/employment relationships are identical across firms. If there is a large amount of heterogeneity in these elasticities across firms, then this will not work.

Natural vehicles to investigate these issues empirically are the MFE and FYFE specifications; they are nested and differ only because of the addition of firm-year fixed effects. We conducted two exercises using these specifications. In the first, we add (log of) firm employment and lagged firm employment to the MFE specification, allowing separate coefficients for each firm.³⁰ The addition of these terms allows us to identify each firm’s wage/employment and wage/lagged-employment elasticities and hence to see

²⁹ Buhai et al. (2014) call the impact of firm employment on wages the firm size effect (although the traditional view of the firm size effect is a steady-state notion). These traditional size effects would typically be absorbed using either firm fixed effects or match fixed effects.

³⁰ As before, we allow for a quartic in tenure and a quadratic in experience and add year fixed effects.

whether these elasticities are heterogeneous.³¹ Note that the addition of lagged employment terms is purely to obtain better estimates of the contemporaneous comovements; equation (A1) in the appendix shows that in the current context, lagged employment is of second-order importance to the bias, as discussed in Section II.D. In the second exercise, we add (log of) firm employment and lagged firm employment with a single (i.e., common across firms) coefficient. The idea here is that if the elasticities we found in the first exercise are homogeneous across firms, then we would expect the addition of these two terms to eliminate much of the RTT bias we found in figures 1 and 2. By contrast, if there is substantial heterogeneity in the elasticities, then the bias will remain. In this case, we might expect the first exercise to eliminate much of the RTT bias. For clarity, we call the specification in exercise 1 the heterogeneous specification and that in exercise 2 the homogeneous specification. Histograms of the contemporaneous wage/employment elasticities obtained from the heterogeneous specification are plotted in figure 3A and 3B for Portugal and Germany, respectively.³²

The figures show two things: first, that the elasticities are far more dispersed across firms in Portugal compared with Germany (the variance in Portugal is three times larger than that in Germany), and second, that the average elasticity—the key determinant of the bias in our analytical model—is much higher in the former than in the latter (0.06 in Portugal vs. 0.01 in Germany).³³ Given the arguments above, we would expect the bias to be larger in Portugal than in Germany—as indeed can be seen in figures 1 and 2. More pertinently for the current discussion, we would expect the RTT profiles obtained from the homogeneous specification to be close to FYFE for Germany but not for Portugal. For Portugal, we would expect only the heterogeneous specification to deliver RTT estimates close to those of FYFE instead.

To examine these statements, we compute the differences between the following RTT profiles for both countries:³⁴ (a) FYFE and the homogeneous

³¹ In the analytical model considered in the appendix, the covariances driving the biases approximate (for small changes in wages and employment) these elasticities.

³² The elasticities with respect to lagged employment were negative on average for Portugal but positive for Germany; histograms are shown in the appendix. The variance was again three times higher in Portugal than in Germany. As noted in the text, the bias formula given in the annex predicts that in neither case do these lagged comovements matter very much in determining the bias.

³³ In the general version of the model, where firms may have different sizes and rates of exit, it is a weighted average of elasticities that matter. Only when firms are the same does the bias depend on the simple average of the elasticities. As we have only large firms in our sample, we might expect them to be close in terms of size and possibly also in terms of quit rates.

³⁴ More precisely, we compute the differences between estimates of $RTT + \gamma$ from the first-stage regressions. Under our assumptions that γ is consistently estimated, these differences are the same as differences in RTT.

FIG. 3.—Contemporaneous firm wage/employment elasticities for Portugal (A) and Germany (B). A color version of this figure is available online.

specification (Hom) and (b) FYFE and the heterogeneous specification (Het). These differences are plotted for Portugal and Germany in figure 4A and 4B, respectively. For comparison purposes, we also add a line representing the bias in MFE (the gap between FYFE and MFE in figs. 1 and 2). In these graphs, the height above the X -axis represents the bias in each respective

FIG. 4.—Returns to firm tenure bias when firm employment is controlled for, for Portugal (A) and Germany (B). FYFE = firm-year interaction fixed effects; Het = heterogeneous specification; Hom = homogeneous specification; MFE = match fixed effects. A color version of this figure is available online.

specification—that is, the extent to which each respective specification fails to match the RTT generated by the FYFE specification. The line showing FYFE minus MFE is the baseline bias, the line showing FYFE minus Het is the bias from the heterogeneous model, and the line showing FYFE minus Hom is the bias from the homogenous model.

