Monopsony for Whom? Evidence from Brazilian Administrative Data

Lee Tucker*

November 2, 2017
Download Latest Version Here

Abstract

The recent empirical analysis on the extent of monopsony in labor markets has been split. While studies of job separation activity suggest that firms hold high degrees of monopsony power, studies of hiring activity provide comparatively little evidence for monopsony and are generally faced with simultaneity concerns. I leverage uniquely well-suited employer-employee matched administrative data from Brazil to study this discrepancy, looking both at whether firms offer higher wages to their new employees at the times when they are growing more rapidly, and the extent to which workers' voluntary separation decisions depend on their own wage. The comprehensiveness of the data allow me to address simultaneity concerns through novel "shift-share" style instruments, as well as the inclusion of local labor market fixed effects. Although my results provide clear evidence that labor markets are imperfect even at hiring, they also strongly suggest that firms hold comparatively little monopsony power over their new hires as compared to their existing workers. I discuss the implications of several models that can potentially explain these results.

*I am grateful to my advisors, Daniele Paserman, Johannes Schmieder, and Samuel Bazzi. Thank you to Marc Muendler for generously sharing data with me. I would also like to thank Kevin Cooke, Matthew Gudgeon, Chinhui Juhn, Sara Machado, Ben Solow, Nicholas Tilipman and Wesley Yin for their helpful comments and assistance, as well as the participants of the Boston University Labor Reading Group and graduate seminars. Any errors or omissions are my own.

1 Introduction

The idea that labor markets may be imperfectly competitive—with large implications for workers' wages—is an old one (Robinson, 1969). Yet, for many years, monopsony was generally considered to be irrelevant to the understanding of the broader labor market.[1] Seminal search-based models of the labor market in the style of Burdett and Mortensen (1998) changed this view. These models showed that the conditions necessary for firms to exert labor market power over employees are both weak and sensible; the frictions associated with dynamic job search can give firms market power over their employees, and these same market frictions support the existence of wage dispersion across firms for otherwise identical workers, even if there are many firms participating in the labor market.[2] Motivated by these findings, a considerable empirical literature has arisen to support the existence of these job search frictions (Bleakley and Lin, 2012; Kuhn and Mansour, 2014; Gan and Li, 2016; Marinescu and Rathelot, 2016; Macaluso, 2017). And, accordingly, the term monopsony is now generally applied to any circumstance in which labor market frictions permit firms to pay employees less than their marginal revenue product.

Monopsony power has wide-ranging implications for the labor market, and it has been offered as an explanation for numerous well-known labor market puzzles. For example, it has been cited as a possible explanation for the lack of large disemployment effects associated with minimum wage increases (Card and Krueger, 2015; Bhaskar and To, 1999), a potential major contributor to both gender and racial wage gaps (Manning, 2003; Lang and Lehmann, 2012; Webber, 2013), and an explanation for why firms invest in their workers' general human capital (Becker, 1962; Acemoglu and Pischke, 1998; Manning, 2003). Additionally, recent empirical work using matched employer-employee administrative data suggests that increases in wage dispersion across firms have been a primary component of rising wage inequality (Card et al., 2012; Song et al., 2015; Barth et al., 2016). Since wage dispersion across firms is typically theoretically motivated by the same frictional forces that motivate monopsony power, this suggests that firms' labor market power could be increasing even as technological advances might seem to reduce the frictional costs associated with job search.

Yet, in spite of the monopsony model's broad interest, the existing empirical research on the topic has been both limited and puzzlingly inconclusive. In his well-known work *Monopsony in Motion*, Manning (2003) suggests two distinct empirical approaches for recovering

[1]For example, Manning (2003) points out that textbooks in labor economics typically made little or no mention of monopsony through at least the late 1990s.

[2]Recent research has weakened these conditions even further. For example, Card et al. (2016) show that even a static model with heterogeneous worker preferences over firm-specific amenities is sufficient to generate monopsonistic behavior.

the labor supply elasticity faced by firms; a direct approach of looking at wage setting behavior by firms, and an indirect approach of looking at workers' job separation behavior. In a simple model of firms' wage setting behavior, either approach should be able to recover the labor supply elasticity that firms face. However, in practice, these approaches have reached very different conclusions. Studies that have looked at job separation behavior have universally found that workers' job turnover decisions are quite insensitive to their own wages. These results suggest that firms hold a high degree of monopsony power over their workers, and that existing workers' wages may be marked down considerably from what they would be in a competitive market. On the other hand, studies on wage setting are much more limited, owing to the particular empirical challenges of ruling out selection and simultaneity concerns. While there is, for example, evidence of considerable firm size wage premia (Oi and Idson, 1999), the most plausible recent studies have suggested that firms either increase wages only slightly or not at all in order to attract new workers, consistent with the notion that labor markets are in fact fairly competitive (Schmieder, 2013; Matsudaira, 2014). The large discrepancy between these two approaches has generally been attributed to the difficulty of addressing the aforementioned empirical issues of selection and simultaneity in wage setting, or to issues related to external validity, rather than an indication of a genuine distinction that is not captured by the underlying model.

In this paper, I address these selection and simultaneity concerns by leveraging the comprehensiveness and unique features of employer-employee matched administrative data from the Brazilian *Relação Anual de Informações Sociais* (RAIS) program. To do so, I first adopt a high-dimensional fixed effects strategy in the style of Abowd et al. (1999). My variant of this specification incorporates local labor market by time fixed effects along with time-invariant worker and establishment fixed effects, and the inclusion of these controls rules out labor market level variation in occupational wages that might otherwise influence estimates of the labor supply elasticity. Under the assumptions of well-specified competitive local labor markets, the inclusion of these local labor market fixed effects is sufficient to mitigate the potential for over-rejection of the competitive market hypothesis due to simultaneity.

In response to remaining concerns about market misspecification, establishment-specific labor supply shocks, and attenuation bias, I then develop three novel instrumental variable strategies based on simultaneous changes in the employment of other labor inputs within the firm. These strategies are based on the recognition that labor markets are both local and occupation-specific, and they rely on a similar logic to the "shift-share" type instrumental variables that are widely used elsewhere in the economics literature to achieve identification from higher-order variation in labor market conditions (Card, 2001; Autor et al., 2013). My preferred strategy also takes advantage of a unique feature of the RAIS data that is not present in alternative data sources that have been used in prior research—the existence of firm identifiers nested within establishment identifiers, allowing me to instrument using growth in only non-local firm labor inputs.

While two-stage least squares analyses using these instruments suggest somewhat larger wage responses to firm-level labor demand shocks than predicted by OLS, these estimates still imply that wages for new workers vary only modestly depending on the rate of establishment level employment growth. Baseline IV specifications imply an average labor supply elasticity faced by firms of between 15 and 70. Further robustness check analyses based on different levels of occupational aggregation suggest that these results are unlikely to be driven by the potential for endogenous substitution across occupations.

Finally, I use an analogous fixed effects estimation approach to construct estimates of labor market separation elasticities in a linear probability model, using longitudinal data on the entire observed labor market histories for the same set of workers that is used to provide estimates of the wage response at hiring. The detailed information contained in the RAIS data allow me to look specifically at rates of voluntary separation, and to control for differences arising from worker heterogeneity, establishment heterogeneity, and local labor market conditions in a similar manner to the approach used in measuring wage responses. Consistent with existing estimates of the wage separation elasticity, I find evidence that workers' rates of voluntary separation are quite insensitive to their own pay, implying an average labor supply elasticity faced by firms of 0.6 to 0.8. Heterogeneous effects regressions further demonstrate that these estimates vary along several dimensions of local labor market conditions that would be predicted to have relevance in search-based models of the labor market. However, estimates of the wage premia offered to new hires in times of growth do not exhibit the same heterogeneity.

In short, this paper confirms the basic finding that estimates of the labor supply elasticity vary considerably depending on the way in which they are measured. It also rules out several alternative explanations related to simultaneity, selection, and external validity. What remains is a puzzle: why do firms behave as if labor markets are fairly competitive when setting wages, even though workers' separation behavior implies that firms have substantial market power over their existing workers? And, if firms have only *ex post* monopsony power, what features of the labor market permit this power to develop?

The paper proceeds as follows: Section 2 provides a brief overview of the existing literature on the topic of monopsony, Section 3 describes both of the primary empirical specifications used in this paper along with the IV strategies used, Section 4 describes the RAIS data along with household survey data used for analyses of heterogeneous effects, Section 5 provides baseline elasticity estimates using both specifications, Section 6 provides a variety of robustness checks of the new worker wage specification, Section 7 reports the results of heterogeneity tests that are explicitly designed to assess the extent to which reported elasticity estimates vary in ways that suggest that they are a function of labor market search frictions, and Section 8 concludes the paper by discussing potential models that may better describe the observed discrepancy in elasticity estimates conducted on new hires vs. existing workers, as well as their implications.

2 Background on Empirical Estimation of Monopsony Power

The traditional static model of labor market monopsony is a direct analogue to the more well-known model of product market monopoly. In this model, monopsonistic firms' marginal labor costs are higher than the equilibrium wage, because firms are assumed to be unable to wage discriminate among workers, and so each additional worker employed serves to increase the market equilibrium wage. Profit-maximizing monopsonistic firms pay each worker a marked down fraction of their marginal revenue product $MRP(L)$, where the extent of the markdown is a function of the labor supply elasticity faced by the firm. In particular, in a standard profit maximization framework with unit labor costs and labor supply elasticity ϵ^L_{Sw}, the equilibrium wage paid to each worker will be:

$$w = \frac{MRP\left(L\right)}{1 + \frac{1}{\epsilon^L_{Sw}}} \tag{1}$$

The static model of monopsony underscores the potential importance of monopsony to labor market outcomes. However, the model also has clear limitations. In particular, the model does not provide clear insights into how one should go about estimating ϵ^L_{Sw}, and the methods that it suggests, such as looking at the firm size wage premium, could plausibly be explained by match premia or other sources of increasing returns to scale (Kremer and Maskin, 1996; Oi and Idson, 1999).

In response to these limitations, since Manning (2003), most models of labor market monopsony have focused instead on a simple dynamic framework. Like the static model, these models begin with the basic assumption that firms set wages in order to achieve a desired firm size.[3] Most of these models use a simple dynamic formulation in which the

[3]The theoretical literature generally does not distinguish between establishments and firms. However, if separate establishments face different local labor markets, then it is potentially more appropriate to consider this phenomenon at the establishment level. In the RAIS data, I observe both establishments and firms, but

present level of a firm's employment L_t is a function of prior employment L_{t-1}, the number of recruits from both employment and unemployment as a function of the wage offered, and the rate of separations from employment to other employment or to unemployment as a function of the wage offered. This can be written as:

$$L_t\left(w_t, L_{t-1}\right) = R\left(w_t\right) + \left[1 - s\left(w_t\right)\right] L_{t-1} \tag{2}$$

where $R\left(w_t\right)$ is the number of recruits to the firm, and $s\left(w_t\right)$ is the corresponding separation rate, each of which is a function of the wage offered to all workers.

This equation implies that the overall labor supply elasticity faced by the firm can be broken into two or more separate elasticities, including an elasticity of new recruits to the establishment, and an elasticity of job separations from establishments. The recruitment and separation elasticities can in turn be broken out into elasticities of recruits from employment and non-employment, and a separation elasticity to employment and non-employment, which may be expected to differ under a Burdett-Mortensen type model with both individual reservation wages and on-the-job search. Indeed, depending on the exact assumptions that one makes with respect to the nature of on-the-job search and steady state firm size and market unemployment, it is possible to construct and consider various formulations of this equation that relate short-term and long-term supply elasticities. In empirical practice, however, most authors have chosen to make simplifying assumptions. The simplest of these is to assume that the firm is at a steady state level of employment, so that $L_t = L_{t-1} = L$. Then, with some algebra, it is straightforward to show that the overall labor supply elasticity is an additively separable function of the recruitment elasticity ϵ_{Rw} and the separation elasticity ϵ_{sw}:

$$\epsilon^L_{Sw} = \left(\epsilon_{Rw} - \epsilon_{sw}\right) \tag{3}$$

If it is also assumed that the overall labor market is in a steady state, then each recruit to a firm is a separation from another firm. This implies that $\epsilon_{Rw} = -\epsilon_{sw}$, and so the above can be simplified to:

$$\epsilon^L_{Sw} = -2\epsilon_{sw} \tag{4}$$

While the steady state assumptions needed to get this particular equation are strong, Equation 4 implies a clear relationship between the elasticity of voluntary separations from establishments and the overall labor supply elasticity faced by those establishments. If the

the unit of observation at which employment growth is measured is always the establishment.

steady state assumptions are relaxed, then the short-run relationship between the separation elasticity and the overall labor supply elasticity may differ somewhat, although the scope for such differences is modest.[4]

In light of this tractable model, since Manning, most empirical tests of monopsony have focused solely on estimating job separation elasticities in order to calculate a single overall labor supply elasticity. That is, these papers test how responsive a firm's existing workers are to variation in their level of pay. Several of these papers take advantage of exogenous sources of wage variation from regulatory changes, such as wage premia offered to teachers in certain schools, or wage differences arising from differences in job title within a firm (Ransom and Oaxaca, 2010; Falch, 2010; Ransom and Sims, 2010). While these papers provide clean identification of separation behavior, they raise particular concerns about external validity, since they typically study heavily regulated and/or unionized industries and occupations which are more likely to have suitably exogenous variation to exploit. More recent work has begun to take advantage of matched employer-employee administrative data, generally using high-dimensional fixed effects approaches similar to the one specified in this paper (Barth and Dale-Olsen, 2009; Hotchkiss and Quispe-Agnoli, 2009). A handful of papers have used other approaches to estimate separation elasticities, such as looking at historical firm-level data or household survey data (Booth and Katic, 2011; Depew and Sorensen, 2013).

Although their exact estimates vary, these papers universally find evidence that workers are at best only somewhat responsive to their own wages in deciding whether to quit their jobs. Labor supply elasticity estimates in the range of 1-4 are most typical, with a few outliers in either direction. Nevertheless, this literature has suggested that such differences may have quite substantial policy implications, arising from profit maximization behavior as noted in Equation 1. For example, several papers that have looked at the differences in separation elasticities between men and women have suggest that large proportions of the gender wage gap may be explained women's lower wage elasticity of separation (Barth and Dale-Olsen, 2009; Ransom and Oaxaca, 2010; Hirsch et al., 2010; Webber, 2013), and similar research suggests that monopsony may also explain much of the wage gap between documented and undocumented immigrants (Hotchkiss and Quispe-Agnoli, 2009).

A much smaller literature has attempted to isolate the elasticity of the labor supply by looking at the relationship between changes in establishment-level employment and the wages offered to employees. Although this may seem to be a more direct method of testing for monopsony, and the value of such an approach has been long acknowledged, few studies exist because of the particular challenges associated with simultaneity in wage setting. To my

[4]Hirsch et al. (2017) provide a useful derivation of the relationship between the separation elasticity and the overall wage elasticity of the firm's labor supply when the steady state constraints are relaxed.

knowledge, only three papers have made attempts to instrument for firm-level labor demand. In his study of the wage premia paid by new firms, Schmieder (2013) finds evidence that this wage premium is primarily attributable to the higher growth rate of these firms. Using firm age as an instrument for growth, he finds evidence of a small upward slope to the labor supply curve; with an estimated elasticity of approximately 46. Matsudaira (2014) studies a policy change in California in which new minimum nurse staffing regulations served as an exogenous shock to firm-level labor demand whose size varied depending on distance from the new legal threshold. This IV strategy provides no evidence for monopsony power, and indeed the point estimates suggest that firms that grew faster lowered wages, though the author suggests that worker selection may be a possible concern. Bellon (2016) develops an IV strategy using French administrative data that instruments using exposure to variation in product demand for exporting firms; his estimates suggest that these firms face a recruitment elasticity of approximately 10, implying an overall labor supply elasticity that is larger than that. Collectively, while these studies provide some evidence for imperfect labor markets, they imply that the labor market is far more competitive than studies of separation behavior would suggest.

