2018

Three essays on labor markets

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

THREE ESSAYS ON LABOR MARKETS

by

LEE TUCKER

B.A., Carleton College, 2006
M.A., Boston University, 2018

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2018

Approved by

First Reader

M. Daniele Paserman, PhD
Professor of Economics

Second Reader

Johannes F. Schmieder
Associate Professor of Economics

Third Reader

Samuel Bazzi
Assistant Professor of Economics

Acknowledgments

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, the Development Reading Group, and graduate seminars.

Finally, thank you to my loving wife Heather, who has helped me through every challenge and has put up with me through it all.

Lee Tucker

Department of Economics

THREE ESSAYS ON LABOR MARKETS

LEE TUCKER

Boston University, Graduate School of Arts and Sciences, 2018

Major Professor: M. Daniele Paserman, PhD
Professor of Economics

ABSTRACT

The recent proliferation of administrative data sources has made it possible to examine numerous longstanding questions related to labor market functions. I make use of these data sources to provide new insights into three such questions; the extent of firms' market power in labor markets, the nature of gains from workers' skill specialization, and the role of job search networks in the locational choices of immigrants.

In Chapter 1, I examine labor market monopsony, the extent to which markets deviate from perfect competition. Prior literature suggests two methods to estimate the extent of monopsony: studying the degree to which firms adjust wages in response to desired changes in employment growth, and measuring the degree to which workers' voluntary separations are sensitive to their own wages. Existing studies have found widely varying answers to these two questions in different contexts. I leverage unique features of Brazilian administrative data to demonstrate that these approaches provide very different results even on the same sample of employees, and I rule out a variety of alternative empirical explanations. These results suggest that labor market monopsony is primarily a function of workers' attachment to their current employers.

In Chapter 2, I study the wage premium associated with skill specialization. While standard models predict that more technologically-advanced firms will hire more specialized workers, I show that higher-ability individuals may actually sort

into less specialized occupations within firms. I test these predictions by constructing occupation-level measures of skill specialization from the U.S. O*NET database, matched to Brazilian administrative data. While I find that specialization among production skills is associated both with higher wages and with employment at higher-wage firms, I find no evidence of specialization premia in cognitive skills.

Finally, in Chapter 3 I study the extent to which job search networks influence new immigrants' decisions to locate in ethnic enclaves. Using detailed data from the New Immigrant Survey, I show that immigrants to the U.S. who arrive without job offers are significantly more likely to locate in enclaves, even after accounting for a wide range of pre-migration and time-invariant characteristics.

Contents

List of Tables

List of Figures

Chapter 1

Monopsony for Whom? Evidence from Brazilian Administrative Data

1.1 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

[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. (2017) show that even a static model with heterogeneous worker preferences over firm-specific amenities is sufficient to generate monopsonistic behavior.

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, Heining, and Kline, 2013; 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 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 chapter, 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, Kramarz, and Margolis (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, Dorn, and Hanson, 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 chapter 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 chapter proceeds as follows: Section 1.2 provides a brief overview of the existing literature on the topic of monopsony, Section 1.3 describes both of the primary empirical specifications used in this chapter along with the IV strategies used, Section 1.4 describes the RAIS data along with household survey data used for analyses of heterogeneous effects, Section 1.5 provides baseline elasticity estimates using both specifications, Section 1.6 provides a variety of robustness checks of the new worker wage specification, Section 1.7 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 1.8 concludes the chapter 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.

1.2 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:

$$w = \frac{MRP\left(L\right)}{1 + \frac{1}{\epsilon^L_{Sw}}} \tag{1.1}$$

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 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:

$$L_t\left(w_t, L_{t-1}\right) = R\left(w_t\right) + \left[1 - s\left(w_t\right)\right] L_{t-1} \tag{1.2}$$

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 ϵ_{Rw}, and an elasticity of job separations from the establishment ϵ_{Sw}. With some straightforward algebra, (1.2) can be written as:

$$\epsilon_{Lw} = \epsilon_{Rw} - \epsilon_{Sw}\frac{L_{t-1}}{L_t} \tag{1.3}$$

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. In particular, this allows for additional flexibility in considering the relationships between each elasticity over the business cycle, at the cost of stronger assumptions about the nature of on-the-job search.[4] However, in empirical practice, most authors

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

[4]Hirsch, Jahn, and Schnabel (2017) provide a useful derivation of the relationship between the

have chosen to make the simplifying assumption that the aggregate labor market is in a steady state, with a constant level of unemployment. On net, this implies that each recruit to a firm occurs one-for-one with a separation from another firm, and so it follows that $\epsilon_{Rw} = -\epsilon_{Sw}$, and so:

$$\epsilon_{Lw} = -\left(1 + \frac{L_{t-1}}{L_t}\right)\epsilon_{Sw} \tag{1.4}$$

If it is additionally assumed that the firm is at a steady state level of employment, so that $L_t = L_{t-1} = L$, then the above can be simplified to:

$$\epsilon_{Lw} = -2\epsilon_{Sw} \tag{1.5}$$

While the steady state assumptions needed to get this particular equation are strong, Equation 1.5 implies a clear relationship between the elasticity of voluntary separations from establishments and the overall labor supply elasticity faced by those establishments. If the 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. In particular, if the average existing establishment is growing slowly in size, then the separation elasticity will actually be closer to the overall labor supply elasticity than would be observed in a steady state.

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

separation elasticity and the overall wage elasticity of the firm's labor supply when the steady state constraints are relaxed.

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 chapter (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 1.1. 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, Schank, and Schnabel, 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 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 chapter 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 chapter. 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 chapter 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, Kramarz, and Margolis, 1999; Abowd, Creecy, and Kramarz, 2002; Card, Heining, and Kline, 2013), 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 1.7 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.

1.3 Overview of Empirical Specifications

The goal of this chapter 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.

1.3.1 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 chapter, the basic equation that I estimate is of the form:

$$\log w_{iomt} = \alpha G_{ot,j(i,t)} + X_{it}\beta + \delta_{j_{(i,t)}} + \theta_i + \psi_{omt} + \epsilon_{iomt} \tag{1.6}$$

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, Haltiwanger, and Schuh (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 employment is zero, and its bounded and symmetric nature addresses potentially large asymmetries between small employers

[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*.

who grow or decline by the same amount.

The specification of Equation 1.6 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 1.6, I am also able to individually isolate these two sources of variation.

[6] The strategy of incorporating overlapping worker and establishment fixed effects, developed by Abowd, Kramarz, and Margolis (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 1·1: Potential Simultaneity Bias in a Monopsonistic Firm Without Local Labor Market Fixed Effects

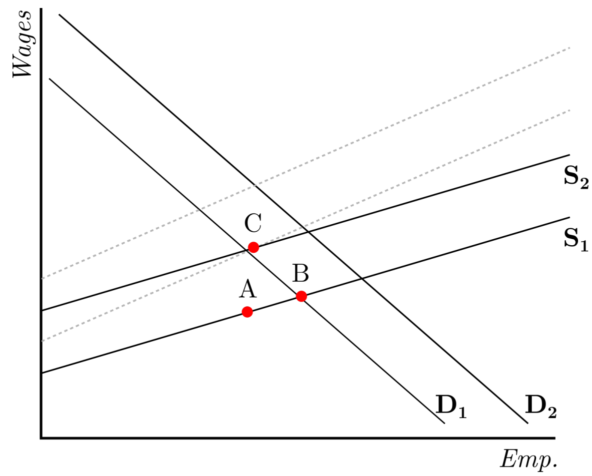

While the restricted sources of variation used to identify the parameter of interest α in this specification help to rule out selection concerns, there may still remain substantial concerns related to simultaneity bias in estimation of firms' labor supply elasticities. Figure 1·1 illustrates this concern 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.

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 1.6 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 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]

[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).

[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.

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 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 1.6 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 chapter, 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.

1.3.2 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:

$$S_{iomt} = \alpha \log w_{iomt} + X_{it}\beta + \delta_{j_{(i,t)}} + \theta_i + \psi_{omt} + \epsilon_{iomt} \tag{1.7}$$

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 1.2 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

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

elasticity that are biased upward (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 chapter 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.

1.4 Description of Data Sources

In this section of the chapter, I describe the primary source of data for the analyses of this chapter, 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.

1.4.1 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

[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 (*PIS - Programa Integração Social / Caixa*).

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, Jahn, and Schnabel, 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 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)

[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.

are shown in Appendix Table A.1. 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 1.7. 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 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 chapter 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 chapter 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

[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 potential 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*.

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 1·2: Map of 10% Random Sample of Brazilian Micro-Regions

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

Figure 1·2 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 (*IBGE - Agência de Notícias*).

Figure 1·3: Histogram of DHS Growth Index for New-Worker Sample

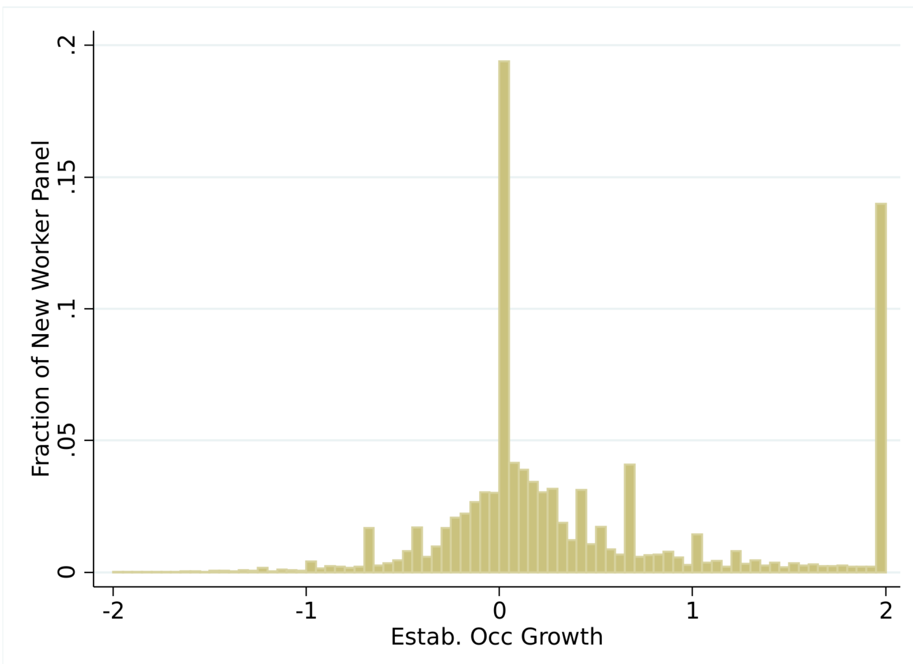

Notes: From RAIS new worker sample, 1995-2014, with restrictions as described in Section 1.4.1. The DHS index is calculated as $2 \times \frac{L_t - L_{t-1}}{L_t + L_{t-1}}$ for each occupation employed within each establishment.