We see that for Germany the homogeneous specification eliminates all of the baseline bias (its profile lies slightly below the X -axis, which we could interpret as removing the bias). The heterogeneous specification overcorrects for the bias and also lies below the X -axis. However, using a bounds test derived in Snell et al. (2016), it is not significantly negative at any tenure value. For Portugal, things are very different. The homogeneous specification has virtually no impact on the bias—the line showing FYFE minus Hom lies practically on top of the baseline. By contrast, allowing heterogeneous comovements has far more leverage than it does in Germany—most of the bias is removed by controlling for heterogeneous (across firm) wage/employment comovements. These results are consistent with what we predicted in the earlier discussion; the extent to which the bias may be removed by adding single-coefficient employment terms depends on how homogeneous cross-firm wage/employment comovements are—where they are heterogeneous, adding employment terms with common cross-firm coefficients has no impact on the bias.

G. Equal Treatment or Unequal Treatment?

Up to now we have focused on firm wage/employment comovements driven by wage shocks that are common to all workers. It is possible, however, that wage components of new hires alone may be causing bias and that these components are not present in incumbent wages. Suppose, for example, that the firm was hiring under conditions of monopsony. Suppose also that when profitability is high, hiring is high and new hires are brought in at a wage above that of incumbents. This would drive up the firm's average wage in that year and drive down the firm's average tenure. Once again we would get downward bias in RTT. But this effect is not an equal-treatment effect—it is driven entirely by comovements between the new-hire wage and employment. This new-hire-only effect works via the same mechanism as our equal-treatment story, but if it were to be the only mechanism behind the bias, it suggests that a more efficient empirical procedure to remove it would focus on new-hire wages only.³⁵

³⁵ If new hires receive a premium or discount in wages that is permanent—as would be consistent with models of full commitment by worker and firm—these will be absorbed in match fixed effects and will not affect estimates of RTT. It is short-lived changes to the wages of a new hire, related to a firm's employment decisions, that we have in mind—e.g., a premium in the first year of employment and thereafter being paid at some standard rate. Contracting models with limited com-

In light of the previous discussion, it would be interesting to see whether augmenting FYFE with firm-year controls for newly hired workers only raises the RTT profile further. We call this augmented specification NHFY for convenience. Unfortunately, NHFY will as a by-product also remove from the RTT profile the effects of wage growth during the first year of tenure. If the RTT gradient is steeper in the first year of tenure than in later years, removing it will move the overall RTT profile upward. To allow for this and to be able to compare like with like, we also strip out the initial tenure effect from the initial FYFE specification by adding a new-hire dummy to it. Estimating FYFE with its new-hire dummy and NHFY produced RTT profiles that were within 0.2% of each other (with NHFY being slightly higher in both countries). The key takeaway of this exercise is that equal-treatment wage components are the main driver behind RTT bias; adjusting additionally for movements in new-hire wages has little incremental effect on the RTT profile.

III. RTT and Implicit Wage Contracts

The purpose of this paper has been to derive unbiased estimates of the causal effects of tenure on wages—that is, the effect on wages of experimentally increasing tenure by 1 year while keeping everything else in the economy (including, e.g., outside options) constant. To achieve unbiasedness, we have shown that it is necessary to control for equal-treatment wage shocks as well as the more traditional unobserved match effects. Wages may vary with tenure for a number of reasons, not least because of returns to experience that are general market returns, but they also respond to internal offers. As discussed above, RTT estimates may be capturing the latter—a reward to the accumulation of specific capital, either human or physical. These returns might accrue to the worker for a variety of reasons: bargaining over quasi rents, for example, or a firm being prepared to respond to outside offers (distributed independently of the value of specific capital) to keep a worker with specific capital in the firm (see, e.g., Lentz and Roys 2015).

However, the existence of quasi rents (due to specific capital accumulation or search frictions) or the ability of firms to commit may allow contracts in which wages do not correspond to marginal products in a time-invariant fashion; a classic example would be back-loaded wages to reduce turnover (as in Holmstrom 1983). If there is no observable variable with which to control for contract-driven wage growth (as would be implied by back-loading, say) then our estimates of RTT will include such effects. If this were true, only the results from a calibrated theoretical model could

mitment, e.g., often have this short-run property (e.g., Beaudry and DiNardo 1991; Rudanko 2009).

separately identify contract and human capital effects from estimated RTT. In this scenario, our bias correction would yield consistent estimates of an $RTT + \text{wage contract effect}$. Even in this scenario, these estimates would be useful raw inputs to a calibration exercise of a theoretical model that attempts to separately identify the respective effects. If there is an observable control for wage contract effects, it should be used to be able to identify pure firm-specific human capital RTT. One class of contract models that does offer observable controls for wage contracts arises from Beaudry and DiNardo's (1991) paper on implicit contracts.