A particular contribution of this paper is to expand the empirical study of imperfect labor markets in the context of a developing country with a large informal sector. Because high-quality employment data in developing economies is scarce, little research exists on the extent of monopsony or other labor market frictions in these contexts. Brummund (2011) estimates the labor supply elasticity for manufacturing firms in Indonesia utilizing structural methods adapted from the industrial organization literature, and Rivera (2013) looks at worker transition behavior using the same data source used in this paper. Satchi and Temple (2009) calibrate a model using Mexican data which suggests a relationship between formal sector workers' bargaining power and the size of the informal sector.

This paper also relies upon, and contributes to two other bodies of literature. The first is the literature on the use of high-dimensional fixed effects models in economics, in particular with respect to the use of more than two classes of such effects to estimate models that were previously computationally infeasible. While this literature began with the goal of examining the degree of assortative matching between workers and firms (Abowd et al., 1999, 2002; Card et al., 2012), recent advances in the algorithms available for use have permitted the consideration of additional dimensions of effects, such as job title or match effects (Torres et al., 2013). Recent algorithmic improvements have also substantially reduced the computation time needed to estimate these models (Correia, 2016). To my knowledge, this is the first paper to use local labor market fixed effects in conjuction with establishment fixed effects to address potential simultaneity issues. The second is the literature on the degree to which human capital is occupation specific (Gathmann and Schönberg, 2010; Guvenen et al., 2015; Macaluso, 2017). In particular, the heterogeneity tests shown in Section 7 are supportive of the hypothesis that workers consider occupation-level labor market conditions in their decision of whether or not to leave their employer voluntarily.

3 Overview of Empirical Specifications

The goal of this paper is not only to estimate the elasticity of the labor supply in the Brazilian context, but also to better understand why estimates of the labor supply elasticity faced by firms vary so considerably depending on whether they attempt to estimate the labor supply directly or via separation behavior. Accordingly, I develop three basic estimation strategies, each of which is outlined below.

3.1 Wage-Setting Specification

The first strategy that I adopt asks the question "do establishments offer higher wages to their new employees when they are growing more quickly?" In this section of the paper, the basic equation that I estimate is of the form:

$$\log w_{iomt} = \alpha G_{ot,j(i,t)} + X_{it}\beta + \delta_{j_{(i,t)}} + \theta_i + \psi_{omt} + \epsilon_{iomt} \tag{5}$$

where w_{iomt} is the log wage income of worker i, employed in occupation o, in local labor market region m, and in year t. The explanatory variable of interest, $G_{ot,j(i,t)}$, is a measure of occupation-level employment growth in the establishment at which worker i is employed at time t, denoted by $j(i,t)$. X_{it} are time-variant worker characteristics, $\delta_{j_{(i,t)}}$ are a full set of establishment fixed effects, θ_i are worker fixed effects, and ψ_{omt} are fully interacted occupation-region-year fixed effects. Because this specification regresses log wages on a measure of overall employment growth, the coefficient of interest, α, can be interpreted here as the inverse of the labor supply elasticity faced by the firm.[5]

In keeping with other literature on establishment-level employment changes, I adopt an occupation-level version of the index of employment change used by Davis et al. (1998), (hereafter referred to as the DHS index) defined as $2 \times \frac{L_t - L_{t-1}}{L_t + L_{t-1}}$. For small changes in employment, this index is approximately equivalent to a more traditional measure of percentage change. However, the index is also defined even when prior-period occupation-level employ-

[5] It must be noted that the estimates provided here do imply some path dependence: a firm that grows quickly in one period and slowly in the next will be predicted to offer different wages in each period than a firm that grows consistently over two periods, *ceteris paribus*.

ment is zero, and its bounded and symmetric nature addresses potentially large asymmetries between small employers who grow or decline by the same amount.

The specification of Equation 5 incorporates several types of fixed effects simultaneously, and this rules out endogeneity arising from time-invariant employer heterogeneity, worker selection, or simultaneity among competitive labor market firms.[6] However, it also means that many types of variation in wages do not contribute to the identification of the parameter of interest α, i.e. the estimated inverse supply elasticity. These include the variation between establishments in their overall average wage levels, the variation between individual workers in their overall average observed wage levels, and the variation in the average wages paid to each occupation in each local labor market in each year.

What remains are two distinct sources of variation which correspond to different sources of variation in establishment-occupation employment growth. The first is variation across time, within occupation, in the growth rate of employment within particular establishments. If, say, occupation o represents accountants, and establishment $j(i,t)$ is an outpost of a particular accounting firm, then an estimate of α will be positive if the firm pays higher wages to new accountants at the time periods in which its employment of accountants is growing most rapidly. The second source of variation on which α is identified is simultaneous variation in the growth rate of different occupations within the same establishment. That is, in the same establishment considered previously, α will also be positive if the wage premium offered to accountants is greater than the wage premium offered to janitors in the time periods when the establishment is growing more quickly in its employment of accountants than janitors, and vice versa. In robustness check regressions presented in Section 6, I am also able to individually isolate these two sources of variation.

While the restricted sources of variation used to identify the parameter of interest α in this specification help to rule out selection concerns, there may still remain substantial concerns related to simultaneity bias in estimation of firms' labor supply elasticities. Figure 1 illustrates this concern in a standard static model. The firm shown in this figure possesses monopsony power, as demonstrated by the upward-sloping labor supply curve that it faces. Our empirical goal is to estimate the elasticity of this firm-specific labor supply curve. And, if we are able to isolate solely shifts in firm-specific demand, then it is straightforward to trace out the shape of the firm's supply curve from equilibrium wages and employment. For example, a shift in firm-level labor demand from D_1 to D_2 would yield an estimate of α based on the line connecting equilibrium points A and B.

[6] The strategy of incorporating overlapping worker and establishment fixed effects, developed by Abowd et al. (1999) and generally referred to as the AKM model of wage setting, has a long history in the literature, and is now typical of recent research conducted using matched employer-employee data.

Figure 1: Potential Simultaneity Bias in a Monopsonistic Firm Without Local Labor Market Fixed Effects

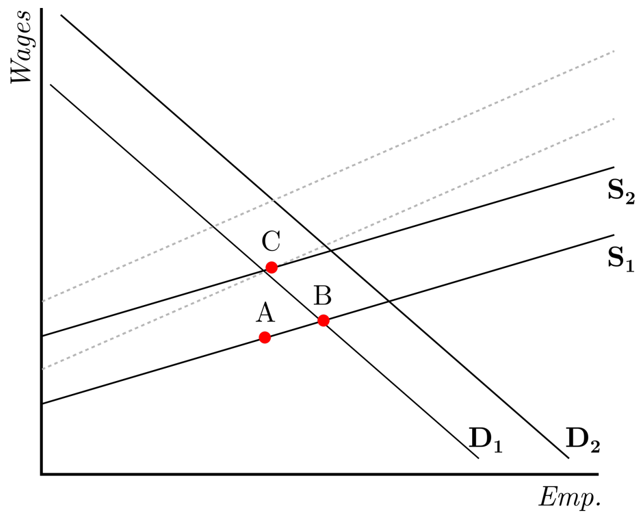

However, in practice, isolating firm-specific shifts in demand is difficult to do in a non-experimental setting for a very simple reason: the labor supply curve faced by the firm is itself a function of all other firms' labor demand. If, for example, there exists a positive correlation between shifts in firm-specific labor demand and the labor demand of other firms in the market (as shown in the figure), then we may observe large changes in wages even with little change in overall employment, leading to upwardly-biased estimates of the labor supply elasticity α. Conversely, a negative correlation between firm-specific supply and firm-specific demand, such as would be found if firms adjust to increase their usage of labor inputs when it is least expensive to do so, would lead to downwardly biased estimates of the firm-specific labor supply elasticity. The ideal test of monopsony power would identify wage responses to purely random variation in firms' product demand, which would yield variation in labor demand that is uncorrelated with local labor market conditions. But, truly exogenous firm-level instruments for labor demand are rare, and so concern about simultaneity has been the major impediment to progress on this empirical question.[7]

To begin to address these simultaneity concerns, the regression specification in Equation 5 incorporates occupation-region-year fixed effects ψ_{omt}, which I will hereafter refer to as "local labor market" fixed effects. Since the administrative data used in this study provide a comprehensive portrait of the formal sector labor market for each occupation o, and because they incorporate suitable geographic detail, the local labor market effects ψ_{omt} simply control flexibly for all variation over time in the average wages paid to new workers locally in that

[7]For example, in *Monopsony in Motion*, Manning opines that "[p]rogress seems to be dependent on finding a good firm-level instrument" (Manning, 2003, p. 96).

occupation. It is straightforward to show that if local labor markets are exactly specified by the fixed effects and if there do not exist establishment-specific labor supply shocks, then the inclusion of these local labor market fixed effects is itself sufficient to preclude over-rejection of the competitive labor market hypothesis due to simultaneity.[8]

However, in practice, even the inclusion of local labor market fixed effects may be insufficient to address all simultaneity issues. For example, occupational labor markets may exist at a finer level of specificity than can be observed in the data, or labor market activity may be more localized than is characterized by the fixed effects. If this is true, then inclusion of coarse fixed effects will fail to root out correlation between firm-specific demand and market-level supply. There could also simply exist establishment-specific labor supply shocks that are correlated with establishment demand shocks. Finally, although it is not a source of simultaneity, there may be considerable classical measurement error in firms' observed employment growth rates relative to their desired growth rates at the time of hire, in part because wages and employment are only measured at the end of each year, rather than at the time of hire. This could be a source of attenuation bias in OLS estimates.

In response to these concerns, I develop three novel instrumental variables strategies, each of which leverages the comprehensiveness of the administrative data used here. Each strategy relies on a similar "shift-share" logic to numerous studies elsewhere in the economics literature, and in particular, these instruments rely on the well-known existence of scale effects in the use of labor inputs. The first strategy is to instrument for establishment occupational employment growth $G_{ot,j(i,t)}$ using $G_{-ot,j(i,t)}$, or the growth in establishment-level employment in occupations other than occupation o. The second strategy, which leverages a particular advantage of these data, is to instrument for employment growth using $G_{ot,f(-m,i,t)}$, where $f(-m,i,t)$ is defined as the firm f of which establishment $j(i,t)$ is a part, excluding any other local establishments. The third strategy simply combines the features of the first two strategies, using $G_{-ot,f(-m,i,t)}$ as an instrument, the growth in non-local employment of other occupations.

Since each of these IV strategies involves using growth in the use of other labor inputs as an instrument for growth in employment of one's own occupation in one's establishment, each of these instruments achieves relevance from the existence of scale effects in production. That is, in response to firm-level product demand or productivity shocks, firms scale

[8] In most traditional static models of the labor market, wage dispersion across firms can arise due to hedonic considerations, but it does not arise as a result of monopsony power. Therefore, if labor markets are competitive, the imposition of appropriate ψ_{omt} fixed effects will ensure that OLS regressions recover an estimated inverse elasticity of zero in expectation. However, depending on the exact assumptions made regarding the nature of equilibrium wage dispersion, OLS estimates could still provide a downwardly biased estimate of the inverse elasticity faced by the firm even if firm-specific labor supply curves are upward sloping. Such results would imply a more elastic labor supply than would actually be the case.

their overall production up or down. The exclusion restriction, meanwhile, requires that the extent to which firms scale their production of other inputs *relative* to the worker's own input is uncorrelated with omitted firm-year specific variables that also influence worker's wages. Notice that the inclusion of establishment and local labor market fixed effects considerably narrows the scope of potential violations of the exclusion restriction. So, for example, firm-level differences in the elasticities of substitution across inputs are not a threat to identification, as long as those firm-level differences are time invariant. Time-variant changes in the relative usage of inputs due to local aggregate labor supply shocks also do not threaten identification.

The primary remaining threat to identification is the potential for endogenous time-varying degrees of input substitution within firms, and there are two primary ways in which such substitution behavior could arise. The first is if employers switch production processes in order to substitute workers in supply-elastic occupations for those in supply-inelastic ones. While, in the short run, the scope of this type of substitution may be limited, this concern is particularly likely to hold in the long run as firms are able to adapt their production processes. To the extent that such substitution is a concern, we should specifically expect it to bias regression estimates of the inverse elasticity downward, implying a more elastic labor supply curve than may actually be the case. However, even if such substitution is a concern, to the extent that my IV estimates capture the overall wage response to changes in product demand, the estimates reported here may be more reflective of the particular margin of labor market adjustment that is most relevant to policymakers.

The second and more concerning threat to identification from endogenous input substitution could arise if supply shocks in substitute inputs are correlated with firm-specific supply or demand shocks for the worker's own occupation. While the inclusion of local labor market fixed effects narrows the scope of such concerns, it is not implausible that such correlations could exist within local labor markets, especially if occupations are defined narrowly. It is for this reason that my preferred identification strategy uses only growth in the employment of non-local labor inputs within the same firm. Robustness check regressions shown in Section 6 further suggest that this type of endogenous substitution is unlikely to be a substantial source of bias in my baseline IV estimates.

Finally, it must also be noted that the AKM model implicitly assumes that wages can be decomposed into additively separable establishment and worker fixed effects. That is, the model assumes that there are no "match effects" by which high-wage workers receive a particular wage premium when they are matched with a high wage or particularly suitable firm. Several recent papers have examined this assumption in various contexts, notably by looking for asymmetry in the changes in wages from workers who move from high-wage firms to low-wage firms. Recent research by Lavetti and Schmutte (2016) has performed these tests using the same Brazilian RAIS data that I use in my analysis over the period 2003-2010, and has found little evidence of match effects. In this paper, the presence of "local labor market" fixed effects further implies that establishment wage premia and any time-specific local labor market wage premium are additively separable, and that time-invariant individual wage premia are additively separable from time-variant local labor market fixed effects. That is, high-wage workers or high-wage firms do not disproportionately benefit from being in a place where wages are high or low at a given point in time.

3.2 Job Separations Specification

In my second empirical strategy, I ask the same question that has been regularly asked in the literature to date: "how responsive are workers to their own pay in their decisions to separate from their current firm?" To answer this question, I estimate a linear probability model of the form:

$$S_{iomt} = \alpha \log w_{iomt} + X_{it}\beta + \delta_{j_{(i,t)}} + \theta_i + \psi_{omt} + \epsilon_{iomt} \tag{6}$$

The binary outcome variable $S_{i,j(i,t),t+1}$ equals 1 if worker i is reported to have voluntarily separated from establishment $j(i,t)$ in the year subsequent to his observation in period t, and the parameter of interest is the coefficient α, estimated. Since this specification regresses a binary outcome on a measure of log wages, the separation elasticity faced by the firm in its choice of wages is calculated as $\frac{\alpha}{\overline{S}}$, where $\overline{S}$ is the mean rate of voluntary separation in the sample. If the steady state assumptions and other assumptions described in Section 2 hold, then this number may be multiplied by -2 to produce an estimate of the overall labor supply elasticity faced by the firm. As with the wage-setting specification, this specification includes worker fixed effects $\delta_{j_{(i,t)}}$, establishment fixed effects θ_i, and local labor market fixed effects ψ_{omt}.

Although prior research has estimated separation elasticities using overlapping establishment and worker fixed effects, the inclusion of local labor market fixed effects again addresses a particular concern with this estimation method. Specifically, overall local labor market conditions are likely to be positively correlated with both wages and with the overall probability of voluntary separation. [9] An estimation strategy that does not address this correlation will produce estimates of the labor separation elasticity that are biased upward

[9] For example, publicly available data in the U.S. on job quits from JOLTS shows quits to be highly pro-cyclical.