Table 1.1 provides some basic summary statistics on the new worker samples that are 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

Table 1.1: 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.

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 1·3 shows the distribution of occupation-level employment growth within the new worker sample, using the DHS index as described in Section 1.3. 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

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

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 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.

1.4.2 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 1·4 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 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

Figure 1·4: Quantile Plot of Monthly Incomes for Formal and Informal Sector Employed (1995)

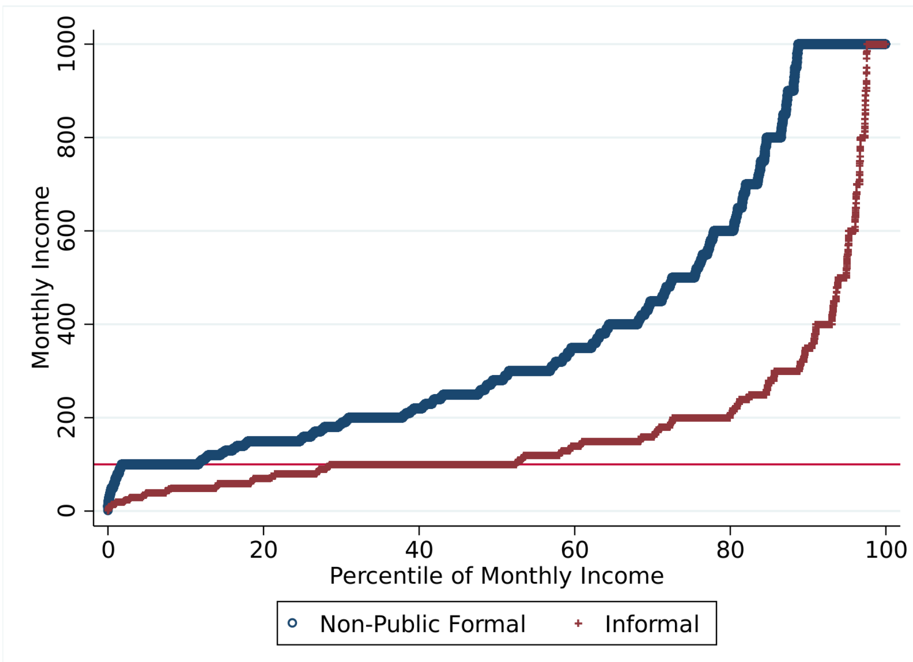

Notes: Income data from PNAD. Non-public formal workers include all individuals who report employment with a *carteira* excluding government employees and military. Informal workers includes workers who report that their employment is without a *carteira*, or who report that their employment is production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

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.

1.5 Results of Both Empirical Specifications

1.5.1 Wage-Setting Specification

Table 1.2: 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	0.00494***	0.0181***	0.0125***	0.0627***	0.00733***	0.0203***
	(0.000358)	(0.00234)	(0.00147)	(0.0167)	(0.00148)	(0.00362)
Implied Elasticity	202.4	55.2	80.0	15.9	136.4	49.3
	[177,236]	[44,74]	[65,104]	[10,33]	[98,226]	[37,76]
Observations	9,844,518	9,293,995	2,458,636	3,195,334	3,234,340	3,036,790
Adjusted R-squared	0.870				0.905	
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
K-P F Stat		324.5	914.6	152.9		1,414

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, quadratic and cubic age profile terms and month of hire controls. Standard errors in parentheses are two-way clustered by firm and micro-region. 95% confidence interval for elasticity estimate in brackets is calculated by the delta method. *** p<0.01, ** p<0.05, * p<0.1

Table 1.2 shows the results of my baseline wage-setting specification, as indicated by Equation 1.6. Column 1 shows the results for the OLS specification, while columns 2 through 4 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

[15]Menezes-Filho, Muendler, and Ramey (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.

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., 2017).[16] Finally, note that in all specifications, standard errors are two-way clustered at the firm and micro-region levels to ensure that inferences are robust to within-firm and within-region 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

[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 A·1 in the Appendix shows the unadjusted age-wage profile observed in the RAIS data, which supports this assumption for my sample.

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.

Figure 1·5: Non-parametric OLS Estimation of New Worker Wage Premia

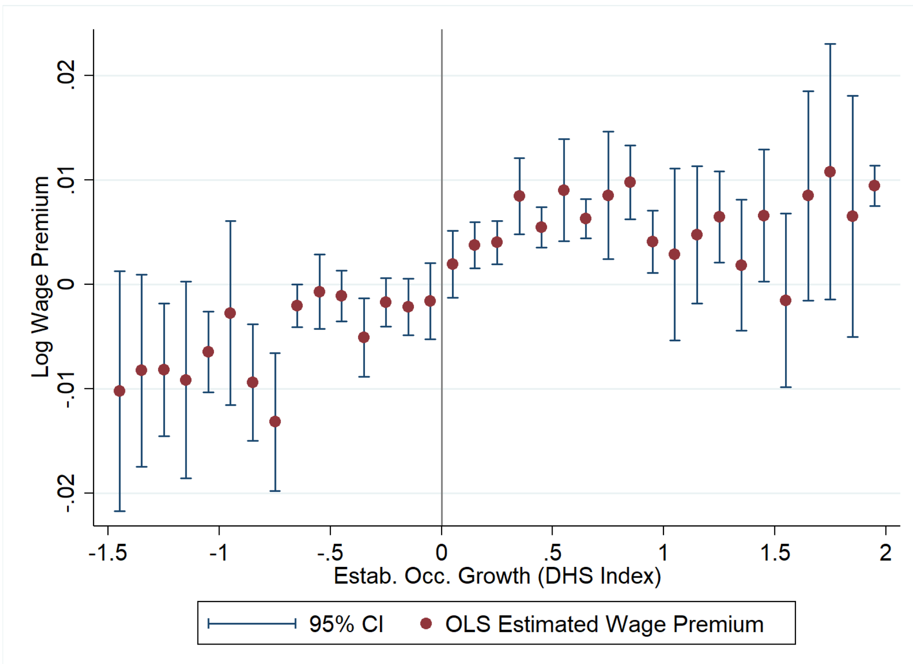

Notes: From the RAIS new worker sample, 1995-2014. Coefficients and confidence intervals are from an OLS regression with 40 indicator variable indicating employment growth in each bin of DHS employment growth index of width 0.1. The omitted category is firms that reported exactly zero growth in occupational employment. All other covariates are as specified in Column 1 of Table 1.2. Standard errors for confidence interval construction are clustered by establishment.

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 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 1·5. 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 1.6, I show these estimates to be fairly robust to a range of other potential concerns.

1.5.2 Job Separations Specification

As discussed in the introduction to this chapter, 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 1.3.2, 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 1.3: Baseline Results of Separations Specification

	(1) OLS	(2) IV: Initial Wage	(3) OLS w/Add'l FEs
Log December Earnings	-0.0175***	-0.00700**	-0.0159***
	(0.00140)	(0.00282)	(0.00157)
Pr(Separation)	0.0507	0.0199	0.0509
Implied Separation Elasticity	-0.344	-0.351	-0.312
	[-0.291,-0.399]	[-0.074,-0.630]	-[0.252,-0.372]
Adjusted R-squared	0.198		0.380
Observations	26,313,032	7,638,242	22,380,842
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		1,135	

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 two-way clustered by firm and micro-region. 95% confidence interval for elasticity estimate in brackets is calculated by the delta method. *** p<0.01, ** p<0.05, * p<0.1

Table 1.3 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 1.2, 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 1.3 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 1.2 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.

1.6 Robustness Checks

While the results presented in the prior sections of this chapter are strongly suggestive, there may still be particular questions arising from the specifications used. In this section of the chapter 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.

1.6.1 Testing for Endogenous Substitution

The results for the IV specifications shown in Table 1.2 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 chapter 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 1.3.1, 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:

$$M_{o,o'} = \frac{\sum_i \mathbf{1}(m_{io} \cap m_{io'})}{\min\{N_o, N_{o'}\}} \tag{1.8}$$

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 1.2 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 1·6 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

Figure 1·6: Ranked Scatterplot of Occupational Pairs by Ever-Transition Probability

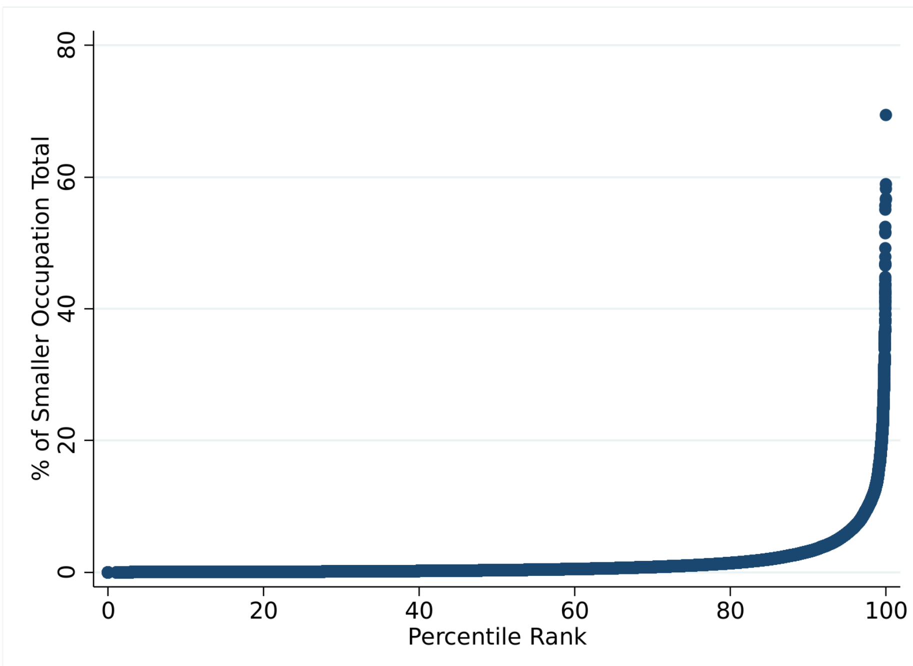

Notes: From RAIS full dataset, 1986-2014. Each pairwise ever-transition probability is calculated as the probability of an individual ever appearing in RAIS as employed in one occupation, conditional on them ever reporting employment in the other occupation. The larger of the two conditional probabilities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

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 1.4: 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 1.2. Standard errors in parentheses are clustered by establishment. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 1.4 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 1.2, 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 1.2 (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.