Beaudry and DiNardo (1991) developed a model where—modulo firm-specific human capital—the minimum unemployment rate since the worker joined the firm (*minu*, for short) was a sufficient statistic for within-firm wage movements.³⁶ This spawned an empirical literature that added *minu* to Mincer equations to assess its importance. The findings in our paper have ramifications for this literature. *Minu* is intrinsically correlated with tenure—it falls in a weakly monotone fashion with it. Failing to control for positive firm wage/employment comovements biases the *minu* coefficient for much the same reasons as it biases RTT: higher firm wages associated with higher firm employment will lead to lower average firm tenure. Given that *minu* is negatively correlated with tenure, it would be tempting to state that this bias is positive (toward 0 for a negative coefficient). But in Snell et al. (2016) we argued that the inclusion of tenure in the regression complicates the bias, and in general it cannot be signed. Nevertheless, using the Portuguese data we showed that adding firm-year fixed effects to a specification such as MFE that also includes *minu* dramatically affects our inferences; the coefficient on *minu* falls (in absolute value) from a highly significant value of -0.93 to a borderline significant value of -0.29 . Whatever the sign of the bias in this context, equal-treatment wage components should be controlled for; at best they are unwanted noise, and at worst they cause bias.

Finally, there is a recent empirical literature that tries to establish the extent to which the contract hiring wage is sensitive to conditions at the time of hiring. In this literature, focus is on the significance of a measure of the state of the labor market (typically aggregate unemployment) and a new-hire dummy. If this variable— δ , for short—is found to have a significantly negative impact on wages, it implies that firms take advantage of poor

³⁶ In their analysis, wages will be weakly increasing with tenure since wages are increasing with the tightest labor market conditions within the current job. Hagedorn and Manovskii (2013) argue that the results are consistent with a match quality interpretation, as opposed to an implicit contract interpretation: better matches, which pay more, are more likely to survive periods of heightened job offers, proxied by cumulative low unemployment rates, and they offer evidence to support this view. Bellou and Kaymak (2016), on the other hand, find evidence for a history dependence in wages for stayers, which suggests that contracts do play a role.

labor market conditions when hiring. As with $\min_u$, Δ is negatively correlated with tenure, and once again failure to control for firm-year fixed effects would cause its estimated effect to be biased. To sum up this discussion, controlling for the effects of wage contracts is crucial to be able to identify human capital RTT, but controlling for equal-treatment components of wages is essential to get good estimates of both.

IV. Conclusion

We have shown in this paper that the positive comovement of equal-treatment wage components and firm employment causes significant bias in RTT. We showed that our equal-treatment shocks are highly persistent (unit root or near-unit root processes) and that controlling for them significantly changes the RTT estimates. This is important, as we would expect only persistent wage shocks to drive firm quits and firm hiring. We concluded that match quality and our equal-treatment shocks are two of a kind—persistent shocks to wages that impact a worker's tenure. Finally, we found that controlling for these two shocks reduces the residual error to (near) white noise. This is consistent with the argument that adding firm-year fixed effects and match fixed effects to Mincer equations is necessary to control for those wage components that are jointly endogenous with tenure. This gives us some confidence that our bias-corrected reduced-form estimates of RTT are causal. We conclude with some additional ad hoc observations arising from our work.

First, if one is purely interested in effects that vary only with year and tenure, then equal-treatment shocks are noise and removing them seems a sensible thing to do. Once match quality is controlled for, only the cross-tenure/year movements in wages are relevant to estimating RTT; components of wages that are common to workers in firm j in year t cannot add information to this.

Second, our FYFE correction allows for the possibility that firms may have heterogeneous wage and employment cotrends. Fast-growing and high-wage-growth firms would have lower average tenure and higher average wages, while slow-growing and low-wage-growth firms would have higher average tenure and lower average wages. This type of issue has been discussed before in the RTT literature, but as far as we know it has not been analyzed.