(toward zero), which would imply that the labor supply to firms is more inelastic than may actually be the case. To date, to my knowledge, most separation elasticity estimates using matched administrative data have either made no direct attempt to control for this source of bias, or they have chosen to instrument for current-period wages using initial-period wages. Since, in this specification, the introduction of local labor market fixed effects addresses any sources of variation over time that are local labor market specific, the results shown in this paper are robust to business cycle considerations, including local and occupation-specific business cycle considerations.

A second empirical concern in the estimation of separation elasticities is the existence of firm-specific wage shocks. If, for example, workers' wages dip in a single period because of a negative firm-level shock, then those workers may be comparatively unlikely to separate voluntarily if they believe that their wages will recover in the next period, because by separating they forego the opportunity to earn future wages as a tenured worker within the firm. The presence of such short-term wage variation is therefore a source of bias toward zero in elasticity estimates relative to workers' long-run sensitivity to their own wages. I adopt two strategies to address this concern. The first is to use initial-period worker wages as an instrument for current period wages, limiting the sample to only individuals who are reported to have more than one year of tenure. The second is to replace establishment fixed effects $\delta_{j_{(i,t)}}$ and local labor market fixed effects ψ_{omt} with a full set establishment-occupation-year fixed effects $\nu_{omt,j_{(i,t)}}$. This second strategy has the effect of identifying the separation elasticity using only variation in wages within establishment-occupation groups, so estimates of α in this specification are based on workers' sensitivity to their own wage relative to what their coworkers are paid.

4 Description of Data Sources

In this section of the paper, I describe the primary source of data for the analyses of this paper, matched employer-employee data from the *Relação Anual de Informações Sociais* (RAIS). I also briefly describe my more limited use of household survey data from the *Pesquisa Nacional por Amostra de Domicilios* (PNAD). Additional details regarding my use of household data are contained in the Appendix.

4.1 RAIS

In Brazil, all firms that are formally registered must report information on their employees in each year to the Ministry of Labor for the provision of an annual wage supplement. This dataset is known as RAIS, and it provides a comprehensive annual census of formal sector employment in Brazil. This information includes a unique identifier that is longitudinally consistent, making it possible to track individual workers over time, even as they switch establishments.[10] It also contains basic demographic information on individuals including their age, sex, nationality, contracted hours per week, and level of education. Job tenure is reported in weeks.

In addition to longitudinal data on workers, the RAIS data include a unique tax identifier for each establishment. This tax identifier nests within itself a firm code, so that both firms and establishments can be identified from the data. Additionally, although the units of observation in the RAIS data are person-years, the data also contain several establishment-year level variables. These include two measures of industry classification (CNAE and IBGE), legal classification that indicates government or private ownership, and geographic information at the level of individual municipalities.[11]

While other analyses of the labor supply to the firm have utilized matched administrative data such as IAB data from Germany (Schmieder, 2013; Hirsch et al., 2017) and LEHD data from the United States (Webber, 2013), the Brazilian RAIS data have key advantages over these datasets. Most notably, the identity code for each establishment in RAIS nests within it the identity code of the firm of which it is a part. This, combined with the detailed geographic data that RAIS provides, allows me to pursue a novel strategy for isolating firm-level labor demand shocks by using employment growth in other non-local establishments of the same firm as an instrument for growth in employment within each establishment.

A second advantage of the RAIS data is that I observe not only each worker's full employment spell over the period 1995-2014, but at the time of their separation from a job spell, I observe a code that indicates their reason for separation. Data on the reasons for job separation are collected because the benefits to which a formally registered employee is

[10]Brazilian private sector workers who are engaged in formalized private sector employment receive a unique identification number through the *Programade Integração Social* (PIS) program. A worker's PIS identifier is used to identify them at all employers and it does not change. As part of the PIS program, employers contribute to a special bank account, administered by the government owned bank Caixa, that is set aside for each employee. The full balance of this account can only be accessed upon retirement, old age, illness, death, or disability. However, formally registered employees who meet certain basic criteria are also eligible to receive an *Abono Salarial*, or annual wage supplement. The supplement is equivalent to an additional month's pay at the minimum wage level, and it is prorated for individuals who were not formally employed for the full year. The agency that administers the *Abono Salarial* uses the data from RAIS to determine eligibility. Therefore, the scope for non-compliance in reporting is very low, and the RAIS data provide a comprehensive census of formal sector employment in each period (cai).

[11] There are a small number of establishments which report different municipalities for different individuals within the same establishment. However, over 99% of establishments report a single municipality for all workers employed for them in each year. Therefore, in calculating establishment-level and establishment-occupation level indices, I use the modal municipality code reported as a measure of the establishment's location.

entitled vary depending on the reason for separation. This allows me to identify voluntary job separations separately from non-voluntary job separations to construct separation elasticity estimates, and therefore to consider the ways in which these two primary estimation methods may suggest very different results, even using the same sample of workers.

The full set of individual job separation codes (both voluntary and involuntary) are shown in Appendix Table 9. In that table, I also list the unconditional probability that a worker is reported to have separated with that separation code reported in the year subsequent to observation, for the longitudinal labor market sample of workers used to calculate separation elasticities. Notably, only about one in four separations is reported to be voluntary in nature. This observation is consistent with the fact that the foregone social welfare benefits from voluntary separation are relatively large, making it more costly for individuals to voluntarily separate to unemployment (Rivera, 2013).

A broader feature of the Brazilian context, one which provides both advantages and disadvantages, is the presence of a large informal sector in Brazil that is not observed in the RAIS data. On one hand, the presence of comprehensive formal sector data alongside the informal sector allows me to examine the extent to which the informal sector influences formal sector labor market responses, as I do in Section 7. On the other hand, I cannot infer that an individual who is not observed in the dataset at any point in time is not employed, as they may have entered informal sector employment. Similarly, I cannot assume that any establishment that first appears in the RAIS data in a given year does so because it is a new establishment, unless I also observe other establishments of the same firm in prior periods. And, as Brazil's labor market has become increasingly formalized, the composition of the formal sector may have changed, restricting my ability to identify changing labor market responses over time. I am able to use household survey data from the PNAD dataset (described below) to observe the local extent of the informal labor market with some geographic and occupational specificity.

In order to limit the potential scope of concerns arising from life cycle factors, gender differences, or other concerns, I limit my sample in the following ways. I include only men, ages 25-54, who are reported to to have been contracted full-time (defined here as 30-50 hours per week), and who are reported to be employed on December 31st of the year in question. If multiple jobs are reported for one individual, I include only the highest-paying job in each year.[12]I drop individuals whose reported income is zero, and individuals without

[12] In my separation elasticity estimates, the measure of separation that I use is based on reported separations, not on non-observation, so these estimates are not affected by the possibility that one could hold two jobs continuously while having a different one report the highest income in each year. My separation estimates should be likewise unaffected by issues related to seasonal separation and rehiring. Additionally, in wage regressions on the new hire sample, I include linear and quadratic terms in tenure to address any po-

a reported unique identifier. I exclude observations from government entities, establishments that are reported to be state owned, and non-profit entities. And, for my sample of new worker observations, I include only individuals who are reported by their employer to have less than one year of tenure at the time of observation. Although I have access to RAIS data from 1986-2014, the above restrictions also require me to restrict my attention to the 1995-2014 period in which all necessary variables are reported.[13]

A full sample of the RAIS data, even with the restrictions indicated above, is very large. Even with the new worker restriction, for example, there are over 120 million new worker observations over the period 1995-2014. Since the methods used in this paper are computationally intensive, I have constructed a random sample of 10% of Brazilian micro-region codes. The choice to sample at the local labor market level is deliberate. The high-dimensional fixed effects regressions that I run in this paper rely on the the movement of individuals across establishments as they change jobs for estimation. Individual and establishment fixed effects are estimated simultaneously on the "connected set" of individuals. So, a sample constructed at either the individual level or at the establishment level would greatly reduce the statistical power of this estimation method by reducing the precision of all fixed effect estimates. In contrast, by conducting a sample at the local labor market level, all individuals who do not relocate across metropolitan areas, or who relocate between labor markets in the sample, are still observed for the full duration of their formal sector employment.

Figure 2 shows a map of Brazil, in which the 48 local labor market regions that I have sampled are indicated in green. While a few of these regions in the sparsely populated Amazon are quite large in geographic area, most are comparatively compact, and they are concentrated in the Northeastern and Southeastern portions of the country in which population is most concentrated. Because the micro-region level of aggregation maps most closely to the traditional definition of a metropolitan area based on overlapping patterns of economic activity, this level of aggregation has also been used by other recent literature looking at regional effects in Brazil (Dix-Carneiro and Kovak, 2017). Of the 48 microcodes sampled by me, the largest in population by far is the micro-region containing the city of Belo Horizonte, which as of 2017 is the 6th largest city and 3rd largest metropolitan area in Brazil (Saraiva).

Table 1 provides some basic summary statistics on the new worker samples that are

tential issues related to seasonality in wages. This method of addressing multiple job holder is also consistent with what has been done in other recent research using RAIS data (e.g. Helpman et al., 2017).

[13] An additional concern with looking at wage setting behavior prior to 1995 is that in the early 1990s, Brazil experienced very high rates of inflation, at times exceeding 1000% per year. In 1994, the *Plano Real* sought to reduce inflation, resulting in the adoption of a new currency that was loosely linked to the U.S. dollar. The use of data from 1995-2014, therefore, explicitly excludes the period prior to the adoption of the *Plano Real*.

Figure 2: Map of 10% Random Sample of Brazilian Micro-Regions

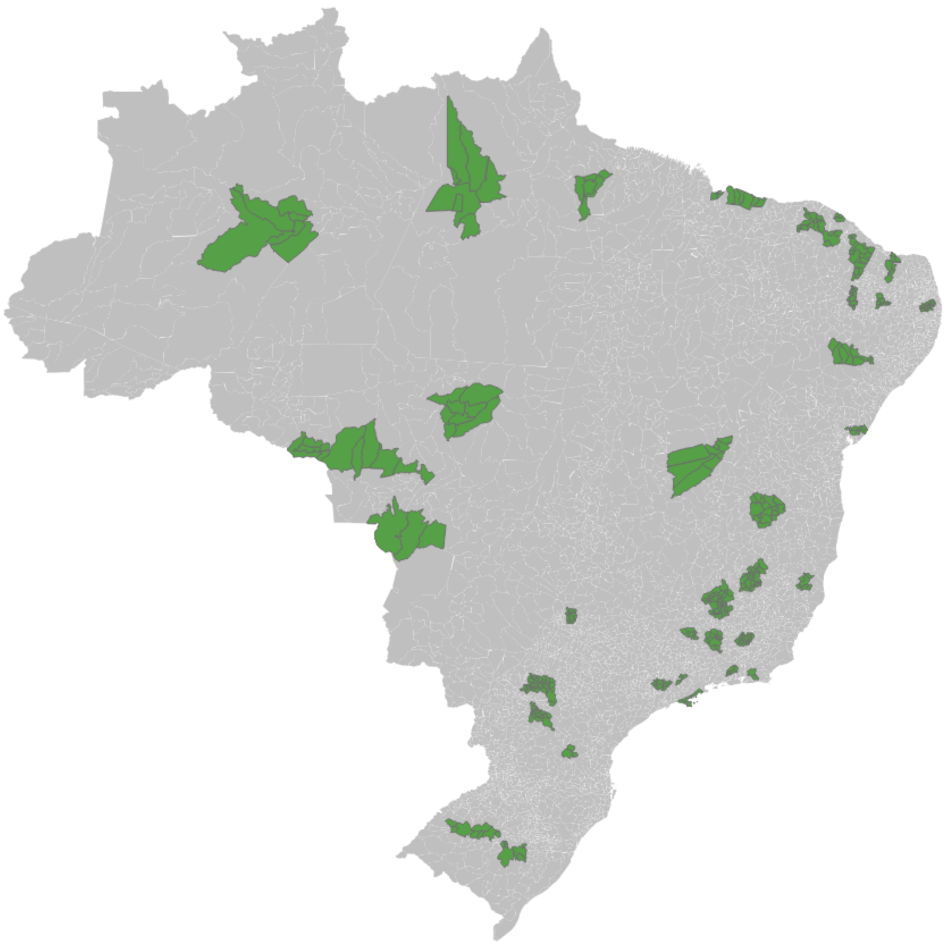

Notes: 48 sampled Brazilian microcode regions are indicated in green.

Table 1: Summary Statistics for New Worker and Work-History Samples

	New Worker Sample		*Work-History Sample*
	All Estabs.	Multi-Region & Estab. Firms	All Estabs.
Observations	12,415,479	4,031,119	30,166,458
Establishments	401,635	88,455	493,446
Firms	309,442	19,053	382,480
Median Estab Size.	96	298	162
Median Estab Occ. Size	16	42	24
December Wage Income	3.159	4.273	4.110
	[4.16]	[5.84]	[6.31]
Age	35.28	35.04	35.25
	[7.83]	[7.72]	[11.49]
Tenure (Years)	0.322	0.345	1.08
	[0.267]	[0.274]	[1.50]
Education			
Less than HS	0.526	0.472	0.479
HS Grad	0.396	0.403	0.401
Some College	0.022	0.032	0.025
College Grad	0.055	0.093	0.095
Location			
Belo Horizonte	0.541	0.515	0.544
São Luís	0.085	0.088	0.087
Londrina	0.058	0.059	0.057
Bragança Paulista	0.048	0.047	0.044
Others	0.269	0.290	0.267
Occupation			
Prof. or Managerial	0.116	0.126	0.142
Techn. or Supervisory	0.147	0.145	0.211
Other White Collar	0.142	0.151	0.148
Skilled Blue Collar	0.450	0.452	0.362
Unskilled Blue Collar	0.145	0.127	0.137

Notes: From RAIS, 1995-2014. New worker sample includes men ages 25-54, with 30-50 hours contracted per week and less than one year of tenure at a private sector establishment. Work-History sample includes all RAIS observations 1995-2014 for individuals who are ever included in the new worker sample. Standard deviations in brackets.