1.6.2 Additional FEs

As discussed in Section 1.3, 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 1.5 shows the results of each of these regression specifications.

Table 1.5: Baseline Regressions with More Restrictive Fixed Effects Specifications

	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.	(4) IV: Other Occs./Estabs.
		Panel A: Estab-Occ. FEs		
Occ. Growth	0.00458***	0.0151***	0.0151***	0.0573***
	(0.000419)	(0.00185)	(0.00266)	(0.0125)
Observations	9,552,645	9,014,500	2,414,326	3,124,873
Adjusted R-squared	0.890			
		Panel B: Estab-Year FEs		
Occ. Growth	0.000921*	0.00184**	-0.00179	-0.00373
	(0.000472)	(0.000693)	(0.00185)	(0.00278)
Observations	9,235,626	8,945,168	2,337,718	3,048,046
Adjusted R-squared	0.888			
		Panel C: Occ.-Ind.-Micro-Year FEs		
Occ. Growth	0.00490***	0.0163***	0.0123***	0.0556***
	(0.000330)	(0.00241)	(0.00354)	(0.0106)
Observations	9,694,195	9,147,793	2,420,991	3,131,404
Adjusted R-squared	0.875			
Individual FEs	Yes	Yes	Yes	Yes

Notes: All results are from regressions on the RAIS new worker sample as in columns 1 through 3 of Table 1.2, 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 two-way clustered by firm and micro-region. *** p<0.01, ** p<0.05, * p<0.1

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 × region × year level could be insufficient to address local labor market heterogeneity if the 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 1.5, I present results in which I replace the occupation × region × year fixed effects with occupation × industry × region × 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.

1.6.3 Using aggregate occupational growth

The approach of this chapter 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 1.2 use occupation-level employment growth as the relevant measure of employment growth, and it uses occupation × region × 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 1.6 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 and with the inclusion of occupation-level local labor market FEs, while columns 2 and 4 apply the IV strategy of instrumenting for establishment em-

Table 1.6: Regressions: Aggregate-Level Employment Changes

	(1) OLS	(2) IV: Other Estabs.	(3) OLS Multi-Estab.	(4) IV: OLS	(5) OLS Other Estabs.	(6) Multi-Estab.
Estab. Growth	0.0101***	0.0512***	0.0136***	0.0105***	0.0433***	0.0133***
	(0.00145)	(0.0187)	(0.00301)	(0.00128)	(0.0129)	(0.00308)
Observations	9,875,692	3,257,557	3,257,557	9,844,518	3,234,340	3,234,340
Adj. R-squared	0.858		0.894	0.870		0.905
Estab. FEs	Yes	Yes	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	No	No	No	Yes	Yes	Yes
K-P F Stat		118.3			250.1	

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 1.2. Standard errors in parentheses are two-way clustered by firm and micro-region. *** p<0.01, ** p<0.05, * p<0.1

ployment growth using non-local firm establishment growth of multi-establishment, multi-region firms, analogously to columns 3 and 4 of Table 1.2. 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 chapter—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.

1.7 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 chapter. While the choice reflects the

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

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:

$$\log\left(N_{m,-f,o,t}\right) = \log\left(N_{m,o,t} - N_{m,f,o,t}\right) \tag{1.9}$$

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

[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.

is engaged in that occupation. That is:

$$Prev_{mo,-ft} = \frac{N_{m,-fot}}{N_{m,-ft}} \Bigg/ \frac{N_{ot}}{N_t} \tag{1.10}$$

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, 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 1·4), 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 1.7 shows the results of these heterogeneous effects regressions on the new worker wage specification, while Table 1.8 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

[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 1.7: 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 1.2, 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 1.8: 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 1.3, 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

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 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 chapter 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 1.7, it may simply be an issue that there is comparatively little variation in wages left to exploit.

1.8 Discussion of Alternative Models

As described in the introduction to this chapter, 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 chapter 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 chapter 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 chapter, 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 chapter strongly rules out the notion that external validity or selection concerns have driven prior results.

The heterogeneous effects results shown in this chapter 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 1.1 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 chapter, 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 1.8, 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

[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

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 1.1.

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 chapter. Furthermore, while matched employer-employee administrative data have significant advantages in their ability to identify the extent of monopsony

information than workers, such a model does not by itself generate the types of heterogeneous effects results shown in Table 1.8.

in the labor market, they lack the kind of detailed within-establishment information such as individual team assignments, 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.

Chapter 2

The Division of Labor and the Labor Market: Are Specialized Worker Skillsets Valued?

2.1 Introduction

In *The Wealth of Nations*, Adam Smith famously describes the productivity gains from specialization in a pin factory. He recounts in detail the manner in which pin-making tasks can be divided among workers to speed the pin-making process, before concluding that "[t]he division of labour ... so far as it can be introduced, occasions, in every art, a proportionable increase of the productive powers of labour" (Smith and Nicholson, 1887). This example has long been cited to describe a world in which economic growth leads to ever-increased specialization, limited only by "the extent of the market."

Yet, Smith does not describe the exact mechanisms by which specialization may be valuable. In a pin factory, it may be reasonable to believe that specialization directly increases the productivity of individual steps, each of which has a well-defined output. Specialization could also reduce the time spent transitioning between steps. However, there are many contexts, particularly in highly-skilled production activites, where the scope for these sources of gains from specialization is less clear.

This chapter considers another possibility: that gains from specialization exist because human capital is in part "task-specific". That is, if specialization allows workers to narrow the breadth of their knowledge and instead increase the depth of their knowledge, then specialization could be a source of increasing returns to scale. While these returns to scale may be limited by forces other than the extent of the market,[1] they still provide an incentive for tasks to be allocated to workers in non-overlapping way so that they may specialize to the greatest feasible extent.

If specialization is a primary source of increasing returns to scale, then we might also reasonably expect that individuals in more specialized occupations will earn higher wages. But why exactly might this be the case, and are there instances in which specialization may actually not be valued? To answer these questions, I turn to the seminal model of specialization and occupational choice of Becker and Murphy (1992). This model relies on the notion that for individuals who specialize, a portion of the human capital that each employee develops is task-specific, and so their overall stock of human capital is multidimensional. The skills associated with human capital may be developed broadly, or they may be developed deeply, but expanding one's skillset in either direction is costly because individuals have limited time to allocate between production and human capital acquisition. I demonstrate that when labor markets are imperfectly competitive, otherwise-identical workers who are specialized in their observable skillsets earn more than workers who are less specialized, because the feasible extent of specialization is increasing in the quality of the firm's production technology. In other words, specialized workers earn more because they work for more productive firms. However, since this model assumes that all workers are identical, it does not make predictions about the wage premia associated with more specialized workers within each firm.

[1] A substantial theoretical literature has sought to consider potential limits to the degree of specialization other than the extent of the market. These papers have focused on such issues as market power (Baumgardner, 1988), coordination and/or communication costs (Becker and Murphy, 1992; Dessein and Santos, 2006), decision implementation costs (Radner and Zandt, 1995), and institutional quality (Costinot, 2009).

In response to this limitation, I extend the model to introduce heterogeous underlying ability of employees–defined here as the amount of task-specific human capital that can be acquired in a period of time. Because there is no inherent division of "high-ability" or "low-ability" tasks in the model, each firm has the unrestricted ability to assign tasks to different types of workers in order to maximize profits. In this simple framework, I find that the positive association between ability and specialization within firms need not hold at all. Rather, the extent to which high-ability and low ability workers are more or less specialized depends on the elasticity of substitution between productive tasks. If, for example, tasks are perfect complements in production, then high-ability workers within a firm will actually be less specialized than lower ability workers, and specialization within firms will be associated with lower wages, not higher ones. Put more succinctly, the observable wage premia to specialized workers don't just depend on the productivity of the firm they work for; they also depend on the trade-offs that firms face in employing workers of different quality.

Testing the theoretical predictions of this model requires information about the human capital composition of different occupations, which I am able to construct from the U.S. Department of Labor's O*NET database. The O*NET database provides information on the level of usage of 35 different skills for each of over 900 occupations, and from these data, I construct measures of skill concentration, using the Gini coefficient of reported skill usage as a baseline measure of specialization. However, it also requires matched employer-employee data in order to distinguish between inter-firm and intra-firm wage patterns. Suitable administrative data are not available for the United States, since U.S. administrative data source do not include data on employees' occupations. So, I turn to detailed matched employer-employee data from the Brazilian *Relação Anual de Informações Sociais* (RAIS) dataset, using existing corcordances to match Brazilian occupations to Census occupations as reported in O*NET. To my knowledge, this is the first paper to construct occupation-level measures of skill specialization, and therefore to attempt to test the gains from specialization empirically. While I do not have access to exogenous variation in occupation or in individual human capital, I am able to use broader occupational group fixed effects and other measures of observed ability to compare the wages of individuals who choose among related occupations.

Surprisingly, when I first regress log wages on occupation-level aggregate mean skill level and degree of specialization in a standard Mincer-style specification, I find no evidence of wage premia associated with specialization either within or across firms. These results are robust to the inclusion of a variety of detailed individual and establishment level controls chosen carefully to mitigate concerns regarding omitted sources of occupation-level differences. Robustness checks constructed using the U.S. Current Population Survey (CPS) provide similar results. In short, these results suggest that while skill is valuable, the extent to which occupations are specialized in their skill usage is not associated with either higher or lower wages. Indeed, these initial results provide no empirical evidence for gains from specialization at all, and they are fundamentally in contrast to the predictions of the Becker and Murphy model.

One concern is that the model may simply be too simple to characterize all production activities well for all types of workers, or that specialization may only be valuable among some skills but not among others for particular types of workers. To address this possibility, I adopt a method of a method of agglomerative hierarchical clustering from the machine learning literature, allowing me to classify the 35 skills listed in O*NET based on the patterns of correlation in the levels of each skill reported across occupations. I demonstrate that observable skills can be readily divided into two basic categories, which I refer to as "Production" skills and "Cognitive" skills, with occupations that require high levels of skill in one group typically requiring low levels of skill in the other. Perhaps unsurprisingly, these skill groups also map closely to the divide between "blue-collar" and "white-collar" occupations.