Third, in this paper we focused on MFE as a baseline specification or method. But in fact we could add firm-year fixed effects to any of the three other methods to control for the bias we have identified.

Finally, the need to control for FYFE would seem to rule out the use of small random samples of workers to obtain unbiased RTT estimates. Such samples are unlikely to contain two workers in the same firm. Just how many workers per firm are required to remove the bias effectively is unclear and a subject for future research.

References

- Abraham, Katherine G., and Henry S. Farber. 1987. Job duration, seniority and earnings. *American Economic Review* 77, no. 3:278–97.
- Altonji, Joseph G., and Robert A. Shakotko. 1987. Do wages rise with job seniority? *Review of Economic Studies* 54, no. 3:437–59.
- Altonji, Joseph G., and Nicolas Williams. 2005. Do wages rise with job seniority? A reassessment. *Industrial and Labor Relations Review* 58, no. 3: 370–97.
- Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review* 104, no. 6:1551–96.
- Battisti, Michele. 2012. Mobility and wages in Italy: The effects of job seniority. In *Labour markets at a crossroads: causes of change, challenges and need to reform*, ed. Henrik Lindberg and Nils Karlson, 205–32. Newcastle upon Tyne: Cambridge Scholars.
- Beaudry, Paul, and John 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, no. 4:665–88.
- Bellou, Andriana, and Baris Kaymak. 2016. Real wage growth over the business cycle: contractual versus spot markets. Unpublished manuscript, Department of Economics, Université de Montréal.
- Bowlus, Audra J. 1995. Matching workers and jobs: Cyclical fluctuations in match quality. *Journal of Labor Economics* 13, no. 2:335–50.
- Buchinsky, Moshe, Denis Fougère, Francis Kramarz, and Rusty Tchernis. 2010. Interfirm mobility, wages and the returns to seniority and experience in the United States. *Review of Economic Studies* 77, no. 2:972–1001.
- Buhai, Sebastian, Miguel A. Portela, Coen N. Teulings, and Aico van Vuuren. 2014. Returns to tenure or seniority? *Econometrica* 82, no. 2: 705–30.
- Burdett, Ken, and Melvyn Coles. 2003. Equilibrium wage-tenure contracts. *Econometrica* 71, no. 5:1377–404.
- Elsby, Michael W. L., and Ryan Michaels. 2013. Marginal jobs, heterogeneous firms and unemployment flows. *American Economic Journal: Macroeconomics* 5, no. 1:1–48.
- Gertler, Mark, and Antonella Trigari. 2009. Unemployment fluctuations with staggered Nash bargaining. *Journal of Political Economy* 117, no. 1: 38–86.
- Hagedorn, Marcus, and Iourii Manovskii. 2013. Job selection and wages over the business cycle. *American Economic Review* 103, no. 2:771–803.
- Hobijn, Bart, and Aysegül Şahin. 2007. Job-finding and separation rates in the OECD. Federal Reserve Bank of New York Staff Report no. 298, New York.

- Holmstrom, Bengt. 1983. Equilibrium long-term labor contracts. *Quarterly Journal of Economics* 98, no. 3:23–54.
- Lentz, Rasmus, and Nicolas Roys. 2015. Training and search on the job. Working Paper no. 21702, National Bureau of Economic Research, Cambridge, MA.
- Martins, Pedro S., Gary Solon, and Jonathan P. Thomas. 2012. Measuring what employers do about entry wages over the business cycle: a new approach. *American Economic Journal: Macroeconomics* 4, no. 4:36–55.
- Moscarini, Giuseppe, and Fabien Postel-Vinay. 2013. Stochastic search equilibrium. *Review of Economic Studies* 80, no. 4:1545–81.
- Rudanko, Leena. 2009. Labor market dynamics under long-term wage contracting. *Journal of Monetary Economics* 56, no. 2:170–83.
- Snell, Andy, Pedro S. Martins, Heiko Stueber, and Jonathan P. Thomas. 2016. Bias in returns to tenure when firm wages and employment comove: A quantitative assessment and solution. IZA Discussion Paper no. 9849, Institute for the Study of Labor, Bonn.
- Snell, Andy, and Jonathan P. Thomas. 2010. Labor contracts, equal treatment and wage-unemployment dynamics. *American Economic Journal: Macroeconomics* 2, no. 3:98–127.
- Topel, Robert. 1991. Specific capital, mobility, and wages: wages rise with job seniority. *Journal of Political Economy* 99, no. 1:145–76.