Figure 3: Histogram of DHS Growth Index for New-Worker Sample

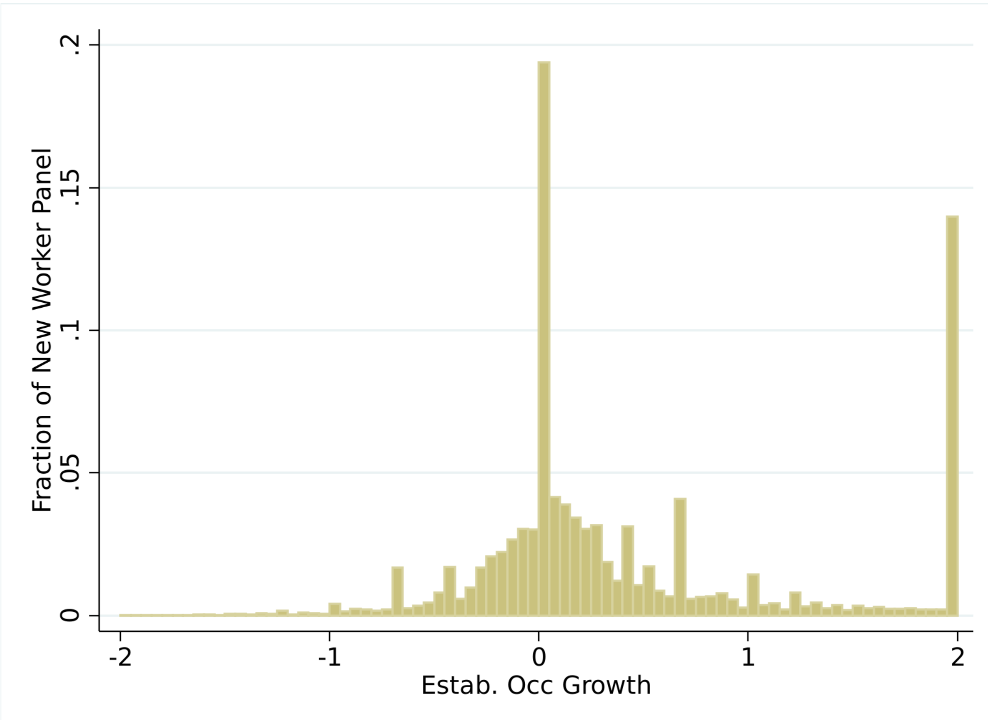

Notes: From RAIS new worker sample, 1995-2014, with restrictions as described in Section 4.1. The DHS index is calculated as $2 \times \frac{L_t - L_{t-1}}{L_t + L_{t-1}}$ for each occupation employed within each establishment.

used for the analysis of hiring and wages (columns 1 and 2), as well as the sample that includes those workers' full work histories (column 3). There are 12.4 million observations in the new worker sample, and they span approximately 400,000 establishments at just over 300,000 unique firms. About one third of those workers are employed by one of roughly 19,000 firms that have multiple establishment in different regions of the country. Several features are of note. Firstly, even though this is the formal sector, levels of education are low by the standards of the developed world; more than half of sampled individuals have less than a high school education, and less than 10% have any education beyond the high school level. Additionally, nearly 60% of the sample is engaged in some form of "blue-collar" occupation, although the considerable majority of these occupations are reported to require some degree of training or skill.[14] Finally, as described above, a slight majority of observations in this sample are from the Belo Horizonte Brazilian micro-region. No other region in Brazil comprises more than 10% of the sample.

Figure 3 shows the distribution of occupation-level employment growth within the new worker sample, using the DHS index as described in Section 3. Unsurprisingly, there are two large spikes in the distribution of this index within the sample. The first occurs at precisely 0, indicating that these new workers are in establishments that have undergone a one-for-one replacement of employees within an occupation. The second spike occurs at an index value of 2, indicating establishments that are employing individuals within an occupation for the first time. There is of course no corresponding spike at -2, simply because there are no worker-level observations within an occupation if an establishment exits employment of that occupation. The occupation-level growth distribution is otherwise relatively continuous and centered near 0, with somewhat greater density at small rates of growth than in corresponding rates of decline.

For all worker-level wage regressions, the dependent variable that I use is log December wage income, reported as a multiple of the Brazilian minimum wage income because the size of the *abono salarial* is determined by the minimum wage. I use the same measure of income as an explanatory variable in job separation regressions. Because all specifications include year fixed effects and I use log wages, results are invariant to the normalization of wage income used in the data.

For regressions of separation behavior, I simply extend the new worker sample to include the entire labor market histories for 1995-2014 of the individuals who ever appear in the original new worker sample. This yields a total of 30.2 million observations over the 20 year period. Several summary statistics in the full work-history sample differ from the new

[14]The division of 343 three-digit CBO occupations into these five broad classifications is from (Menezes-Filho et al., 2008).

worker sample. Reported mean wage incomes are higher, which is at least in part an expected result of returns to job tenure. However, more notably, the work-history sample is weighted toward workers with more education and who are employed in white collar occupations. These distinctions appear to be a function of the informal sector in Brazil, in which it is likely that there are comparatively more opportunities for blue-collar employment.

4.2 PNAD

Like many middle-income countries in Latin America, Brazil has a large informal sector labor market in addition to its formal sector. An important question concerns the extent to which this informal sector influences the competitiveness of formal sector labor markets. However, the RAIS data only include information on the formal sector. So, to analyze this question, I incorporate statistics constructed from an annual survey of Brazilian households, the *Pesquisa Nacional por Amostra de Domicilios* (PNAD). Among numerous other topics, the PNAD survey asks individuals about the details of their employment status that includes whether or not they are in possession of a *carteira de trabalho assinada* for their primary employment. Only workers in formally-registered establishments are eligible to receive this document, and issuance of the document is mandatory for formal-sector workers because the document is used to obtain the benefits associated with the PIS program. Thus, I am able to use these data to construct a measure of the proportion of the labor force that is engaged in formal sector employment.

As discussed previously, there is ample evidence that the labor market is becoming increasingly formalized over time in Brazil. According to PNAD microdata, in 1995, the initial year of this analysis, 29.7% of the labor force reported that they were in possession of a *carteira* from a non-government, non-military entity, while another 22.3% reported that they did not have such a contract, 6.9% reported that they were in government or military employment, and 41.4% were self-employed, employed in production for own consumption, or otherwise had a status that could not be determined. By 2014, the final year of this analysis, fully 41.2% of the labor force reported that they were privately employed and in possession of of a formal sector contract, with the proportion of the labor force explicitly reporting no contract having declined to 19.1%, and the balance of the increase in the proportion of formal sector employment arising from declines in other categories.

There is also clear evidence that informally-employed workers earn considerably lower salaries, on average, than formally employed ones. Figure 4 shows a quantile plot of monthly income in 1995 for individuals who are reported to be employed in the private sector with formal contracts, compared against the quantile plot of income of those who are informally

Figure 4: Quantile Plot of Monthly Incomes for Formal and Informal Sector Employed (1995)

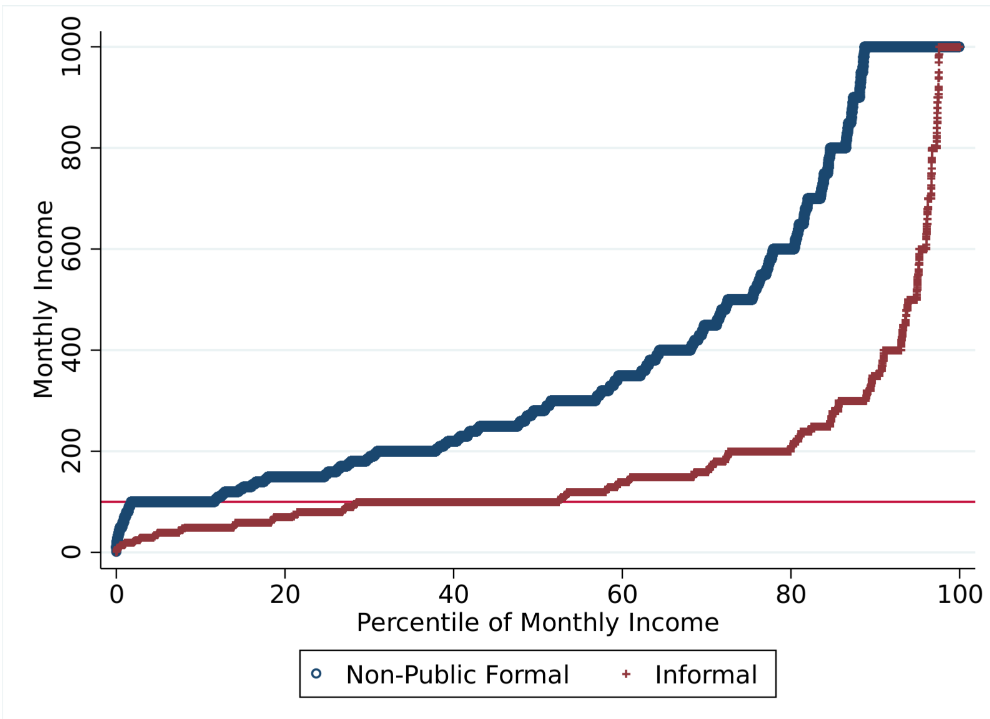

Notes: Income data from PNAD. Non-public formal workers include all individuals who report employment with a *carteira* excluding government employees and military. Informal workers includes workers who report that their employment is without a *carteira*, or who report that their employment is production for own consumption. The red line represents the Brazilian minimum wage for 1995.

employed. From the graph, it is clear that formal sector incomes first-order stochastically dominate informal sector incomes. Indeed, in 1995, more than half of informal sector workers were paid a monthly wage that was at or below the statutory minimum wage for Brazil, while only about 10% of formal sector workers reported incomes at or below the minimum wage. This suggests that the effects of labor market informality on market competitiveness may be expected to be analogous to the effects that a large stock of unemployed workers would have in search models with equilibrium unemployment such as in Burdett and Mortensen (1998).[15]

Further details regarding the construction of my measure of labor market formality are discussed in the Appendix.

5 Results of Both Empirical Specifications

5.1 Wage-Setting Specification

Table 2: Baseline Regression Results: New Worker Wages

	All Establishments		*Multi-Region, Multi-Estab. Firms*			
	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.	(4) IV: Other Est./Occs.	(5) OLS	(6) IV: Other Occs.
Occ. Growth	0.00509***	0.0182***	0.0131***	0.0643***	0.00729***	0.0172***
	(0.000452)	(0.00186)	(0.00486)	(0.0213)	(0.00110)	(0.00252)
Observations	9,844,177	9,293,749	2,458,597	3,195,268	2,458,597	2,316,113
Adjusted R-squared	0.870			0.907		
Estab. FEs	Yes	Yes	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Num. Clusters	251,555	193,598	49,709	60,107	49,709	31,325
K-P F Stat		4,373	532.5	101.6		5,082

Notes: From RAIS new worker sample. Dependent variable is log December wages; Occ. Growth is the occupation-specific DHS index of establishment employment growth from the prior year. Columns 3 through 6 are restricted to workers in multi-establishment, multi-region firms as in Column 3. All specifications include education group controls, education group by year controls, nationality controls, quadratic and cubic age profile terms and tenure controls. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

Table 2 shows the results of my baseline wage-setting specification, as indicated by Equation 5. Column 1 shows the results for the OLS specification, while columns 2 through 4

[15]Menezes-Filho et al. (2008) use a selection model to analyze selection into formal sector status in Brazil using RAIS data. They find suggestive evidence that a larger informal sector is associated with a larger formal sector wage premium for otherwise comparable workers.

show results using the specified IV strategies of using growth in other occupations within the establishment and growth in employment at other non-local establishments within the same firm. Because the restriction necessary to produce the IV specification in columns 3 and 4 exclude both single-establishment firms and firms with no non-local establishments, the sample sizes reported for these column are considerably smaller than the sample sizes reported in columns 1 and 2. Accordingly, in columns 5 and 6 I reproduce the OLS and own-establishment IV results, restricted to the same of establishments used in column 3. In all wage regression specification, time varying worker covariates include four education group indicators, a full set of education group by year indicators, polynomial terms in $(age - 40)$, linear and quadratic tenure effects, and a full set of nationality controls. Other than the addition of nationality and tenure controls, this specification is largely identical to that used in the recent literature using overlapping establishment and worker fixed effects (Card et al., 2016).[16] Finally, note that in all specifications, standard errors are clustered at the establishment level to ensure that inferences are robust to within-establishment correlation in the unexplained component of wages.

As indicated by the table, the baseline regression specification implies very small non-zero inverse elasticity estimates. These are consistent with a labor market that is not perfectly competitive, but is not strongly monopsonistic either. The OLS estimates in columns 1 and 4 suggest that a firm seeking to double its employment level in a particular occupation may be expected to offer a wage premium to new workers of about 0.5% relative to what they would offer in a period of no employment growth. This implies a labor supply elasticity to the firm of approximately 200. As expected, the three IV strategies each provide somewhat larger inverse elasticity estimates. However, these estimates still imply a labor supply elasticity faced by the firm at the time of hiring of between 15 and 75, with my preferred specification in column 4 implying the most inelastic firm-level labor supply. These estimates are not infinite. Nonetheless, they are much larger than the range of estimates provided from most studies of job separation activity, and they are suggestive of a labor market that is reasonably competitive.

One concern in looking at the baseline specification may be that the wage premium could be highly nonlinear. In particular, the results could be driven by establishments that are new entrants into a particular occupational market, or solely by establishments that are

[16]In particular, the use of quadratic and cubic terms in $(age - 40)$, rather than the more traditional use of linear terms in age and experience, has been used to address a well-known problem in identifying worker effects (and in particular, their presumed cohort component) when linear age effects are included in a regression along with year fixed effects. Instead, in this specification, the identifying assumption is that the age-wage profile is flat at age 40. Figure 7 in the Appendix shows the unadjusted age-wage profile observed in the RAIS data, which supports this assumption for my sample.

Figure 5: Non-parametric OLS Estimation of New Worker Wage Premia

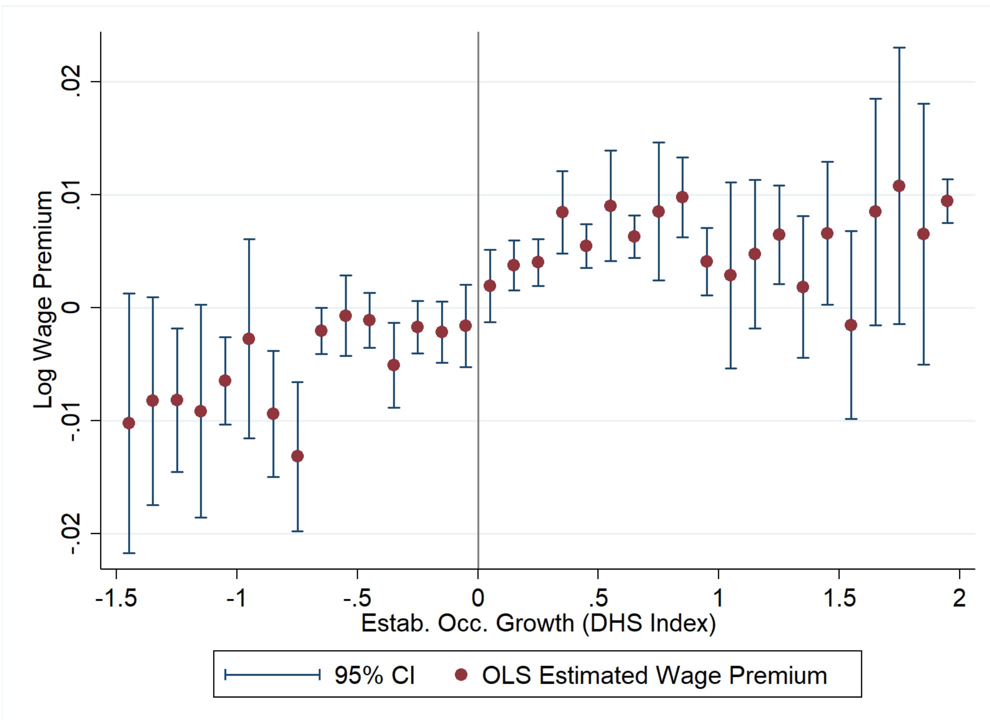

Notes: From the RAIS new worker sample, 1995-2014. Coefficients and confidence intervals are from an OLS regression with 40 indicator variable indicating employment growth in each bin of DHS employment growth index of width 0.1. The omitted category is firms that reported exactly zero growth in occupational employment. All other covariates are as specified in Column 1 of Table 2. Standard errors for confidence interval construction are clustered by establishment.

experiencing rapid growth or decline. Although the DHS index is a bounded and symmetric index that may in part address these outlier type concerns, it may still be beneficial to examine this wage behavior in a non-parametric way. Because of the size of this dataset, even with my region-level sample I am able to construct non-parametric OLS estimates of the establishment log wage premia offered to workers at various binned levels of employment growth. These estimates are presented graphically as Figure 5. Each point in this figure represents the point estimate for a bin of width 0.1 in the DHS measure (approximately 10% growth or decline for values close to 0), plotted relative to the point of zero growth. As the figure demonstrates, the pattern of estimated wage premia is relatively linear; if anything, the gradient of wage premia is steepest for small levels of growth or decline. Non-parametric estimates become more imprecise with greater growth and decline because comparatively few establishments exhibit such rapid changes in employment, even at the individual occupation level. However, the largest bin, which is comprised primarily of new entrants to a particular occupation's labor market, is comparatively precisely estimated, and the results suggest a wage premium that is consistent with the broader pattern. Indeed, for no levels of growth or decline is the wage premium predicted to vary by more than 2% from what would be predicted under zero employment growth.