When I regress log wages on the mean skill level and the degree of specialization within each group, I find much more evidence to support the notion that there are gains from specialization, but only among Production skills. In my preferred baseline specifications, a one standard-deviation increase in the Gini coefficient of production skills is associated with a 4-6% increase in earnings, even after controlling for a variety of individual characteristics, after controlling for occupations' mean level of skill, and after using fixed effects to control flexibly for differences in broader occupational categories. When I incorporate establishment fixed effects, I find evidence that specialization within production skills is associated with wage premia even within establishments, though the results are somewhat smaller than would be suggested by looking across firms. Results limited to blue-collar workers (who tend to be relatively more skilled in production skills) suggest even larger benefits to specialization in production skills than the baseline results would suggest. Conversely, specialization in one's cognitive skillset is not associated with higher wages either within or across firms, either for all workers or for white collar workers specifically. Again, these results are largely supported by robustness checks using the CPS.

Finally, since the Becker and Murphy model predicts that gains from specialization arise primarily from sorting of specialized individuals to highly productive firms, I construct a more direct test of sorting patterns using a two-step procedure based on the computational method of of Abowd, Kramarz, and Margolis (1999) (AKM). In the first stage of this procedure, I identify both individual worker and establishment fixed effects for approximately 95% of the worker-level observations in my baseline wage sample. Under the assumptions of the baseline model, the estimated establishment fixed effects identify the firm-level component of wages and are closely related to total factor productivity differences across firms. In the second step, I then regress these firm fixed effects on the occupation-level measures of skill content associated with each individual's occupation. While I find that individuals in occupations with more specialized production skillsets work at higher wage firms, as predicted by the model, I also find that employment at a high-wage firm is associated with employment in an occupation with a broader cognitive skillset. Again, these results do not appear to be driven by differences in observable worker characteristics, worker mean reported skill levels, or broader occupational controls.

These results strongly suggest that while this standard model of specialization may accurately describe the role of specialization in production for blue-collar workers and their skillsets, it does not effectively characterize the relationship between skill specialization and wages well for white-collar workers or the cognitive skills which they tend to possess. Moreover, the evidence provides little to support the notion that there are gains from specialization within one's cognitive skills at all.

The chapter proceeds as follows: Section 2.2 provides a brief theoretical and empirical overview, Section 2.3 develops the modified Becker and Murphy model and outlines its empirical predictions, Section 2.4 describes the empirical specifications in use to look at both worker wage premia and worker sorting, Section 2.5 describes briefly the datasets and methods that are used, Section 2.6 provides regression results on both worker wage premia and on worker sorting, and Section 2.7 concludes the chapter.

2.2 Theoretical and Empirical Background

In considering the relationship between human capital acquisition and the gains from specialization, we typically wish to focus on the gains from specialization in production within firms, rather than across firms. While the traditional Ricardian framework of gains from trade based on comparative advantage characterizes potential gains from specialization at the firm level (i.e. in the scope of production), as well as specialization across countries or within households, it does not translate naturally to characterizing the gains from task specialization among workers within firms. That is, Ricardian comparative advantage alone does not explain why, say, some firms that perform construction work choose to hire plumbers and electricians, while others choose to hire more general "handymen." For this, more detailed models of firms' production structure and the trade-offs involved are needed. So, models of the gains from specialization in production within firms generally require three basic features:

1. A source of increasing returns to scale in production. This source may take many different forms, such as input complementaries, increased leverage of comparative advantage by workers, or task-specific hunan capital. However, if there are no increasing returns to scale, then there is in general no reason for anyone to prefer a multi-worker firm (in which specialization is feasible) to a single-worker firm in which within-firm specialization is infeasible by definition.

2. A countervailing force that bounds increasing returns to scale. Examples in the theoretical literature abound, including coordination, communication, and enforcement costs (Becker and Murphy, 1992; Costinot, 2009), adaptation costs (Dessein and Santos, 2006) market power of workers (Baumgardner, 1988), and capital market imperfections (Fishman and Simhon, 2002). Absent any other such force, the only feasible equilibrium is monopoly within each product

market, i.e. the degree of specializiation being limited by the extent of the market.

3. A reason for the firm or for individuals within the firm to prefer some degree of specialization in their productive activities to simply involving all individuals in all aspects of production. This may be the same as the source of increasing returns to scale itself (e.g. higher productivity arising from specialization or increased focus on one's comparative advantage), or it may be a result of features of the countervailing force (e.g. reduced coordinaton costs through specialization).

There are likely numerous ways to construct a model that embodies all three of these features. However, and importantly, none of these requirements explicitly demands that what we typically think of as human capital play any role in specialization at all. For example, we might expect plumbers and electricians to be more productive than handymen simply because the former spend less time transitioning between tasks, gathering different sets of tools, or hauling a wide range of needed supplies. While these explanation imply gains from specialization, they also imply that if two equally able workers who are involved in two different portions of the production process, reassigning the tasks of each worker to the other is likely to be costless. On the other hand, it could also be the case that specialized workers are able to develop some form of task-specific human capital. In this case, reassigning the tasks of workers once human capital is acquired might greatly reduce productivity.

There is reason to suspect that this latter skill-based notion of gains from specialization is important. Firstly, a growing empirical literature supports the notion that human capital is a composition of skills, and that occupations are often distingished by the particular combination of skills that they require and tasks that they perform rather than by the presence of unique skills (Poletaev and Robinson, 2008; Peri and

Sparber, 2009; Kambourov and Manovskii, 2009; Deming, 2015). Additionally, recent evidence suggests that individuals' labor market success depends in part on the extent to which their prior experiences or initial aptitude are similar to those of the job they seek or the market in which they participate (Gathmann and Schönberg, 2010; Guvenen et al., 2015; Macaluso, 2017). While many of these papers have considered either the role of strength in particular skill dimensions or of distance across all skill dimensions. However, to my knowledge, this chapter is the first that has looked specifically at the extent to which occupational skills are concentrated or dispersed.

Perhaps the theoretical model that most neatly presents these underlying features necessary to look at skill specialization is the model of Becker and Murphy (1992). In this model, each firm's production process consists of a unit measure of unique tasks, and so a worker who is specialized uses their time to perform only a subset of tasks, spending a greater fraction of their time on each. Becker and Murphy do not explicity define the notion of "task-specific human capital" in their paper, but the notion is implicit in the fact that workers increase their productivity only by spending a portion of their time engaging in a task-specific skill investment.

While the Becker and Murphy model is tractable and provides clear predictions, it also has notable shortcomings from the perspective of any empirical analysis. Most notably, the model does not explicitly define the notion of an occupation; if an occupation is a collection of tasks, then there may be infinitely many occupational definitions corresponding to different collections of these tasks. That is, a firm which needs both electrical work and plumbing work to be performed might hire an electrician and a plumber, or it might hire a "handyman" with some lesser skill in each of these tasks. Since most empirical analyses of workers' skills are based on the skill content of their occupations, the lack of a clear mapping to occupations is unintuitive. Additionally, there is no explicit notion of occupational choice in this model; task assignment is implicitly the same thing as occupational choice, and since the model is static, task assignment and human capital acquisition occur simultaneously. Furthermore, the model implies a one-to-one mapping between production tasks and skills, with little possibility for the existence of skills which are useful in multiple tasks. Finally, the Becker and Murphy model of specialization makes the clear predicition that no two workers will perform the same task within the same firm, and this is seemingly at odds with the observation that many firms report having numerous employees who are engaged in the same occupation.

Nonetheless, as a stylized model of the gains from specialization, this model is attractive in that it generates clear empirical predictions solely from differences in firms' production technologies and without imposing particular notions of comparative advantage. Accordingly, I present a modified version of this model in the next Section, in order to provide key testable predictions regarding the relationship between skill specialization and wages.

2.3 Theoretical Model

Below, I briefly present a slightly modified version of the original Becker and Murphy model. Relative to the original, I relax the assumption of Leontief production by firms to a more general CES formulation. I also explicitly assume that labor markets are imperfectly competitive in the style of Card et al. (2017). In general, the assumption of imperfectly competitive labor markets is necessary ensure that the compensation of otherwise-identical individuals varies across firms in a manner that is correlated with the degree of specialization. More specifically, the formulation shown in this chapter is able to generate the type of linearly separable log wage empirical specification used in Section 2.6.4. All other assumptions and predictions are as in the original paper, including the initial assumption that all employees are identical. Later, in

Subsection 2.3.2, I relax the assumption of identical workers to model the introduction of heterogeneous ability.

2.3.1 Basic Model Setup

A single final consumption is produced in a competitive product market and sold at a unit price P. Each firm produces this final consumption good by applying a firm-specific technology parameter A_f to a unit measure of intermediate goods (also referred to as tasks), indexed by s, with a constant elasticity of substitution among intermediate goods. These intermediate goods are produced by contracting with individual laborers, referred to as employees. The firm is also subject to a coordination cost $C(n)$ that is increasing and convex in n, the number of workers employed by the firm. Then, each firm's production function is:

$$Y_f = A_f \left(\int_0^1 x_{fs}^{\eta} ds \right)^{\frac{1}{\eta}} - C(n) \tag{2.1}$$

where $\frac{1}{1-\eta}$ represents the elasticity of substitution between any two intermediate goods inputs in final goods production, and n is the number of employees of the firm. For readability, I will now omit the subscripts f, as it suffices to focus on a single firm in this model. And, I will rule out the case of perfect substitutes by assuming that $\eta < 1$.[2]

Each employee (i.e. each intermediate goods producer), has one unit of time which they devote solely to labor market activities associated with particular task inputs. This unit of time is spent either working (T_w) or investing time to learn the productive

[2]Because of increasing returns to specialization, when all inputs are perfect substitutes, then any equilibrium in which workers perform a non-zero measure of tasks or in which firms use a non-zero measure of task inputs is suboptimal. To avoid complications, I will restrict attention to circumstances where η is "low enough" relative to the extent of gains from specialization that the representative firm always finds it optimal to use the full measure of task inputs in production. That is, I assume that $\eta < \frac{1}{1+\theta}$, where $\theta > 0$ is as indicated in (2.6).

skills specifically associated with tasks performed in production (T_h). There is a one-to-one mapping between tasks performed and observable skills acquired. And, since it cannot be optimal for an employee to spend time learning skills that they will not use in production, I refer to the skills acquired as "task-specific human capital."[3]

Then, for each intermediate good s and for each employee i:

$$T_i(s) = T_{ih}(s) + T_{iw}(s) \tag{2.2}$$

$$\sum_s T_i(s) = 1 \tag{2.3}$$

As in Becker and Murphy, I will for now assume that all workers are ex ante identical. That is, prior to workers' investment decisions, all workers are equally adept at performing all tasks. Individuals' productivity differences across different inputs arise only because these individuals rationally choose to spend time learning the skills corresponding only to the production tasks that they perform. Productive output of worker i in intermediate good s depends on the worker's efficiency $E(s)$ and on the time spent working in production of good s. In turn, the worker's efficiency is a function of both exogneously determined general human capital H and of time spent investing in task-specific human capital $H_T(s)$. That is:

$$Y_i(s) = E_i(s)\, T_{iw}(s) \tag{2.4}$$

$$E_i(s) = dH^{\gamma} H_{iT}(s) \tag{2.5}$$

$$H_{iT}(s) = T_{ih}^{\theta}(s) \tag{2.6}$$

[3]Although the setup of this model appears to imply that task-specific human capital is acquired on the job, it need not be the case. For example, if individuals learn task-specific skills through higher education, attendance at trade school, etc., then they forego the opportunity to produce during the time that they are on the job. The model as written does not permit discounting based on the fact that skills must generally be acquired before they are used to produce, but this extension can be added.

Here, the parameter $\gamma > 0$ determines the value of general human capital investment, while the parameter $\theta > 0$ determines how large the productivity gains are from investing time in task-specific human capital, and d is a general productivity shifter. Importantly, the assumption that $\theta > 0$ implies that there are increasing returns to time spent on individual production tasks as long as the employee spends some time in human capital investment. Since individuals have a finite allotment of time that must be split between skill investment and production, these increasing returns are bounded by individual workers' limited time. All individuals will allocate a portion of their time to task-specific human capital acquisition in this formulation, because otherwise they can produce nothing.[4] And, finally, I assume that each employee's output is observable and contractible, so that the managers of firms can always enforce an allocation of workers' time across the tasks that they are assigned in order to maximize the employee's overall output.

Proposition 1. *For any given task, if an employee chooses to invest in task-specific human capital, then they spend a fraction* $\frac{\theta}{1+\theta}$ *of their time in that task on human capital investment, and a fraction* $\frac{1}{1+\theta}$ *of their time in production.*

Proof. Substitute (2.2) and (2.5) into (2.4) and take the first order condition with respect to $T_{ih}(s)$ to show that $T_{ih}(s) = \frac{\theta}{1+\theta}T_i(s)$ and therefore that $T_{iw}(s) = \frac{1}{1+\theta}T_i(s)$. □

This first proposition implies a constant ratio of task-specific human capital investment to task-specific work that depends only on the productivity of those task-specific human capital investments. It also implies, by substituting these ratios back into (2.5) and (2.4), that conditional on choosing to invest in task-specific human capital, the marginal product of the employee's time spent on any task s will be:

[4] This assumption can be relaxed. For example, if we redefine $E_i(s) = dH^{\gamma}(1 + H_{iT}(s))$, then employees may choose to produce in autarky without specializing unless θ and A are high enough and $C(n^*)$ is low enough at equilibrium firm employment n^* to allow i to produce at least dH^{γ} by optimally specializing. However, for simplicity, I will use the same formulation as Becker and Murphy here.

$$Y_i(s) = dH^{\gamma}\theta^{\theta}\left(1+\theta\right)^{-(\theta+1)}T_i(s)^{\theta+1} \tag{2.7}$$

$$\frac{\partial Y(s)}{\partial T_i} = dH^{\gamma}\left(\frac{\theta}{1+\theta}\right)^{\theta}T_i\left(s\right)^{\theta} \tag{2.8}$$

Taking the derivative of (2.8), it is also straightforward to show that as long as employees are willing to engage in task-specific human capital investment, each individual has increasing returns to time spent on each task. This follows directly from the assumption that $\theta > 0$.

Proposition 2. *Any employee who is engaged in task-specific human capital investment will be the only employee engaged in the production of that intermediate good for the firm.*

Proof. See Appendix B.1. □

Proposition 3. *If $\eta < \frac{1}{1+\theta}$, then for each intermediate good that they produce, employees who invest in task-specific human capital will devote an equal fraction of time. That is, $T_i\left(s\right) = \frac{1}{|s_{if}|}$ for all $s \in \{s_{if}\}$, where $\{s_{if}\}$ is the set of intermediate goods that employee i produces at firm f and $|s_{if}|$ is its measure. And, since all employees are identical, $\frac{1}{\left|s_{if}\right|} = n_f^*$ for all employees within each firm.*

Proof. See Appendix B.2. □

Proposition 2 implies that there is no overlap among the tasks performed by any two workers. So, the level of output of intermediate production task s that is produced by the firm, x_s, is equal to the output of that task for the worker who produces it, $Y_i(s)$. Proposition 3 implies further that an employee who invests in task-specific human capital for any task will always invest equally across all of the tasks that they perform, since it is always preferable to reallocate time from tasks with lower marginal product of time to tasks with a higher marginal product of time as long as tasks are not strong gross substitutes. Together, these propositions imply that $x_s = x_i$ for all tasks s performed by worker i, and therefore that the total output of each worker i is simply $x_i \cdot |s_i|$.

The combined effect of these two propositions is that each employee's output can be characterized solely as a function of the shared parameters θ, d, H, and γ, as well as the optimally chosen firm size n^*, a function of both A_f and of $C(n)$. That is, substituting n_f^* for $T_i(s)$, it is the case that:

$$Y_{if} = d\theta^\theta (1+\theta)^{-(1+\theta)} H^\gamma \left(n_f^*\right)^\theta \tag{2.9}$$

Plugging (2.9) into (2.1) yields the following equation for firm output:

$$Y_f = A_f d\theta^\theta (1+\theta)^{-(1+\theta)} H^\gamma n^{1+\theta} - C(n) \tag{2.10}$$

While this completes the setup of the basic Becker and Murphy model, and it is sufficient to show that more specialized workers will be more productive, it is not quite sufficient to demonstrate that wages will vary with the extent of specialization.[5] For this, we need labor markets to be imperfect, so that wages vary across firms with marginal product, instead of (as might otherwise be assumed) ex-ante identical workers earning identical wages across firms regardless of specialization because of Bertrand wage competition.

So, assume that labor markets are imperfectly competitive in the manner of Card et al. (2017). In this model, each employee receives utility from the combination of wages and amenities that each firm provides, with heterogenous preferences over firm-specific amenities. That is, each individual's indirect utility takes the form:

$$u_{if} = \beta \log(w_{if} - b) + a_f + \epsilon_{if} \tag{2.11}$$

[5]Becker and Murphy themselves do not consider the wage implications of their model, though it is straightforward to show that with perfectly competitive labor markets or Bertrand wage competition, workers' wages do not vary with their degree of specialization. The model shown in this paper nests this perfectly competitive solution in the limit.

where $\beta > 0$, b is a common outside option, a_f is the utility from exogenous firm-specific amenities that are valued equally by all individuals and are costless to produce, and ϵ_{if} is an individual preference drawn from an independent Type I extreme value distribution and is unobservable to the firm. Firms post wages in order to attract employees, and so each firm f chooses its wage to maximize profits by solving:

$$\max_{w_f} P \cdot Y_f\left(n\left(w_f\right)\right) - w_f n\left(w_f\right) \tag{2.12}$$

where the observability of individual workers' output, the lack of utility or disutility in particular task assignments, and the constraint that $n\left(w_f\right) = \frac{1}{|s_f|}$ ensure that Y_f is as in (2.10).[6]

Since workers choose a firm freely based both on posted wages as well as their amenity preferences, and because of the functional form assumption on ϵ_{if}, the probability in equilibrium that an individual chooses a particular firm is represented by a logit choice probability (McFadden, 1973). With a large number of firms, this is closely approximated by an exponential probability. So, as in Card et al. (2017), each firm faces an upwardly sloping labor supply curve that depends on the offered wage:

$$\ln n\left(w_f\right) = \ln\left(\mathcal{N}\lambda\right) + \beta \ln\left(w_f - b\right) + a_f \tag{2.13}$$

where $\mathcal{N}$ is the total number of workers in the labor market and λ is a constant that is common to the supply curve of all firms. By calculating the elasticity of this labor supply with respect to wage and plugging it into the first order condition of (2.12) with respect to wages, it can be shown that the equilibrium wage offered by each firm

[6] As in Card et al. (2017), I assume here that firms choose wages directly, but they choose employment levels only indirectly through posted wages. Since there is no uncertainty in this model, such a formulation is consistent with a standard cost minimization problem in which firms offer the lowest wage that is needed to achieve a desired level of output.

is discounted from its employees' marginal revenue product according to the function:

$$w_{if} = \frac{1}{1+\beta}b + \frac{\beta}{1+\beta}A_f P\left(d\left(\frac{\theta}{1+\theta}\right)^{\theta} H^{\gamma} n^{\theta} - C'\left(n_f^*\right)\right) \tag{2.14}$$

Notice that if employees care only about their wage (i.e. as $\beta \to \infty$), the wage that is paid to each worker approaches their marginal revenue product. Alternatively, if employees care only about their amenities, then the wage that is paid to each worker approaches their outside option b. And, since the degree of skill specialization among workers in a firm is positively related to its employees' marginal product, greater specialization is predicted to be associated with higher wages in equilibrium.

2.3.2 Heterogeneous Ability

Up to this point, we have assumed that all workers are identical. This implies that there is no absolute or comparative advantage in this model whatsoever. While this assumption is attractive for its simplicity, it is also quite unrealistic. In particular, much of the human capital literature is based on the idea that individuals are heterogeneous, either in their exogenously determined productivity, their cost of education, or both.

In this subsection, I relax the assumption of identical workers by allowing for different workers to have heterogeneous ability in learning task-specific skills. That is, some workers can learn more in the same amount of time than others, giving them absolute advantage in all tasks, but no comparative advantage in any particular tasks. For tractability, assume that there are only two types of individuals, $i \in \{H, L\}$, and that type H workers learn more effectively than type L workers, so that $\theta_H > \theta_L$. Since the proof of Proposition 1 does not depend on other workers, it still follows that if worker i learns task-specific human capital for a given task, they allocate $\frac{\theta_i}{1+\theta_i}$ of their task-specific time to investment and $\frac{1}{1+\theta_i}$ of their time to work. Similarly, Propositions 2 and 3 also hold, so workers will always divide their time equally among the tasks that they perform, and specialized workers will be non-overlapping in the tasks that they perform. This implies that the unit interval of production tasks can be divided into two subsets of tasks, performed by type H and L workers, respectively, with an equal level of output of each task within each subset, and with total output of each worker of each type as $Y_i = x_i |s_i|$.