Overall, these results are fairly consistent with the exceedingly small literature that has credibly estimated the labor supply elasticity faced by firms by looking directly at wage setting, which has estimated the elasticity of the labor supply curve faced by firms at greater than 10. Later, in Section 6, I show these estimates to be fairly robust to a range of other potential concerns.

5.2 Job Separations Specification

As discussed in the introduction to this paper, most contemporary literature on firms' monopsony power has not chosen to estimate the elasticity of firms' labor supply curves directly, primarily because of prior concerns about simultaneity. Instead, most papers have built on the aforementioned dynamic model specification of Manning to estimate the elasticity of workers' separations with respect to their own wage, and they have then used these estimates to infer the labor supply elasticity faced by firms under the assumptions of the original model. In this subsection, I conduct my own analysis of this type, using the regression specification described in Section 3.2, which incorporates local labor market fixed effects into a linear probability model that is otherwise similar to what has been previously estimated in the literature. The education indicators, education by year indicators, and nationality indicators are as specified in the previous section. In place of quadratic and cubic terms in $(age - 40)$, separations specifications include quadratic and cubic terms in $(age - 18)$, because 18 is the age in the sample at which the mean voluntary separation rate is highest. In keeping with the typical empirical specifications used in this literature, I omit own job tenure from the controls used in the baseline separations specification.

Table 3: Baseline Results of Separations Specification

	(1) OLS	(2) IV: Initial Wage	(3) OLS w/Additional FEs
Log December Wage	-0.0184**	-0.00701**	-0.0159***
	(0.000821)	(0.00322)	(0.000654)
Pr(Separation)	0.0507	0.0199	0.0509
Implied Separation Elasticity	-0.362	-0.352	-0.312
Adjusted R-squared	0.198		0.380
Observations	26,312,112	7,638,089	22,380,111
Num. Clusters	389,628	122,167	243,027
Estab. FEs	Yes	Yes	No
Individual FEs	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	No
Occ-Micro-Estab-Year FEs	No	No	Yes
K-P F Stat		1,402	

Notes: The outcome in all regressions is a binary indicator of voluntary job separation, and Log December Wage is the worker's own reported wage. All specifications include education group controls, education group by year controls, and nationality controls. Column 2 includes only workers with greater than one year of tenure reported. Standard errors in parentheses are clustered by establishment. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3 shows the baseline results of this specification. Below each specification, I report both the mean voluntary separation rate in the sample, and the implied separation elasticity estimate of that regression. Column 1 shows the results of OLS estimation, while column 2 shows the results from using initial wages as an instrument for current-period wages, and column 3 shows the results using OLS, but including establishment-occupation-year fixed effects in place of separate establishment and occupation-region-year fixed effects. Overall, these results suggest that workers are indeed highly unresponsive to their own wages in their decisions of whether or not to separate from their current employment, with estimated separation elasticities of -0.3 to -0.4. Under the steady state assumption and the model of firm wage setting described in Section 2, these would be presumed to correspond to a labor supply elasticity to the firm of 0.6 to 0.8, exceedingly far from the estimated elasticities of 15 to 76 estimated by IV regressions that apply the direct wage setting specification.

As with the wage regressions on the new hire sample, the results from Table 3 are reasonably consistent with prior estimates in the literature, in spite of the methodological improvements made here. If anything, these estimates suggest that Brazilian workers may be somewhat less responsive to wages in their separation decisions than workers in the other contexts in which this question has been studied, which tend to find estimates of the labor supply elasticity from 1 to 5. There are many potential explanations that one might imagine for why this might be the case in this particular setting, including numerous explanations regarding the presence of additional labor market frictions in a developing country context. However, explanations based on the standard dynamic model presented in Section 2 cannot explain the large difference between estimates produced using these two methods on the same sample of workers, even with relatively large deviations from the steady state assumptions that are typically imposed.

6 Robustness Checks

While the results presented in the prior sections of this paper are strongly suggestive, there may still be particular questions arising from the specifications used. In this section of the paper I discuss and show results for several robustness check strategies that may address specific concerns along these lines. In the Appendix, I show several additional robustness check results that may be of interest as well.

6.1 Testing for Endogenous Substitution

The results for the IV specifications shown in Table 2 identify the elasticity of the labor supply of each occupation based on two key assumptions:

1. In the presence of firm-level product demand or productivity shocks, there exists a scale effect such that *a priori*, the firm desires to increase in the short run its use of other labor inputs.

2. In the presence of firm-level product demand or productivity shocks, firms do not endogenously change the extent of their substitution of one labor input for the other labor inputs that are being used as instruments as a result of any factor that is not fully captured by the fixed effects included in the empirical specification.

The first assumption implies the relevance of the IV strategies described here, and unsurprisingly the input instruments used in this paper are all quite strong. However, the second assumption, that patterns of substitution are not endogenous to time-varying firm-level shocks, could be plausibly violated in some circumstances. For example, as discussed in Section 3.1, if local labor market fixed effects are specified too finely relative to actual local labor markets, and if occupational inputs within the same labor market are then included in the growth measure that is used as an instrument, then simultaneity bias could arise from a correlation between the labor demand of one's own firm and the labor demand of other firms within the same local labor market. In IV strategies 2 and 3, in which I instrument using only non-local inputs of the same firm, the scope for this type of correlation is greatly reduced, but if local labor market conditions are strongly correlated with non-local conditions, there may still be scope for some form of endogenous substitution bias.

A common feature of all these stories is that they require firms to engage in substitution. Yet, many labor inputs are very poor substitutes for one another. Consider, for example, a hospital that employs doctors, nurse practitioners, and janitors. One might readily imagine a situation in which, in response to an establishment-specific supply shock, the firm engages in more or less substitution of nurse practitioners for doctors at different points in time. However, janitors are almost surely not a substitute for either doctors or nurse practitioners in the production of medical care. So, while an estimate that uses employment growth of nurse practitioners and janitors as an instrument for employment growth of doctors might be subject to endogenous substitution bias, an estimate that uses only employment growth of janitors will not be subject to the same bias.

It is, of course, not feasible to model the individual production functions of each firm, and I will not seek to do so. However, I have constructed the following straightforward robustness check, based on the simple assumption that short term substitutability across occupations is likely to be strongly correlated with labor market transition behavior over time.

Let $m_{oi} = 1$ if worker i ever reported employment in occupation o over the full period observable in the RAIS data, 1986-2014, including the full geographic sample. Then, for each pair of occupations o and o', one may construct the following:

$$M_{o,o'} = \frac{\sum_i \mathbf{1}(m_{io} \cap m_{io'})}{\min\{N_o, N_{o'}\}} \tag{7}$$

This formula, which I will refer to as the "ever-transition" probability, is simply the larger of the two conditional probabilities that, given that a worker is ever observed in occupation o, they are ever observed in occupation o', and vice versa. Notably, these ever-transition probabilities are computed without any assumption that observed employment in occupations o and o' occurs in adjacent years, that it occurs in the same establishment, firm, or location, or that it occurs in any particular order. Using the larger of the two conditional probabilities addresses situations in which one of the two occupations under consideration employs many more workers than the other. With 343 occupations listed in the RAIS dataset, this yields 58,653 occupational pairs whose ever-transition rates can be calculated and ranked.

Then, let $\tilde{o}$ be the set of occupations such that $M_{o,o'} < \bar{M}$. So, one may calculate $G_{\tilde{o},t,,j(i,t)}$ as the growth rate for the set of occupations such that $M_{o,o'} < \bar{M}$ for all $\tilde{o}$, i.e. the growth rate for the set of pairwise low-transition occupations only. If substitution bias is leading the IV inverse elasticity estimates shown in Table 2 to be biased downward, then the use of low-transition occupations as an instrument will recover estimated inverse elasticities that are larger. That is, they will find the labor supply curve faced by the firm to be more inelastic.

Figure 6: Ranked Scatterplot of Occupational Pairs by Ever-Transition Probability

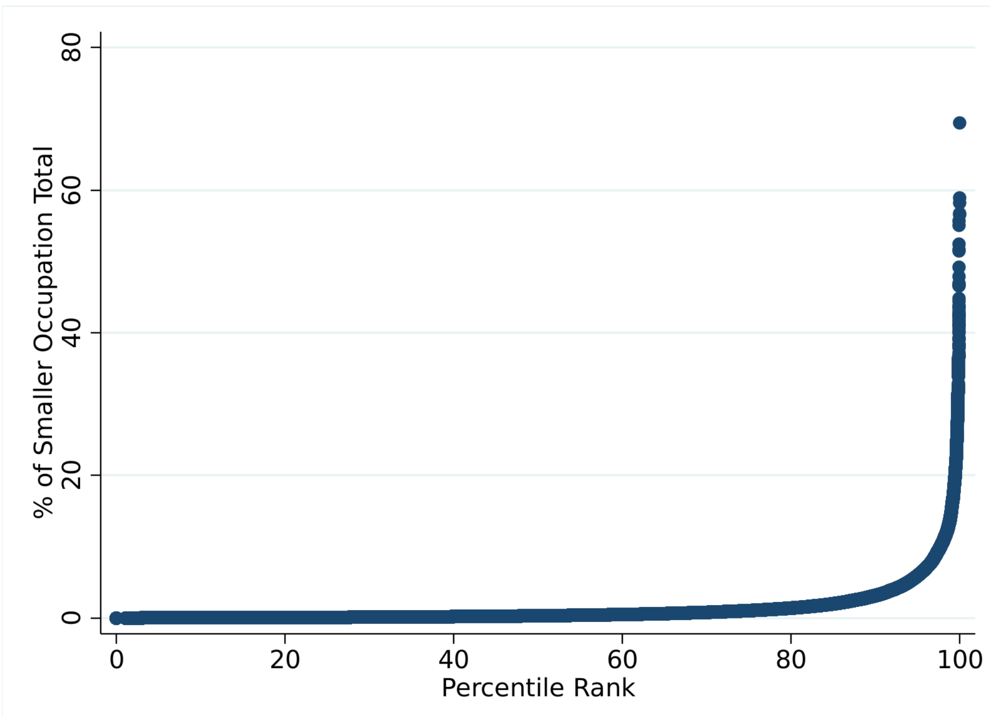

Notes: From RAIS full dataset, 1986-2014. Each pairwise ever-transition probability is calculated as the probability of an individual ever appearing in RAIS as employed in one occupation, conditional on them ever reporting employment in the other occupation. The larger of the two conditional probabilities for each pair is reported, and same-occupation pairs are excluded. Median pairwise transition probability is 0.3%.

Figure 6 shows an ordered scatterplot of ever-transition rates for all 58,653 occupational pairs, constructed using the entire RAIS dataset from 1986-2014. It is perhaps unsurprising to see that a relatively small fraction of occupational pairs exhibit high ever-transition rates. In contrast, most occupational pairs have very low ever-transition rates, and there are many examples in the data of occupational pairs in which no individuals ever transitioned between the two occupations in nearly 30 years of formal sector observation. The median occupational pair exhibits a transition rate of approximately 0.003, implying that over this 29 year period, only three in one thousand individuals who report employment in the smaller of the two occupations in the median pair were ever employed in the other occupation. Investigation of individual pairs, unsurprisingly, also shows that seemingly closely related occupations tend to rank more highly than seemingly unrelated occupations. For example, economists are relatively likely to have ever been accountants; this pair is in the 99th percentile, with an ever-transition probability of approximately 14%. However, economists have only a median pairwise ever-transition probability of also being mining supervisors (0.3%), and almost no individuals who are ever economists also report ever being employed as agricultural machine operators.

Table 4: Results of IV Regressions Using Only Growth in Low-Transition Occupations

	Within-Estab. Growth		*Non-local Firm Growth*	
	(1) Below Median Switchers	(2) Below 25th Pctile. Switchers	(3) Below Median Switchers	(4) Below 25th Pctile. Switchers
Occ. Growth	0.0201***	0.00880	-.006534	.008618
	(0.00475)	(0.00900)	(0.01644)	(0.01432)
Observations	2,221,523	957,926	1,132,760	585,612
Estab. FEs	Yes	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes	Yes
Num. Clusters	27,759	8,212	26,590	12,832
Kleinbergen-Paap F Stat	857.2	354.3	202.5	106.9

Notes: Each column reports the result of an IV specification in which employment growth in the worker's own occupation is instrumented for using growth only in low ever-transition occupations within the worker's same establishment. All other specification details are as in columns 2 and 4 of Table 2. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

Table 4 shows the results of a regression using an IV specification in which establishment-level growth in occupation o is instrumented for by growth in occupation $\tilde{o}_o$, either within establishment or in non-local establishments of the same firm. The instrument in column 1 is establishment-level growth in below-median ever-transition occupations in the same establishment, while the instrument in column 2 is establishment-level growth in occupations whose ever-transition probabilities are below the 25th percentile in the same establishment. In column 3, the instrument is below-median ever-transition occupational growth in non-local establishments of the firm, while in column 4 this instrument is again restricted to growth in below 25th percentile occupations within non-local establishments of the same firm. All other details of these specifications are as in Table 2, columns 2 and 4. Notice that in all columns of this table, the sample size is reduced considerably. This is because many establishments and firms simply do not employ individuals in occupations that are sufficiently transition-distant from one another, and so the individuals in these establishments are dropped. The inverse elasticity estimate shown in the first column of this table is larger from both of the coefficients from the analogous specifications in Table 2 (columns 2 and 5), but the difference is quite small and is not statistically distinguishable from those baseline estimates. Attempting to limit the set of low-transition occupations further, as I do in column 2, provides results with little power owing to the greatly reduced sample size. Similarly, columns 3 and 4 provide no clear evidence of endogenous substitution bias. In all, the results shown in this subsection provide no evidence to suggest that endogenous substitution bias is a major concern in the baseline estimates.

6.2 Additional FEs

As discussed in Section 3, the fixed effects specification used for my baseline regression results provides an estimate of the labor supply elasticity faced by firms that is identified from two types of variation: simultaneous variation across occupations within each establishment, and variation over time within occupation-establishment groups. However, by including either a full set of establishment-occupation fixed effects or a full set of establishment-year fixed effects, it is possible to isolate each of these sources of variation. Table 5 shows the results of each of these regression specifications.

The results of these regressions suggest that the baseline elasticity estimates are driven primarily by variation in wages over time, rather than within-year variation in the growth rate of different occupations. More specifically, the inclusion of establishment-occupation fixed effects in Panel A provides estimates that are extremely similar to the baseline estimates in both OLS and IV specifications. In contrast, the inclusion of establishment-year fixed effects leads to point estimates that are both very small and statistically insignificant. This suggests that most of the variation in wages, at least once local labor market conditions are controlled for, occurs at the establishment-year level rather than at the establishment-occupation level. Broad wage-setting policies, such as policies based on rent sharing, could also be consistent with these findings.

Another, somewhat different concern is that the inclusion of fixed effects at the occupation $\times$ region $\times$ year level could be insufficient to address local labor market heterogeneity if the

Table 5: Baseline Regressions with More Restrictive Fixed Effects Specifications

	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.
		Panel A: Estab-Occ. FEs	
Occ. Growth	0.00470***	0.0152***	0.0158***
	(0.000516)	(0.00169)	(0.00565)
Observations	9,552,340	9,014,280	2,414,287
Adjusted R-squared	0.890		
		Panel B: Estab-Year FEs	
Occ. Growth	0.00113***		-0.00175
	(0.000392)		(0.00339)
Observations	9,235,334		2,337,681
Adjusted R-squared	0.888		
		Panel C: Occ.-Ind.-Micro-Year FEs	
Occ. Growth	0.00511***	0.0166***	0.0123**
	(0.000464)	(0.00172)	(0.00559)
Observations	9,693,900	9,147,854	2,421,238
Adjusted R-squared	0.875		
Individual FEs	Yes	Yes	Yes

Notes: All results are from regressions on the RAIS new worker sample as in columns 1 through 3 of Table 2, but with different fixed effects specifications. Panel A replaces establishment FEs with establishment × occupation ones. Panel B replaces establishment FEs with establishment × year ones. Panel C replaces occupation × region × year fixed effects with occupation × industry × region × year ones. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

local labor market conditions faced by a firm who seeks to hire new workers are driven by time-variant industry-level considerations in addition to occupation-level considerations. If, for example, it is less costly for firms to attract workers from their own industry than from other industries, then the wage that firms choose to offer could depend on the local industry wage premium as well as the local occupational wage premium. In such a case, failure to control for local industry conditions could allow for the continued presence of simultaneity.

In Panel C of Table 5, I present results in which I replace the occupation $\times$ region $\times$ year fixed effects with occupation $\times$ industry $\times$ region $\times$ year fixed effects. This specification allows for industry-level conditions to impact firm behavior, and also flexibly permits interaction between occupation and industry-level labor market conditions in wage setting. The results are little different from the baseline regression. This suggests that industry-level conditions may not be a large simultaneity concern as long as occupation-level conditions are accounted for.

6.3 Using aggregate occupational growth

The approach of this paper is largely based on the principle that labor market activities occur primarily at the occupation level. There is ample support for this assertion in the recent literature on labor market transition behaviors (Gathmann and Schönberg, 2010; Guvenen et al., 2015; Macaluso, 2017), and it aligns with the intuitive notion that many occupations are not ready substitutes for one another in light of their vastly different skillsets. When an establishment seeks to hire new workers with particular skills, it faces the labor market for the occupation that possesses those skills. Accordingly, the baseline new hire wage regression results shown in Table 2 use occupation-level employment growth as the relevant measure of employment growth, and it uses occupation $\times$ region $\times$ year local labor market fixed effects as the measure of local labor market conditions.

However, the limited body of work that has looked at monopsony in wage setting to date has not, in general, taken an occupation-based approach consistently. In studies where a particular occupation's employment is subject to an exogenous shock (e.g. Matsudaira, 2014), that shock is typically used as an instrument for an occupation-level measure of changes in employment. In contrast, strategies that use firm- or market-level instruments (e.g. Schmieder, 2013; Bellon, 2016) have typically instrumented for establishment-level changes in employment. Therefore, it may be particularly useful to understand the extent to which these approaches may be expected to differ.

In Table 6 I present results in which workers' wages are regressed on an establishment-wide measure of employment growth. Columns 1 and 3 show OLS regression results without

Table 6: Regressions: Aggregate-Level Employment Changes

VARIABLES	(1) OLS	(2) IV: Other Estabs.	(3) OLS	(4) IV: Other Estabs.
Estab. Growth	0.0103***	0.0525***	0.0106***	0.0448***
	(0.00107)	(0.0198)	(0.000855)	(0.0146)
Observations	9,875,353	3,257,487	9,844,177	3,234,270
Adjusted R-squared	0.858		0.870	
Estab. FEs	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	No	No	Yes	Yes
Num. Clusters	252,768	62,165	251,555	61,413
Kleinbergen-Paap F Stat		147		156.8

Notes: Columns 1 and 3 provide OLS estimates of worker wages on the DHS index of total employment growth in the worker's own establishment. Columns 2 and 4 provide IV estimates on workers in multi-region, multi-establishment firms, using the DHS index of total employment growth in non-local establishments as an instrument for total employment growth in the worker's own establishment. All other specification details are as in Table 2. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

and with the inclusion of occupation-level local labor market FEs, while columns 2 and 4 apply the IV strategy of instrumenting for establishment employment growth using non-local firm establishment growth of multi-establishment, multi-region firms, analogously to columns 3 and 4 of Table 2. The inverse elasticity estimates obtained using this approach are within the range of estimates produced by the occupation-level approach, although they are at the lower end of the range, as is result from my preferred IV strategy of using non-local growth in other occupations. Notably, the range of inverse elasticity estimates shown in this table (0.0448 - 0.0525) is also quite close to the estimated inverse elasticity that Schmieder (2013) obtains when using establishment age as an instrument for employment growth in Germany (0.046). These IV estimates would correspond to an overall labor supply elasticity to the firm in the range of 19 to 23, instead of the 15 to 76 implied by the baseline IV regression results.

Overall, this table suggests that the key result of the paper—that differences between separation elasticity estimates and hiring wage elasticity estimates are not driven by sample selection or simultaneity—is not greatly impacted by whether employment growth is measured at the occupation level or at the aggregate level. I argue that as a more direct reflection of labor market behavior, a more occupation-level approach has much to recommend it. However, occupational data may not be available in all datasets, and so it is helpful to know that these results are not primarily driven by the strategy of using occupations as different labor inputs.

7 Analysis of Heterogeneous Effects

Although the baseline elasticity estimates are themselves interesting and suggestive, there may be particular value in understanding the ways in which these elasticity estimates vary over well-known demographic groups. The large observed difference between new hire and separation elasticities underscores this concern; if firms' face such large differences in labor supply elasticity for new and existing workers, then it is quite reasonable to expect that these same elasticity estimates will vary in predictable ways depending on the characteristics of the workers or markets analyzed.

The limited evidence on separation elasticities has suggested that they vary based on gender, and also over the business cycle.[17] Both of the above results can readily be explained as a function of differences in labor market search frictions. However, as evidence for the importance of search frictions, these are a fairly indirect test. A more direct test would consider whether individuals who are observably identical are more or less sensitive to their own wage depending on a more direct measure of the the probability of finding equivalent employment in their local labor market, outside of their own firm. Because of the comprehensiveness of the Brazilian RAIS data, I am able to construct intuitive measures of this probability, and therefore I am able construct simple tests of the hypothesis that labor market frictions are an important determinant of both separation behavior and of firms' wage setting.

The estimates shown in this section are all OLS estimates, and all specifications include the same combination of worker, establishment, and local labor market fixed effects as are been used elsewhere in this paper. While the choice reflects the particular challenge of credible IV estimation with interaction terms, it is not clear that the use of OLS should have a particular impact on the significance of these estimates, because the particular sources of simultaneity that IV specifications can address do not have an obvious correlation with these market-level measures.

The first measure that I construct is a simple measure of local occupation-specific market size, excluding one's own firm. The specific measure that I adopt here is:

$$\log(N_{m,-f,o,t}) = \log(N_{m,o,t} - N_{m,f,o,t}) \tag{8}$$

[17] To my knowledge, no research has tested for the existence of heterogeneous effects in wage setting behavior.

The estimated interaction term coefficient in this regression specification can be roughly interpreted as the change in the predicted elasticity (or inverse elasticity) associated with a doubling of the number of individuals employed in one's own occupation in the local labor market outside of one's firm. Under the assumption that potential employment opportunities are closely tied to existing patterns of employment, this statistic captures the measure of local labor market opportunities that would be predicted to be most important in a directed search model of the labor market where the rate of job finding depends on the number of sufficiently suitable jobs. [18] In such a model, firms face a more elastic labor supply curve when the ex-firm local labor market is larger, because employees have more opportunities for ex-firm matches and therefore face lower search costs.

The second measure that I construct is a measure of the relative prevalence of an occupation in the ex-firm local labor market. The local relative prevalence ratio for an occupation is simply the proportion of employment outside the firm that is engaged in the occupation, relative to the proportion of nationwide employment that is engaged in that occupation. That is:

$$Prev_{mo,-ft} = \frac{N_{m,-fot}}{N_{m,-ft}} \bigg/ \frac{N_{ot}}{N_t} \tag{9}$$

In a labor market search model with undirected job search (such as a model with exogenous arrival of random job offers), this measure captures the measure of local labor market opportunities that would be predicted to be most important, again under the assumption that potential employment opportunities are closely tied to existing patterns of employment. In such a model, firms face a more elastic labor supply curve when the ex-firm local labor market has a high local labor market prevalence of the occupation in question, because each arrival of a job offer to the worker is more likely to be of the same occupation and therefore exceed the worker's reservation threshold for accepting an offer.

The third measure that I construct is the proportion of employed men in each occupation and state that are employed in the formal sector, defined in PNAD as being in possession of a *carteira de trabalho assinada*. As described in the Appendix, formal sector status can be inferred for most but not all respondents to PNAD. For example, individuals who report that they are self-employed or employers themselves cannot be determined to be formally or informally employed. Such individuals are not counted as employed in the formal sector for the measure constructed here. Additionally, because changes over time in the degree of formality may pick up the effect of changes in local labor market conditions,

[18] An even better measure of local labor market conditions would use job vacancy data. However, I am unaware of any data on job vacancies in Brazil that would be available at a level of geography or occupational specificity such that they would be usable here.

for this specification I use only data on the degree of labor market formalization prior to the beginning of my analysis, in the period 1992-1995.

Unlike the first two measures that I construct, the theoretically predicted relationship between labor market formality and the labor supply elasticity is not immediately evident. However, in light of the large differences in wages offered to formal and informal sector workers (see Figure 4), it may be most appropriate to think of informal sector workers as in essence underemployed. Given this, one compelling hypothesis may be that, since firms can readily draw upon the pool of underemployed informal sector workers, that a larger informal sector makes the supply curve faced by the firm more elastic. But, even with this simplification, the Burdett and Mortensen (1998) model that has influenced most contemporary studies of job separation behavior does not yield such immediate predictions regarding the relationship between equilibrium unemployment/underemployment and the labor supply elasticity without further assumptions regarding the frictional parameters that characterize the model. Broadly speaking, the model predicts that labor markets with large informal sectors will have more elastic labor supply when the arrival rate of formal sector job offers is relatively low in those markets.[19]

Table 7 shows the results of these heterogeneous effects regressions on the new worker wage specification, while Table 8 shows the same interactions effects as applied to the separations specification. The two tables provide show markedly different results. Specifically, there is no evidence here that firms offer larger or smaller wage premia in relation to growth in response to the local labor market conditions represented by these measures. All interaction coefficients are very small and statistically insignificant. In contrast, there is substantial evidence that these workers' degree of sensitivity to local labor market conditions is related to the conditions in their local labor market.

When workers have more local opportunities for alternative employment in their current occupation outside of their own firm, they are more willing to voluntarily separate in response to low pay. This observation holds regardless of whether one uses the measure that would

[19] More specifically, in the simplest version of the Burdett and Mortensen model, the degree of wage dispersion depends on two key frictions, the rate of arrival of job offers to the unemployed, and the rate of arrival of job offers to the employed. Employment to unemployment transitions only occur because current job matches are destroyed at an exogenous rate. While the exact extent of wage dispersion (and therefore the supply elasticity) depends on both frictions, the unemployment rate/underemployment rate depends primarily on the arrival rate of job offers to the unemployed alone. It can readily be shown that both the unemployment/underemployment rate and the extent of equilibrium wage dispersion are decreasing functions of the offer arrival rate for unemployed workers, but the extent of equilibrium wage dispersion is an increasing function of the offer arrival rate for employed workers. So, an equilibrium in which there is a large informal sector and also lower wage dispersion is characterized by a low arrival rate of formal sector job offers for informal sector workers, but also a low arrival rate of new job offers for workers who are currently employed in the formal sector.

Table 7: Heterogeneous Effects in New Worker Wage Regressions

	(1) Ex-Firm Occ. Size	(2) Ex-Firm Occ. Prevalence	(3) Local Occ. Formality
Occ. Growth	0.00751***	0.00506***	0.00548***
	(0.00166)	(0.000457)	(0.000708)
Growth × Log Ex-Firm	-7.85e-05		
	(0.000225)		
Growth × Prevalence		-0.000291	
		(0.000390)	
Growth × Formality			-0.00120
			(0.00107)
Local Log Occ. Ex-Firm	-0.00147**		
	(0.000729)		
Local Occ. Prevalence (Std.)		0.0147*	
		(0.00782)	
Observations	8,203,784	9,844,177	9,152,119
Adjusted R-squared	0.874	0.870	0.870
Estab. FEs	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes
Num. Clusters	225,262	251,555	246,152

All regressions use the same OLS specification as in column 1 of Table 2, but with the additional interaction terms as shown. Ex-Firm Occ. Size is the log of prior period employment in non-firm local establishments in the same occupation. Ex-Firm Occ. Prevalence is a standardized ratio of the proportion of ex-firm employment that is in the same occupation to the ratio of national employment in that occupation. Local Occ. Formality is the percentage of male employment in the same state and occupation that reported a formal sector contract in 1992-1995, from PNAD. The own term for formality is omitted because it is absorbed into local labor market fixed effects. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

Table 8: Heterogeneous Effects in Existing-Worker Separation Regressions

	(1) Ex-Firm Occ. Size	(2) Ex-Firm Occ. Prevalence	(3) Local Occ. Formality
Log December Earnings	-0.0103***	-0.0175***	-0.0198***
	(0.00154)	(0.000796)	(0.00113)
Earnings × Ex-Firm Occ.	-0.000939***		
	(0.000174)		
Earnings × Prevalence		-0.00146***	
		(0.000374)	
Earnings × Formality			0.00384***
			(0.00136)
Log N Ex-Firm Occ.	0.00211***		
	(0.000528)		
Local Occ. Prevalence (Std.)		0.00222	
		(0.00310)	
Observations	24,570,881	26,312,112	24,559,117
Adjusted R-squared	0.198	0.198	0.200
Estab. FEs	Yes	Yes	Yes
Individual Controls	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes
Num. Clusters	381,869	389,628	381,953
P(Separation)	0.0511	0.0507	0.0504

All regressions use the same OLS specification as in column 1 of Table 3, but with the additional interaction terms as shown. Ex-Firm Occ. Size is the log of prior period employment in non-firm local establishments in the same occupation. Ex-Firm Occ. Prevalence is a standardized ratio of the proportion of ex-firm employment that is in the same occupation to the ratio of national employment in that occupation. Local Occ. Formality is the percentage of male employment in the same state and occupation that reported a formal sector contract in 1992-1995, from PNAD. The own term for formality is omitted because it is absorbed into local labor market fixed effects. Standard errors in parentheses are clustered by establishment. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

be considered most relevant in a directed model of search (the number of ex-firm employed in the occupation) or if one uses the measure that would be considered most relevant in an undirected model of job search (local ex-firm occupational prevalence). The magnitude of each of these heterogeneous effects is quite small, but they are consistent with the notion that search frictions play a role in these labor market decisions.

Additionally, higher degrees of labor market formality are associated with more wage-inelastic separation behavior. This finding is consistent with the idea that firms can more quickly attract workers in markets with a large informal sector. It is also consistent with a Burdett and Mortensen equilibrium in which the formal sector job offer arrival rate for informal sector workers is relatively low, and the job offer arrival rate for workers currently employed in the formal sector is also relatively low. However, there may be other plausible channels for this finding as well. For example, formal sector experience could be itself considered a valuable worker trait or a signal of worker quality, leading formal sector employees in heavily informalized labor markets to be more sensitive to their own level of pay relative to the labor market. It could also be the case that formalization is associated with greater dispersion in firm-specific amenities that might lead individual workers to want to remain with their employer even if they are comparatively low-paid.