So, letting $\overline{s}$ be the interval of tasks assigned to the set of type L workers employed, the firm's production function can simply be written as:

$$Y = A\left(\overline{s}x_L^{\eta} + (1-\overline{s})x_H^{\eta}\right)^{\frac{1}{\eta}} - C\left(n\right) \tag{2.15}$$

Notice, however, that unlike in most standard CES formulations, the task share parameter $\overline{s}$ is not fixed. Rather, it is chosen by the firm, along with individual worker shares of each type $|s_i|$ and wages of each worker type, in order to maximize profits, subject to the constraint that $\sum_i n_i |s_i| = 1$. That is, the firm now solves:

$$\max_{\overline{s}, w_H, w_L} PY\left(\overline{s}, n_H\left(w_H\right), n_L\left(w_L\right)\right) - w_H \cdot n_H\left(w_H\right) - w_L \cdot n_L\left(w_L\right) \tag{2.16}$$

where Y is as in (2.15). As in the baseline model, I assume that each individual's indirect utility takes the form of (2.11); this implies that each firm faces an upwardly sloping labor supply curve for each type of worker as in (2.13). The labor supply to the firm is a functiion of posted wages, but it does not depend on task assignment, and therefore it is not a function of $\overline{s}$.

Before proceeding with the maximization problem, notice that modifying (2.9) to reflect the fact that $n_L^* = \frac{\overline{s}}{|s_L|}$ yields the following equation for the output of each employee of each type:

$$Y_i = dH^{\gamma}\theta_i^{\theta_i}\left(1+\theta_i\right)^{-(1+\theta_i)}|s_i|^{-\theta_i} \tag{2.17}$$

Then, taking the ratio Y_H/Y_L from (2.17), letting $\theta_L = \theta$ and $\theta_H = \kappa\theta$ for some $\kappa > 1$, and substituting in that $Y_i = x_i\,|s_i|$ to solve for the ratio of intermediate good output levels $\frac{x_H}{x_L}$ yields:

$$\frac{x_H}{x_L} = \phi\frac{|s_L|^{\theta+1}}{|s_H|^{\kappa\theta+1}} \tag{2.18}$$

where $\phi = \kappa^{\kappa\theta}\theta^{\theta(\kappa-1)}\frac{(1+\theta)^{1+\theta}}{(1+\kappa\theta)^{1+\kappa\theta}}$ is a constant function of parameters θ and κ that is less than one and decreasing in both its arguments.

Since x_i is strictly increasing in the time that a worker of type i spends developing task-specific human capital, equation (2.18) indicates the implicit ratio between the breadth and depth of each type's skill development, which arises because of the limited time that each type of worker faces, and because of the difference in the productivity of each type.[7]

To proceed with the maximization problem, I restrict my attention to interior cases in which the firm hires one or more workers of each type and chooses $\overline{s} \in (0,1)$.[8] This implies that the first order conditions of (2.16) must hold simultaneously. Taking the first derivative of (2.16) with respect to $\overline{s}$ implies that $Y_{\overline{s}} = 0$, where $Y_{\overline{s}}$ is the first derivative of output with respect to $\overline{s}$. That is, as long as $\frac{\partial n_i}{\partial \overline{s}} = 0$ for $i \in \{H, L\}$, the firm chooses $\overline{s}$ to maximize its output conditional on the number of employees it hires. So, taking the derivative of (2.15) with respect to $\overline{s}$, plugging in (2.17) and its derivative and simplifying yields the following equilibrium relationship that is solely a function of exogenous parameters and that must hold in any equilibrium where the

[7]In particular, this equation also makes it clear that type H workers are always more productive than type L workers. To see this, set $|s_L| = |s_H| = |s|$. Then $x_H/x_L = \phi\,|s|^{\theta(1-\kappa)}$, which is greater than 1 as long as $\kappa > 1$ and $\theta > 0$.

[8] I do not rule out the existence of non-interior cases in which the firm employs only type H or type L workers, but the empirical conclusions of this paper with respect to intra-firm wage differentials are of interest only when firms employ workers of both types. Future revisions of this paper will include a charactertization of interior and non-interior solutions to the firm's maximization problem with two types of workers.

firm employs and assigns tasks to both types:

$$\frac{x_H}{x_L} = \left(\frac{\eta(\theta+1)-1}{\eta(\kappa\theta+1)-1}\right)^{\frac{1}{\eta}} \tag{2.19}$$

For $\eta < \frac{1}{\kappa\theta+1}$, this optimal ratio is strictly increasing in η and greater than 1, with $\lim_{\eta\to-\infty} \frac{x_H}{x_L} = 1$. In other words, given any feasible productivity parameters κ and θ, there exists a unique optimal ratio in the depth of skill that is developed by the two types. When production tasks are perfect complements, the depth of each worker's task-specific human capital acquisition is identical. On the other hand, if production tasks are more substitutable, then the depth of high-type workers' task-specific human capital acquisition will be greater relative to that of low-type workers within the firm. Intuitively, the equilibrium ratio of skill depth does not depend on the firm's levels of employment or on relative wages precisely because the firm has the ability to adjust the scope of tasks that are performed by each of type of worker by varying $\overline{s}$. That is, the firm is able to achieve exactly the ratios of skill depth and breadth that maximize output given the trade-off between imperfectly substitutable inputs and gains from skill specialization. Notice, however, that if inputs are too substitutable (i.e. if $\eta > \frac{1}{\kappa\theta+1}$), this equilibrium relationship breaks down.

Furthermore, there is a unique relationship between the productivity parameters, the elasticity of substitution, and the optimal breadth of tasks performed by workers of each type. This can be shown by equating the right-hand rides of (2.18) and (2.19). Doing this, and letting $|s_H| = \gamma |s_L|$ so that the value of γ represents the task and skill breadth of high-ability workers relative to that of low-ability workers, it can be shown with some algebra that:

$$\gamma = \left(|s_L|^{\theta(1-\kappa)}\, \phi \left(\frac{\eta(\theta+1)-1}{\eta(\kappa\theta+1)-1}\right)^{-\frac{1}{\eta}}\right)^{\frac{1}{\kappa\theta+1}} \tag{2.20}$$

Depending on the various parameter values and on the optimal value of $|s_L|$, the value of γ may be greater than or less than one, indicating that type H workers perform either a narrower or a broader set of tasks. However, it is straightforward to verify that if tasks are even moderately gross complements, that unless $|s_L|$ is very large (that is, unless coordination costs are high enough that the optimal firm size is very small) or θ and κ are very small, γ will be greater than 1.[9] And, from (2.19), the difference in skill depth between type H and type L workers declines with η as well. In other words, the more complemetary production tasks are, the greater the extent to which high-ability workers develop a broader, less specialized skillset than low-ability workers within the same firm. In the extreme case of perfect complements, any equilibrium where $\gamma > 1$ is sufficient to ensure that high-ability workers hold less specialized skills than low-ability workers under any sensible metric.

Figure 2·1 provides a stylized example of a firm that employs both high and low ability workers in production. Each rectangle in the interval $[0, \bar{s})$ represents the task breadth and output of a low-ability worker, while each rectangle in the interval $[\bar{s}, 1]$ represents the task breadth and output of a high-ability worker. These depths and breadths are also directly reflective of the extent of each worker's task-specific human capital. Although the heights of each set of rectangles are a function of the optimal firm size, as shown in (2.19), the ratio between the heights of each set of rectangles is a function solely of underlying parameters and the elasticity of substitution between production tasks, and this ratio is achieved by the firm by adjusting the value of $\bar{s}$, which dictates the proportion of tasks that is performed by each type of worker.

[9] While (2.20) is a nonlinear function of parameters κ, θ, and η and of equilibrium employment of type L workers which determines $|s_L|$, it is possible to verify numerically that both $\gamma > 1$ and $\gamma < 1$ are possible. For example, with $\eta = -1$ (i.e. an elasticity of substitution across inputs of $\frac{1}{2}$), $\theta = 1$ and $\kappa = 2$ it is straightforward to verify that $\gamma < 1$ only if $|s_L| > \frac{4}{9}$. So, given these parameter values, any firm that would hire three or more workers if only type L workers were available to it will choose to assign a broader set of tasks to any type H workers that it hires. In contrast, with $\eta = -9$ (i.e. an elasticity of substitution across inputs of $\frac{1}{10}$), even a two-person firm that hires one worker of each type will choose to assign tasks such that $\gamma > 1$ when $\theta = 1$ and $\kappa = 2$, and any firm that would hire three or more low-type workers (i.e. with $|s_L| \leq \frac{1}{3}$) will assign tasks such that $\gamma > 1$ for any $\kappa > 1$ when $\theta = 1$.

Figure 2·1: Stylized Example of a Firm Employing High and Low Ability Workers

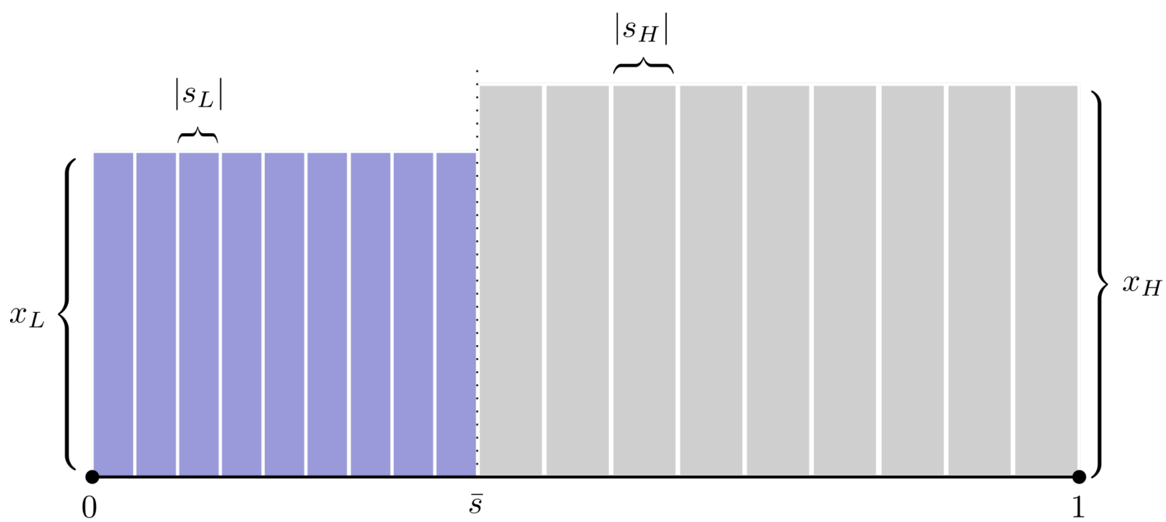

2.3.3 Summary of Empirical Predictions

The Becker and Murphy model provides a rich framework for considering the sources of wage premia associated with specialization. Since this chapter focuses on the potential wage premia associated with specialization, I will focus on two key empirical predictions:

1. Across firms, those firms with access to more productive technology (or lower coordination costs) will employ more specialized workers. If workers' wages vary with their productivity because of imperfectly competitive labor markets, then we will observe that more specialized workers earn higher wages than less specialized workers. The observed specialization premium across firms will also be observable as a firm size wage premium, because the degree of within-firm specialization is limited by the extent of the firm.