Overall, these findings are quite supportive of the suggestion that separation decisions are a function of workers' search costs. These results also give credence to the notion that the differences in separation elasticity estimates among demographic groups shown elsewhere in the literature (such as between men and women) may also be a function of differences in search costs. However, while these heterogeneous effects results should be considered to be suggestive, they are not causal estimates. Although the fixed effects specifications used in this paper can address many sources of variation in the level of wages or the rate of voluntary separation, they do not rule out alternative explanations for the variation in those estimates across groups. More restrictive controls, such as including additional dimensions of controls for heterogeneous trends across occupation or region, could potentially rule out some of these alternative explanations. Further research is surely necessary, but the computational demands of high-dimensional fixed effects methods such as this one do place limits on the speed of such progress.

An additional relevant concern in the interpretation of these heterogeneity result may be that the patterns of overlapping fixed effects used in these specifications are not entirely analogous. Individual effects, establishment effects, and local labor market effects can explain a much higher proportion of variation in wages than they can explain variation in voluntary separations. In considering the heterogeneous effects results on wage setting shown in Table 7, it may simply be an issue that there is comparatively little variation in wages left to exploit.

8 Discussion of Alternative Models

As described in the introduction to this paper, prior results in the monopsony literature have suggested somewhat dramatically different labor supply elasticities depending on the method used to estimate it. Specifically, while studies that examine separation behavior have typically suggested that labor markets are highly monopsonistic, the few studies that have looked at wage setting on new hires have found, at best, only very modest evidence of firms' labor market power.

The results shown in this paper do not overturn these basic findings. Rather, by taking advantage of the particular depth and comprehensiveness of data in the Brazilian setting, the results shown in this paper rule out several of the various alternative explanations that have been given for why elasticity estimates might vary depending on the methodology used. The inclusion of local labor market fixed effects, along with the new instrumental variables strategies employed by this paper, help to rule out concerns about simultaneity in the labor market when looking at wage setting behavior of establishments. Additionally, by constructing elasticity estimates using both new hire wages and voluntary separation behavior on an identical sample of workers, this paper strongly rules out the notion that external validity or selection concerns have driven prior results.

The heterogeneous effects results shown in this paper also provide new insights into the circumstances in which establishments have the greatest monopsony power by leveraging the comprehensiveness of the RAIS dataset. Put simply, workers' separation decisions are more sensitive to their own wages when they have ample outside opportunities, as measured both by the number of jobs in their occupation at other local firms, and as measured by the relative prevalence of their own occupation in their local market. Perhaps most interestingly, results also suggest that in a developing economy setting such as Brazil, the presence of informal labor market opportunities may influence workers' sensitivity to their own wages. Yet, the evidence does not suggest that these local labor market conditions influence wage setting behavior on new hires in the same way.

Taken collectively, these results suggest that the large observed differences between wage setting decisions by establishments and separation decisions by workers are in fact indicative of large differences in behavior. Workers, once employed, are quite insensitive to their own wages in their decision of whether or not to voluntarily separate from a firm. However, at best, firms choose to increase wages only modestly in order to attract new workers.

Both the canonical static model of monopsony and the richer dynamic model of monopsony used in recent literature imply that each firm faces a single labor supply curve against a homogeneous labor input that is used in production. While the parsimony of these models is a strength, they may be inadequate to describe the nature of firms' monopsony power, which varies not only on aggregate local market conditions, but also on whether these firms are recruiting new workers or compensating existing workers, and on the characteristics of the workers themselves. The assumption of a homogeneous labor input is in contrast with the assumptions made in several other prominent lines of literature in economics, and these alternative models, applied to wage setting, may provide insight as to why labor markets appear to be fairly competitive for new workers, but much less so for existing workers.

The first potential explanation, one with a long history in labor economics, is the existence of firm-specific human capital (Becker, 1962). The well-known standard theoretical prediction of the Becker model is that firms pay only for firm-specific human capital, not for human capital whose applicability is general to all firms. While subsequent work has shown that firms' monopsony power can incentivize them to invest in their workers' general human capital training (Acemoglu and Pischke, 1998; Manning, 2003), a lesser-known but equally important implication of the Becker model is that in the presence of firm-specific human capital, firms must pay workers who acquire such a wage premium over what they would obtain elsewhere, to disincentivize those workers from quitting, which would be costly to the firm even if labor markets are frictionless. In contrast, firms are indifferent regarding turnover of generally trained employees in the Becker model. This implies a negative relationship between wages and turnover, as is observed in all studies of job separation behavior, including this one. However, it does not also imply that workers' wages in the presence of this relationship are marked down from their marginal revenue product in proportion with their separation elasticity, as described in Equation 1 and as is typically assumed in recent studies of heterogeneity in separation elasticities. Similarly, this model makes no assumptions of labor market frictions related to hiring activity, implying that wage setting for new hires is likely to be comparatively elastic.

Although the firm-specific human capital model can neatly explain the distinction between estimated new hire wage elasticities and separation elasticities, it is not clear that the magnitude of the differences observed here can be rationalized by such a model alone. For example, in Becker's firm specific human capital model, assuming that firms have knowledge of the probability of circumstances that lead to quits and layoffs, a separation elasticity of -0.4, as observed in this paper, would imply that a firm is willing to pay a 25% wage premium to reduce its probability of voluntary separation in a given year by 10%, suggesting extremely high turnover costs faced by the firm, even for low-wage workers. Additionally, the firm-specific human capital model cannot readily explain the results of heterogeneous effects regressions as shown in Table 8, all of which are supportive of search-based models of labor market frictions.

A second potential explanation may be the presence of firm-specific amenities that are *ex ante* unobservable to workers. If a particular worker cannot observe at the time of hiring how much he will enjoy a particular work environment or manager, then the firm that hires him may have a limited ability to hire him at a wage that is below his marginal revenue product in the presence of competition from other potential employers. [20] However, firms may be able to infer that workers who have been employed at the same establishment for a long period of time have preferences for firm-specific amenities and are therefore willing to accept lower wages, giving those firms more market power and allowing them to pay existing workers less than the value of their marginal product. Examples of amenities that may be ex ante difficult for workers to observe might include the quality of or degree of personal compatibility with one's manager or coworkers, the disutility associated with a new commuting pattern, the utility or disutility associated with performing job-specific tasks, or the utility associated with particular details of a job's benefits package.

The policy implications of a model of monopsony based on post-hiring revealed amenities are considerable. In particular, like the more well-known static hedonic model of wage setting, this model suggests that much of the wage dispersion observed across firms is ex-post efficient because it arises as a result of differences in worker preferences. On the other hand, such a model also implies that public policies that reduce the degree of firm-specific dispersion in amenities, such as the government provision of benefits that are otherwise provided heterogeneously by firms, could make labor markets less monopsonistic, and therefore increase the wages of tenured workers. Conversely, policies that make firm-specific amenities more observable prior to hiring, such as rules on the disclosure of benefits, might make workers more responsive to their own wages in their separation decisions (if searching for amenity-bearing firms is a source of search costs), but would have a lesser impact on wage dispersion than would be predicted by the static model relationship of Equation 1.

Identifying the most appropriate model of the labor market for studying variation in firms' monopsony power is a considerable task, and one that goes well beyond the scope of this paper. Furthermore, while matched employer-employee administrative data have significant advantages in their ability to identify the extent of monopsony in the labor market, they lack the kind of detailed within-establishment information such as individual team assignments,

[20] Acemoglu and Pischke (1998) construct a model of ex post monopsony power in which firms learn about the ability of workers only after they are hired to incentivize firms' investments in general human capital training. While a model such as that one in which workers ex ante have more information than firms may generate several similar predictions to a model in which firms have more information than workers, such a model does not by itself generate the types of heterogeneous effects results shown in Table 8.

training histories, or establishment-level benefits information that might be beneficial for studying these questions. Nonetheless, the evidence presented here makes it clear that while firms may have considerable ex post monopsony power, they have comparatively little such power ex ante. Given the wide-ranging implications of monopsony power in labor markets, it is clear that considerably more work is needed in order to understand where, when, and why this power is most prevalent.

References

PIS - Programa Integração Social — Caixa. `http://www.caixa.gov.br/beneficios-trabalhador/pis/Paginas/default.aspx`.

John M. Abowd, Francis Kramarz, and David N. Margolis. High Wage Workers and High Wage Firms. *Econometrica*, 67(2):251–333, March 1999. ISSN 1468-0262. doi: 10.1111/1468-0262.00020.

John M. Abowd, Robert H. Creecy, and Francis Kramarz. Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. April 2002.

Daron Acemoglu and Jorn-Steffen Pischke. Why Do Firms Train? Theory and Evidence. *The Quarterly Journal of Economics*, 113(1):79–119, 1998. ISSN 0033-5533.

David H. Autor, David Dorn, and Gordon H. Hanson. The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6):2121–2168, October 2013. ISSN 0002-8282. doi: 10.1257/aer.103.6.2121.

Erling Barth and Harald Dale-Olsen. Monopsonistic discrimination, worker turnover, and the gender wage gap. *Labour Economics*, 16(5):589–597, October 2009. ISSN 0927-5371. doi: 10.1016/j.labeco.2009.02.004.

Erling Barth, Alex Bryson, James C. Davis, and Richard Freeman. It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States. *Journal of Labor Economics*, 34(S2):S67–S97, February 2016. ISSN 0734-306X. doi: 10.1086/684045.

Gary S. Becker. Investment in Human Capital: A Theoretical Analysis. *Journal of Political Economy*, 70(5):9–49, 1962. ISSN 0022-3808.

Matthieu Bellon. The characteristics of worker flows by firm growth: Empirical evidence from a matched firm-worker dataset from France. 2016.

V. Bhaskar and Ted To. Minimum Wages for Ronald McDonald Monopsonies: A Theory of Monopsonistic Competition. *The Economic Journal*, 109(455):190–203, April 1999. ISSN 1468-0297. doi: 10.1111/1468-0297.00427.

Hoyt Bleakley and Jeffrey Lin. Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities. *Journal of urban economics*, 72(2-3):87–103, 2012. ISSN 0094-1190. doi: 10.1016/j.jue.2012.04.003.

Alison L. Booth and Pamela Katic. Estimating the Wage Elasticity of Labour Supply to a Firm: What Evidence is there for Monopsony?*. *Economic Record*, 87(278):359–369, September 2011. ISSN 1475-4932. doi: 10.1111/j.1475-4932.2011.00728.x.

Peter Brummund. Variation in monopsonistic behavior across establishments: Evidence from the indonesian labor market. In *Working Paper*. 2011.

Kenneth Burdett and Dale T. Mortensen. Wage Differentials, Employer Size, and Unemployment. *International Economic Review*, 39(2):257–273, May 1998. ISSN 0020-6598. doi: 10.2307/2527292. ArticleType: research-article / Full publication date: May, 1998 / Copyright © 1998 Economics Department of the University of Pennsylvania.

David Card. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64, 2001.

David Card and Alan B. Krueger. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton University Press, Princeton, NJ, twentieth-anniversary edition edition, December 2015. ISBN 978-0-691-16912-5.

David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of West German wage inequality. *NBER working paper*, (w18522), 2012.

David Card, Ana Rute Cardoso, Jörg Heining, and Patrick Kline. Firms and Labor Market Inequality: Evidence and Some Theory. 2016.

Sergio Correia. A feasible estimator for linear models with multi-way fixed effects. *Duke University Preliminary Version. URL: www. scorreia. com/research/hdfe. pdf*, 2016.

Steven J. Davis, John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. The MIT Press, Cambridge, Mass. London, reprint edition edition, January 1998. ISBN 978-0-262-54093-3.

Briggs Depew and Todd Sorensen. The elasticity of labor supply to the firm over the business cycle. *Labour Economics*, 24(C):196–204, 2013. ISSN 0927-5371.

Rafael Dix-Carneiro and Brian K. Kovak. Margins of labor market adjustment to trade. Technical report, National Bureau of Economic Research, 2017.

Torberg Falch. The Elasticity of Labor Supply at the Establishment Level. *Journal of Labor Economics*, 28(2):237–266, 2010. ISSN 0734-306X. doi: 10.1086/649905.

Li Gan and Qi Li. Efficiency of thin and thick markets. *Journal of Econometrics*, 192(1): 40–54, May 2016. ISSN 03044076. doi: 10.1016/j.jeconom.2015.10.012.

Christina Gathmann and Uta Schönberg. How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics*, 28(1):1–49, January 2010. doi: 10.1086/648416.

Fatih Guvenen, Burhanettin Kuruscu, Satoshi Tanaka, and David Wiczer. Multidimensional Skill Mismatch. Working Paper 21376, National Bureau of Economic Research, July 2015.

Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J. Redding. Trade and inequality: From theory to estimation. *The Review of Economic Studies*, 84(1):357–405, 2017.

Boris Hirsch, Thorsten Schank, and Claus Schnabel. Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany. *Journal of Labor Economics*, 28(2):291–330, 2010.

Boris Hirsch, Elke J. Jahn, and Claus Schnabel. Do Employers Have More Monopsony Power in Slack Labor Markets? *ILR Review*, page 0019793917720383, July 2017. ISSN 0019-7939. doi: 10.1177/0019793917720383.

Julie L. Hotchkiss and Myriam Quispe-Agnoli. Employer Monopsony Power in the Labor Market for Undocumented Workers. SSRN Scholarly Paper ID 1459045, Social Science Research Network, Rochester, NY, December 2009.

Michael Kremer and Eric Maskin. Wage Inequality and Segregation by Skill. SSRN Scholarly Paper ID 3631, Social Science Research Network, Rochester, NY, August 1996.

Peter Kuhn and Hani Mansour. Is Internet Job Search Still Ineffective? *The Economic Journal*, 124(581):1213–1233, December 2014. ISSN 1468-0297. doi: 10.1111/ecoj.12119.

Kevin Lang and Jee-Yeon K. Lehmann. Racial Discrimination in the Labor Market: Theory and Empirics. *Journal of Economic Literature*, 50(4):959–1006, December 2012. ISSN 0022-0515. doi: 10.1257/jel.50.4.959.

Kurt Lavetti and Ian Schmutte. Estimating Compensating Wage Differentials with Endogenous Job Mobility. *Labor Dynamics Institute*, August 2016.

Claudia Macaluso. Skill remoteness and post-layoff labor market outcomes. 2017.

Alan Manning. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, Princeton, N.J, March 2003. ISBN 978-0-691-11312-8.

Ioana Marinescu and Roland Rathelot. Mismatch Unemployment and the Geography of Job Search. Working Paper 22672, National Bureau of Economic Research, September 2016.

Jordan D. Matsudaira. Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations. *The Review of Economics and Statistics*, 96(1):92–102, 2014.

Naércio Aquino Menezes-Filho, Marc-Andreas Muendler, and Garey Ramey. The structure of worker compensation in Brazil, with a comparison to France and the United States. *The Review of Economics and Statistics*, 90(2):324–346, 2008.

Marc-Andreas Muendler, Jennifer Poole, Garey Ramey, and Tamara Wajnberg. Job Concordances for Brazil: Mapping the Classificaçao Brasileira de Ocupaçoes (CBO) to the International Standard Classification of Occupations (ISCO-88). *University of California, San Diego, unpublished manuscript*, 2004.

Walter Y. Oi and Todd L. Idson. Firm size and wages. In *Handbook of Labor Economics*, volume 3, Part B, pages 2165–2214. Elsevier, 1999.

Michael R. Ransom and Ronald L. Oaxaca. New market power models and sex differences in pay. *Journal of Labor Economics*, 28(2):267–289, 2010.

Michael R Ransom and David P. Sims. Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri. *Journal of Labor Economics*, 28(2):331–355, 2010. ISSN 0734-306X. doi: 10.1086/649904.

Jason M. Rivera. Pay Structure and Turnover: Evidence from Brazilian Employer-Employee Matched Data. SSRN Scholarly Paper ID 2356720, Social Science Research Network, Rochester, NY, December 2013.

Joan Robinson. *The Economics of Imperfect Competition, 2nd Edition*. Palgrave Macmillan, London, 2nd edition edition, July 1969. ISBN 978-0-333-10289-3.

Adriana Goncalves Saraiva. IBGE - Agência de Notícias. http://2.agenciadenoticias.ibge.gov.br/agencia-sala-de-imprensa/2013-agencia-de-noticias/releases/16131-ibge-divulga-as-estimativas-populacionais-dos-municipios-para-2017.html.

Mathan Satchi and Jonathan Temple. Labor markets and productivity in developing countries. *Review of Economic Dynamics*, 12(1):183–204, January 2009. ISSN 1094-2025. doi: 10.1016/j.red.2008.09.001.

Johannes Schmieder. What Causes Wage Dispersion: Evidence from New Firms. Technical report, October 2013.

Jae Song, David J. Price, Fatih Guvenen, and Nicholas Bloom. Firming Up Inequality. Technical report, National Bureau of Economic Research, 2015.

Sónia Torres, Pedro Portugal, John T. Addison, and Paulo Guimaraes. The Sources of Wage Variation: A Three-Way High-Dimensional Fixed Effects Regression Model. IZA Discussion Paper 7276, Institute for the Study of Labor (IZA), 2013.

Douglas A. Webber. Firm-Level Monopsony and the Gender Pay Gap. SSRN Scholarly Paper ID 2254196, Social Science Research Network, Rochester, NY, April 2013.

A Description of the PNAD Household Survey Data

The *Pesquisa Nacional por Amostra de Domicilios* (PNAD) is an annual survey of Brazilian households that is conducted in all non-Census years. It asks questions of household inhabitants regarding a wide range of topics, including detailed information on current employment. For the purposes of this paper, I seek to understand the impact that labor market informality has on the competitiveness of formal sector labor markets. Therefore, I use the PNAD data to construct an occupation by state measure of labor market formality as follows.

By definition, a formal sector worker is in possession of a *carteira de trabalho assinada* for their primary employment. Therefore, I define individuals in PNAD who report that they possess a *carteira* as formalized. Because my analysis focuses on private sector employment, I exclude government and military employees. I also exclude individuals who report did not report that they possess a *carteira*, but who report that they were self-employed or an employer themselves. For the purposes of Figure 4, public sector and self-employed workers are classified as neither formal nor informal.

Like many public use microdata, the PNAD data report households at a relatively coarse geographic level to inhibit identification of individuals. In this case, data are reported at the state level, which I am able to map to the Brazilian micro-region level. In keeping with the notion that labor market decisions are heavily both local and occupation-specific, I also seek to exploit variation in occupation-specific conditions, and so I allow this measure of variation to vary at the state by occupation level. Unlike the RAIS data, which report occupational codes using the Brazil-specific *Classificaçao Brasileira de Ocupaçoes* (CBO) classification, the PNAD microdata report occupation information using the internationally standard ISCO-88 classification system. I map two-digit ISCO-88 codes to Brazilian CBO codes using a concordance constructed by Muendler et al. (2004). Since this does not provide a one-to-one occupational mapping, where multiple ISCO codes map to a single CBO code, I weight the data by the proportions of respondents in the PNAD data who report each ISCO occupation. Finally, because it is likely that changes in the degree of formalization over time are reflective, I use only initial degree of labor market formality in heterogeneous effects regressions, incorporating data from the 1992, 1993, and 1995 PNAD surveys to reduce noise arising from the small sample sizes in occupation groups.[21]

Therefore, my final measure of labor market formality is:

$$Formality_{s,o} = \frac{N_{formal,s,o}}{N_{s,o}} \tag{10}$$

[21] There was no PNAD survey conducted in 1994 because the Brazilian Census was conducted in that year.

where $N_{formal,s,o}$ is the number of formal-sector workers in state s and occupation o in 1992-1995, and $N_{s,o}$ is total reported employment in that same occupation and state at that time.

B Additional Figures and Results

Figure 7: Unadjusted Age Wage Profile in New Worker Sample, All Years

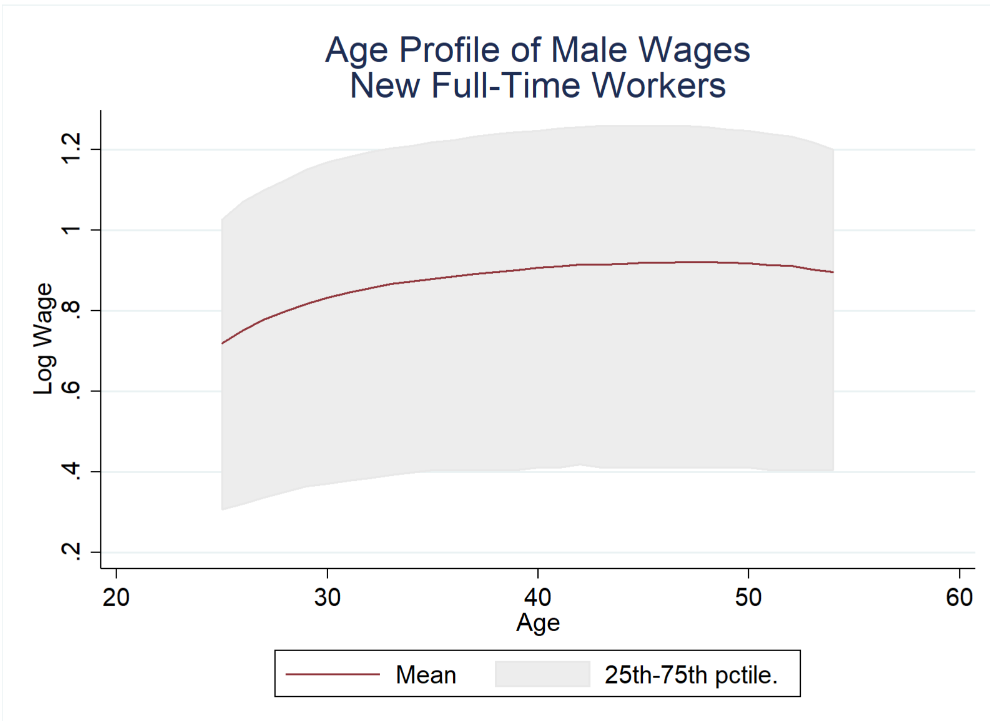

Notes: From the RAIS new worker sample, 1995-2014.

Figure 8: Tenure Wage Profiles, Indexed by New Local Worker Average Occupational Wage

Notes: From the RAIS work history sample, 1995-2014. Mean wage and interquartile range are indexed, where 100 is the mean wage for new (< 1 year of tenure) workers of the same occupation in the same Brazilian micro-region in that year.

Figure 9: Probability of Voluntary and Non-voluntary Separation by Years of Tenure

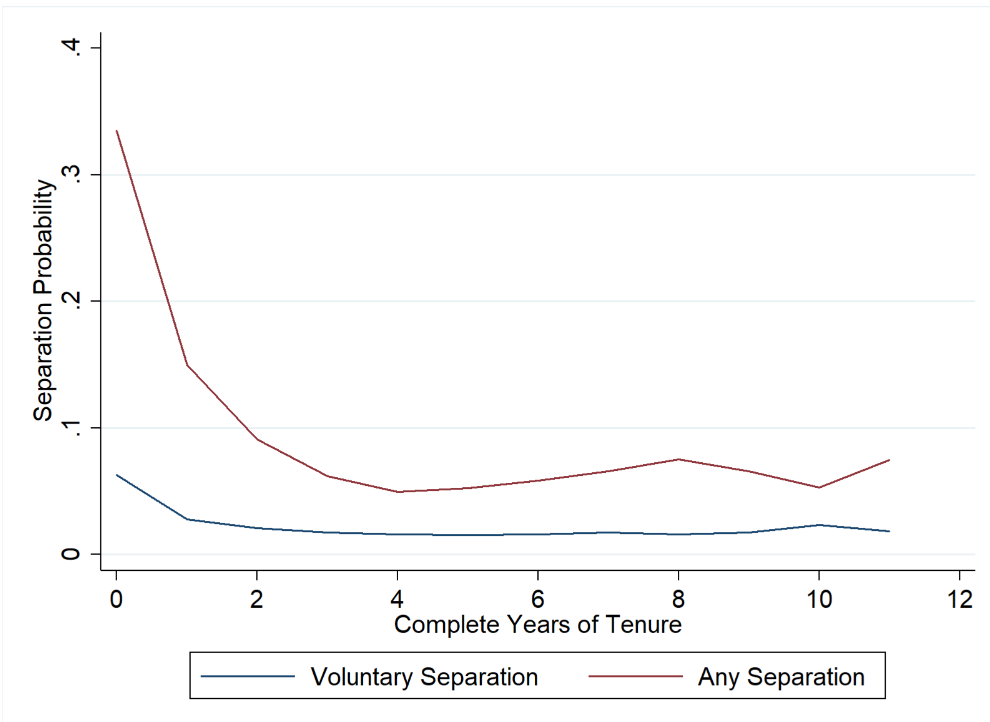

Notes: From the RAIS work history sample, 1995-2014. Voluntary and Non-Voluntary separation are as categorized in Table 9

Table 9: Job Separation Codes Reported in RAIS, 1995-2014

#	Label	Translation	% of Separations
Voluntary Separations			
20	desl com jc	resigned with just cause	0.23%
21	desl sem jc	resigned without just cause	20.11%
31	trans s/onus	transfer with cost to worker	5.36%
71	apos ts sres	retirement - length of service without contract termination	0.17%
78	apos id sres	retirement - age without contract termination	0.01%
80	apos esp sre	retirement - special without contract termination	0.01%
All Voluntary Separations			25.88%
Non-Voluntary Separations			
10	dem com jc	terminated with just cause	1.52%
11	dem sem jc	terminated without just cause	52.78%
12	term contr	end of contract	18.03%
30	trans c/onus	transfer with cost to firm	0.71%
40	mud. regime	change of labor regime	0.06%
50	reforma military	reform - paid reserves	0.06%
60	falecimento	demise, death	0.33%
62	falec ac trb	death - at work accident	0.01%
63	falec ac tip	death - at work accident corp	0.00%
64	falec d prof	death - work related illness	0.00%
70	apos ts cres	retirement - length of service with contract termination	0.28%
72	apos id cres	retirement - age with contract termination	0.04%
73	apos in acid	retirement - disability from work accident	0.02%
74	apos in doen	retirement - disability from work illness	0.02%
75	apos compuls	retirement - mandatory	0.04%
76	apos in outr	retirement - other disability	0.05%
79	apos esp cre	retirement - special with contract termination	0.01%
All Non-Voluntary Separations			73.97%
Unknown/Other			
-1, 22, 32-34, 90		Unknown/Other/No description available	0.14%

Notes: Percentages are calculated from all reported job separation events in RAIS, 1995-2014. English code translations from Lavetti and Schmutte (2016).

Table 10: OLS Estimates of the Establishment Size Wage Premium in Brazil

VARIABLES	(1) Log Earnings	(2) Log Earnings	(3) Log Earnings	(4) Log Earnings
Log Estab. Size	0.0690***	0.00270	0.0512***	0.0722***
	(0.00471)	(0.00289)	(0.00235)	(0.00333)
Observations	12,092,453	12,017,918	10,793,031	12,065,801
Adjusted R-squared	0.341	0.629	0.815	0.580
Estab. FEs	No	Yes	No	No
Individual FEs	No	No	Yes	No
Occ-Micro-Year FEs	No	No	No	Yes
Num. Clusters	397,556	323,021	358,852	396,372

Notes: Each result is from a regression of individual log December wages on establishment size, with columns 2-4 including one dimension of fixed effects as specified. All other controls are as in Table 2. Standard errors in parentheses are clustered by establishment. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 11: First-Stage Regressions for Baseline Wage-Setting IV Strategies

	(1) IV: Other Occs	(2) IV: Other Estabs.	(3) IV: Other Estabs./Occs.	(4) IV: Other Occs.
Other Input Growth	.3941232	.2571539	.1031151	.6412474
	(.0059599)	(.0111438)	(.0102324)	(.0089956)
Individual FEs	Yes	Yes	Yes	Yes
Establishment FEs	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes	Yes
K-P F Stat	4,373	532.5	101.6	5,082

Notes: Dependent variable in all regressions is DHS index of own-establishment own-occupation employment growth. In column 1, "Other Input Growth" is DHS index of other occupations within the same establishment, in columns 2 and 4 it is the same occupation within the same establishment, and in column 3 it is other occupations within the same establishment, corresponding to columns 2, 3, 4 and 6 of Table 2. Standard errors in parentheses are clustered by establishment.

Table 12: New Worker Wage Regressions Over Shorter Periods

	All Establishments		*Multi-Region, Multi-Estab. Firms*		
	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.	(4) OLS	(5) IV: Other Occs.
			Panel A: 1995-2001		
Occ Growth.	0.00330***	0.00865*	0.0439**	0.00154	0.00463
	(0.00103)	(0.00463)	(0.0191)	(0.00313)	(0.00666)
Observations	2,097,657	1,943,924	434,604	434,604	403,655
Adjusted R-squared	0.894			0.923	
Num. Clusters	88,371	63,786	15,697	15,697	9,522
Kleinbergen-Paap F Stat		1,392	126		2,694
			Panel B: 2002-2008		
Occ. Growth	0.00464***	0.0182***	0.0164**	0.00791***	0.0192***
	(0.000691)	(0.00367)	(0.00781)	(0.00191)	(0.00579)
Observations	3,167,182	2,979,514	769,224	769,224	719,355
Adjusted R-squared	0.894			0.922	
Estab. FEs	Yes	Yes	Yes	Yes	Yes
Num. Clusters	120,122	91,323	22,609	22,609	13,680
Kleinbergen-Paap F Stat		1002	207.6		708.7
			Panel C: 2008-2014		
Occ Growth.	0.00603***	0.0186***	0.00472	0.00754***	0.0178***
	(0.000703)	(0.00306)	(0.00764)	(0.00162)	(0.00384)
Observations	4,408,498	4,201,443	1,177,606	1,177,606	1,122,249
Adjusted R-squared	0.885			0.911	
Num. Clusters	153,907	122,015	27,891	27,891	18,728
Kleinbergen-Paap F Stat		1,996	192.5		2,595
Estab. FEs	Yes	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes	Yes	Yes

Notes: All regression specifications shown are as in the baseline regressions of Table 2, but with the panel period restricted to the seven-year interval shown. Standard errors reported in parentheses are clusterered by establishment. *** p<0.01, ** p<0.05, * p<0.1

Table 13: New Worker Wage Regressions Excluding Fixed Effects

	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.
Panel A: Excluding Individual FEs			
Occ. Growth	0.00250***	0.0239***	0.00363
	(0.000639)	(0.00255)	(0.00814)
Observations	11,066,654	10,505,238	3,041,017
Adjusted R-squared	0.734		
Num. Clusters	269501	211686	59435
K-P F Stat		3592	547.3
Panel B: Excluding Establishment FEs			
Occ. Growth	0.00205***	0.00954***	-0.00462
	(0.000470)	(0.00161)	(0.00484)
Observations	9,894,760	9,329,360	2,470,962
Adjusted R-squared	0.829		
Num. Clusters	292876	222060	58719
K-P F Stat		8391	647.4
Panel C: Excluding Local Labor Market FEs			
Occ. Growth	0.00862***	0.0145***	0.0209***
	(0.000543)	(0.00237)	(0.00713)
Observations	9,875,353	9,324,335	2,475,758
Adjusted R-squared	0.858	-0.247	-0.346
Num. Clusters	252768	194529	50353
K-P F Stat		3550	448.8

Notes: All regression specifications shown are as in the baseline regressions of Table 2, but with one category of fixed effects excluded from the regression specifications in each panel. Standard errors reported in parentheses are clusterered by establishment. *** p<0.01, ** p<0.05, * p<0.1