2. Within firms, high ability workers may be more or less specialized than low ability workers. In particular, the extent of *relative* specialization or non-specialization depends on the extent to which individual task inputs are com-

plementary in production. Again, in the presence of imperfect labor markets such that workers' wages vary with their productivity, specialization may be associated with either a positive or a negative wage premium relative to other employees within the same firm.

Finally, a quick note on general human capital as it pertains to this model. Although general human capital is treated as an exogenous productivity shifter here, its existence (and observability) is actually very important for empirical estimation. To see why, consider what would be measured in the data if general human capital were unobserved. Although there would be variation in the observable concentration of skills (i.e. skill specialization), and although this would still be a reflection of underlying productivity differences across firms, the degree of skill specialization would also be perfectly positively collinear with the overall observed depth of skill.

Of course, in practice, we expect (and observe) that there is distinct variation between the average level of an individual's skill usage and the concentration of their skills. The extent of an individual's general human capital might vary because of their institutional environment, because of the role of ability in acquiring general human capital, or for other reasons. These differences also influence wages, and in In fact, in the O*NET data used in this chapter, the mean level of skill reported for an occupation tends to be negatively correlated with its concentration, not positively. However, for these reasons, in all empirical specifications presented in this chapter, I include controls for the mean level of skill reported alongside measures of the concentration of those skills.

In the sections that follow, I take these key empirical predictions to the data, considering both inter-firm and intra-firm wage patterns as they relate to the observable degree of specialization.

2.4 Empirical Specification

2.4.1 Wage Premia

The basic model described in Section 2.3 provides two key testable implications, which are the subject of the empirical analysis of this chapter:

1. Workers who are more specialized in their skillsets will sort to higher-wage firms and will earn higher wages across firms.

2. Workers in who are more specialized in their skillsets will not necessarily earn higher wages within firms, as high-ability individuals may actually be less specialized than low-ability individuals within firms.

In order to look at wage premia, we can look directly at the association between wages and occupation level measures of skill and specialization. Accordingly, my baseline empirical specification is:

$$\ln(w_{iot}) = \alpha + \beta_0\mu_o + \beta_1 G_o + \gamma X_{iot} + \epsilon_{it} \tag{2.21}$$

where the values of μ and G are standardized measures of the mean skill percentile and Gini of skills of each individual's occupation o and X_{it} are standard individual-level covariates such as age, education, and experience. To examine specifically within-establishment wage patterns, I also include establishment-level fixed effects in certain specifications. The estimated value of β_1 reflects the predicted wage premium associated with a one standard deviation increase in the degree of specialization of occupational skills, holding constant both the level of skill and differences in wages arising from observable compositional differences in the individuals choosing each occupation.

Two natural concerns arise from a specification such as this. The first is that with a time-invariant occupation-level measure of skill content, there may be correlated occupation-level considerations that might lead specific occupations to have higher or lower wages than would be predicted by their skill content and specialization. In particular, hedonic considerations may lead some occupations to be paid more than others, as may policies such as licensing or unionization that restrict labor supply to certain occupations or influence workers' bargaining power. While the nature of the O*NET data limits what can be done to remedy this concern, in most specifications I include broader occupational group fixed effects, so that workers' wages are compared only to those in similar occupations.

A second, more challenging concern arises from the fact that these predictions are arrived at via a highly stylized model that is intentionally agnostic to the actual content of workers' skills and tasks. It suggests that workers lack sources of comparative advantage, and that all production activities are simply represented as a slice of a unit interval. While these abstractions allow the model to provide useful insights, it is not *ex ante* obvious that an empiricial specification that measures skill breadth or depth across all skills is appropriate.

In particular, the model suggests that wage premia associated with specialization are observable because of the process by which firms assign different workers to different production tasks in a single labor market. Yet in the real world, different occupations may exist in very different labor markets, and skills that are valued in some of these markets may be of little or no use in others. For example, a measure of specialization that includes skills that are clearly not in the scope of white collar employment will be an imprecise measure of those occupations' specialization. The same is true for blue-collar skills among white-collar workers.

To address this particular concern, in most regression specifications, I replace the scalars μ_o and G_o with vectors that measure mean skill and the concentration of skills within two specific subsets of skills observed in O*NET, which I term "Production" skills and "Cognitive" skills. These skill subsets are determined via an agglomerative hierarchical clustering procedure, and their definitions map closely to the skills used most heavily in white-collar versus blue-collar occupations. Further details of the procedure by which these skills groups are determined are provided in Section 2.5.4.

2.4.2 Worker Sorting

The second key set of predictions of the model in Section 2.3 relate to the sorting of workers. In particular, while substitutability considerations may lead firms to assign higher-ability workers to less specialized occupations within firms, we should still expect that highly specialized workers will be consistently sorted into more productive firms. This prediction is a direct implication of the notion that specialization is valuable. And, if labor markets are imperfect, this implies that highly specialized workers will also be sorted into high-wage firms as well.

To investigate this further, I turn to the method of multidimensional fixed effects initially proposed by (Abowd, Kramarz, and Margolis, 1999), and commonly known as the AKM method. This method has been widely adopted in the recent literature on sorting between workers and firms, in particular on the contribution of sorting of workers to increased inequality over time. For my analysis, I will not consider changes over time, because of the selection issues described in Section 2.5.2. Instead, I will construct a simple two-step procedure to establish the extent to which sorting to high-wage firms is associated with different, potentially correlated characteristics.

The method begins with the assumption that log wages can be decomposed into additive worker and establishment fixed effects, once life-cycle considerations are flexibly controlled for.[10] The baseline model of skill specialization described in Section

[10] The standard methodology in recent practice controls for the age profile of earnings using higher-order polynomials in age, interacted with indicators for educational attainment groups. For a more complete overview of the methodology used, see (Card et al., 2017). In my analysis, I adopt the same first-stage methodology characterized in that paper, and used in Card, Heining, and Kline (2013).

2.3 meets this criterion, because firms offer wages that are a fraction of their marginal product of labor in response to taste-based heterogeneity in firm-specific amenities. So, I can run regressions of the form:

$$\ln(w_{it}) = \alpha_0 + \delta_i + \psi_{\mathbb{J}(i,t)} + \gamma_0 X_{it} + \epsilon_{it} \tag{2.22}$$

The function $\mathbb{J}(i,t)$ indicates the identity of the establishment at which worker i is employed at time t. The values of X_{it} include year indicators as well as higher-order polynomial terms in age and potential labor force experience. Under the assumptions specified above, both establishment and worker fixed effects are identified using this method. In particular, the model implies that the establishment wage premium identified by each establishment fixed effect is an increasing function of the associated firm's level of technology.

Then, in the second step, I run regressions of the form:

$$\psi_{\mathbb{J}(i,t)} = \alpha_1 + \beta_0 \mu_o + \beta_1 G_o + \gamma_1 \tilde{X}_{it} + \epsilon_{it} \tag{2.23}$$

In this step, the values of $\tilde{X}_i t$ include similar individual-level covariates as in my baseline wage regressions. By regressing the identified establishment fixed effect against individual level covariates, I can test for worker sorting along multiple dimensions simultaneously, and can therefore determine the extent to which worker sorting patterns are related to the observed degree of skill specialization, while also controlling for sorting that is driven by observable education or other correlated factors.[11]As noted in the previous subsection, μ_o and G_o can be either scalar values, or a vector of values for the mean skill and Gini of skills within subsets of the overall set of ob-

[11]Another method of examining the degree of worker sorting is the method of Kremer and Maskin (1996). While that method provides certain advantages, such as the ability to examine changes in the degree of sorting across time, it is limited to examining a single dimension of sorting. In contrast, the two-step AKM method allows me to consider sorting along multiple dimensions simultaneously, to ensure that the sorting I observe is not driven by other correlated factors.

servable skills. If, as predicted, high-wage firms employ more specialized workers, we should expect that $\beta_1 > 0$.

A key feature of the AKM method is that establishment and worker fixed effects are identified simultaneously using the connected set of individuals and establishments in the data. Establishments that only ever employ a single worker are dropped, as are establishments who never have at least one employee transition to another employer. Identification of establishment fixed effects comes in particular from individuals who are observed repeatedly in the sample at different employers at different points in time. However, since I am able to incorporate approximately 95% of observations from the analysis of wage premia using this method, these results can still be interpreted to measure sorting patterns across the aggregate labor market.

2.4.3 Occupation Switchers

A final prediction of the theoretical model relates specifically to the source of any wage premia associated with specialization. That is, the degree to which workers are able to specialize is limited by the extent to which firms can accommodate specialization in production, and this is in turn a function of those firms' production technologies. Similarly, when workers of different ability levels take more or less specialized employment within firms, their degree of specialization is a reflection of their underlying productive capacity. In other words, the model suggests that *any* wage premium associated with specialization is either a reflection of an individual characteristic or a firm characteristic, not a function of the occupation itself. Once these time-invariant individual and firm characteristics are accounted for, there should be no discernable wage patterns associated with the degree of workers' specialization at any point in time.

To consider this prediction, I seek to look specifically at individuals who are observed repeatedly in different occupations with different skill usage. So, instead of the two-step AKM procedure described above, I run a single-step regression that includes individual and establishment fixed effects along with the measures of mean skill and the concentration of skills as before:

$$\ln(w_{iot}) = \alpha + \beta_0\mu_o + \beta_1 G_o + \delta_i + \psi_{\mathbb{J}(i,t)} + \beta X_{it} + \epsilon_{it} \tag{2.24}$$

In this specification, the values of β_0 and β_1 represent the average net increase or decrease in wage associated with a one standard derivation change in the level of skill usage and the degree of skill specialization, respectively. And, the prediction of the model is that each of these should be zero, because all relevant determinants of employees' wages are still captured by the combination of individual and establishment fixed effects.

2.5 Data Sources and Measurement

In this section, I characterize the various data sources and data methods used for my analysis. These include descriptions of my primary data sources, the O*NET database and the Brazilian RAIS matched employer-employee dataset, as well as the U.S. Current Population Survey data used for robustness checks and the hierarchical clustering method used to classify occupational skills.

2.5.1 O*NET

The O*NET database is a standard source of information on occupational content from the United States that has been used throughout recent literature on the skill content of occupations. It replaces the Dictionary of Occupational Titles, which also has a long history of application in the literature on occupations.

For each of 963 SOC-classified occupations, responses are collected from a random sample of incumbent workers in each occupation, which are selected from a random sample of businesses. The O*NET database is regularly updated as individual occupations are periodically re-surveyed over time; the analyses presented in this chapter use version 21.0 of the O*NET database, released in August 2016 (*O*NET Resource Center - O*NET 21.0 Production Database*).

For each of 35 skills, respondents are asked "What level of [skill name] is needed to perform your current job," and are asked to respond on a 1-7 categorical scale.[12] I henceforth refer to this question as the "Level Question." If respondents have previously responded that this skill is not important to the performance of their current job, then they are directed not to respond to the Level Question, and their response is coded as 0. Publicly available data report the mean response to each Level Question for each occupation, but not the individual workers' responses.[13]

Figure 2·2 shows the distribution of raw scores for each of the 35 skills defined in the O*NET database. The extent of each box marks the 25th and 75th percentile of occupation-level mean responses to the Level question; the median reported occupation-level mean is also marked. It is clear from this figure that both median scores and the distribution of scores differ considerably among different skills in the data. While most occupations report relatively low raw scores in such skills as Programming, in other skills such as Reading Comprehension, a majority of occupations report comparatively high scores. Although these distributions may appear to be

[12]The publicly available O*NET dataset includes mean response data on several dimensions of occupational heterogeneity, including Abilities, Tasks, and Knowledge, as well as Skills. In practice, prior research has often used data from several of these questionnaires, aggregating responses from specific questions that are argued to reflect specific types of skill content (Autor, Levy, and Murnane, 2003; Autor, 2013). Because this chapter seeks to consider the overall breadth of a workers skills, I have opted to use only responses to the O*NET Skill questionnaire, and to include all 35 skills in my measures.

[13]To encourage consistent response patterns, the O*NET Skill questionnaire frames certain categorical choices using descriptions of what that level represents. For example, a 2 in Reading Comprehension may represent a level of skill needed to "Read step-by-step instructions for completing a form," while a response of 6 may represent a level needed to "Read a scientific journal article describing surgical procedures." Because this scale has no cardinal content, I convert all raw score data to percentile form as described below.

Figure 2·2: Boxplot of Raw Score Data for Each of 35 Skills in O*NET Database

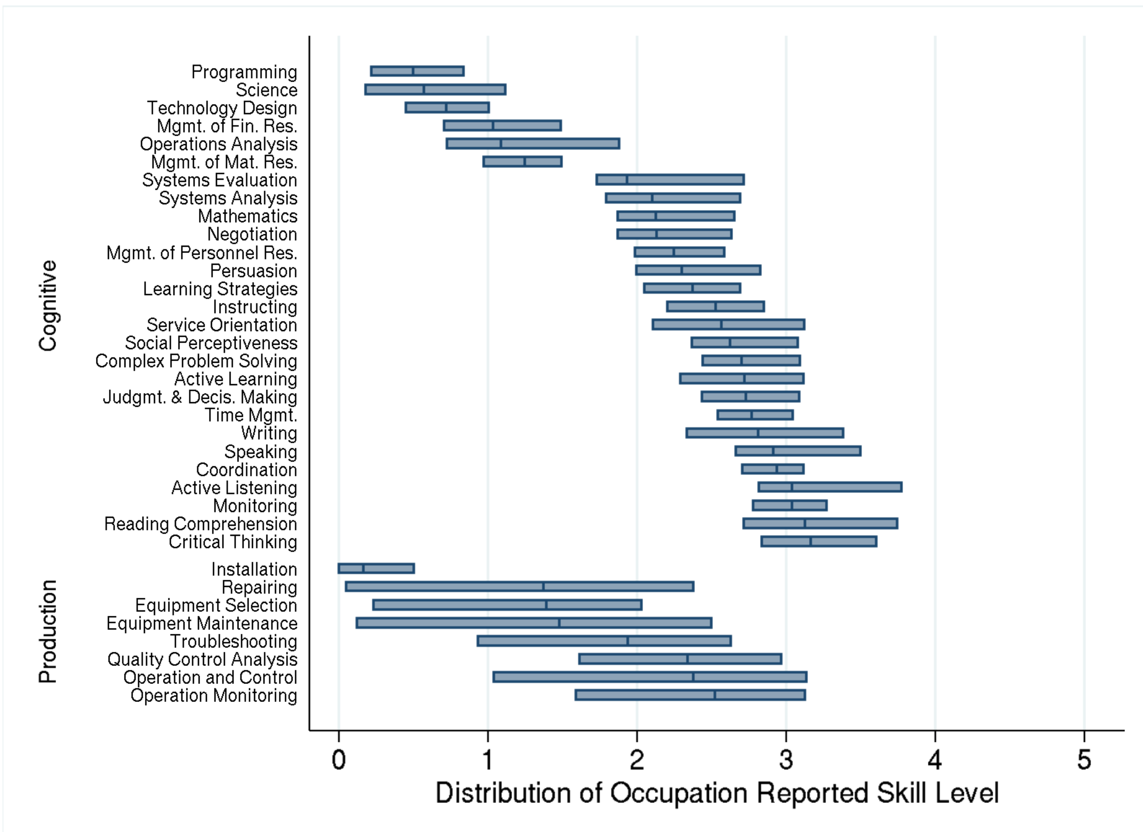

Notes: From O*NET, matched to Brazilian CBO occupational codes. Boxes represent the 25th, 50th, and 75th percentiles of occupation-level responses.

suggestive of the relative scarcity or abundance of different skills, it is important to recognize that raw responses to the Level question have no inherent cardinal meaning.

To address this issue, I adopt a variant of the strategy suggested by Autor (2013). I convert all Level question raw scores to percentiles of the overall occupational distribution. This conversion ensures that my measures of occupational skill measure skill relative to the set of other occupations, rather than in absolute terms. Implicitly, this also weights each skill equally in all aggregate skill calculations, rather than weighting most highly the skills for which the mean raw response is greatest.

Figure 2·3: Reported Skill Percentiles of Economists (SOC 19-3011.00)

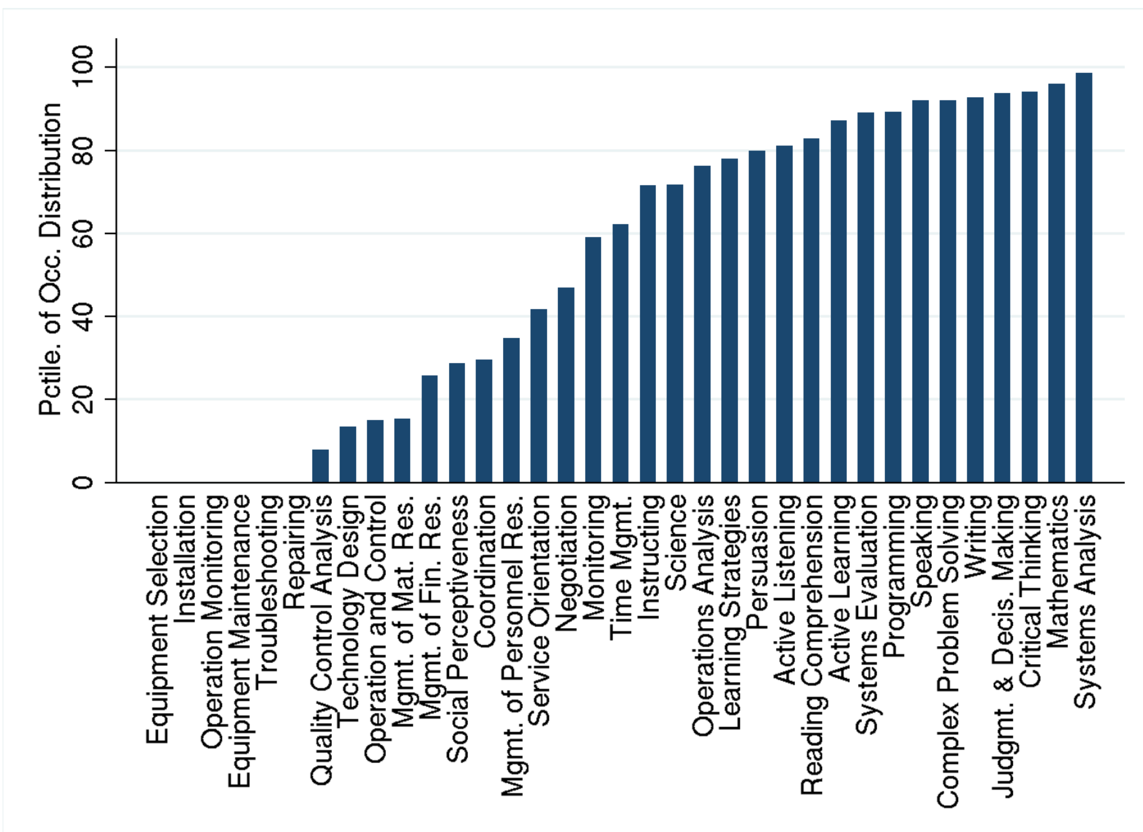

Notes: From O*NET. Percentiles are reported relative to all other occupations in the database.

As an example, Figure 2·3 graphs the average skill percentiles reported by Economists

(SOC 19-3011.00). As shown, Economists are not only relatively skilled, but their skills are relatively broad; they report skill levels that are at or above the 80th percentile of occupations sampled in 13 of 35 skills reported.[14]

Figure 2·4: Stylized Example Of Gini Coefficient for Two Occupations

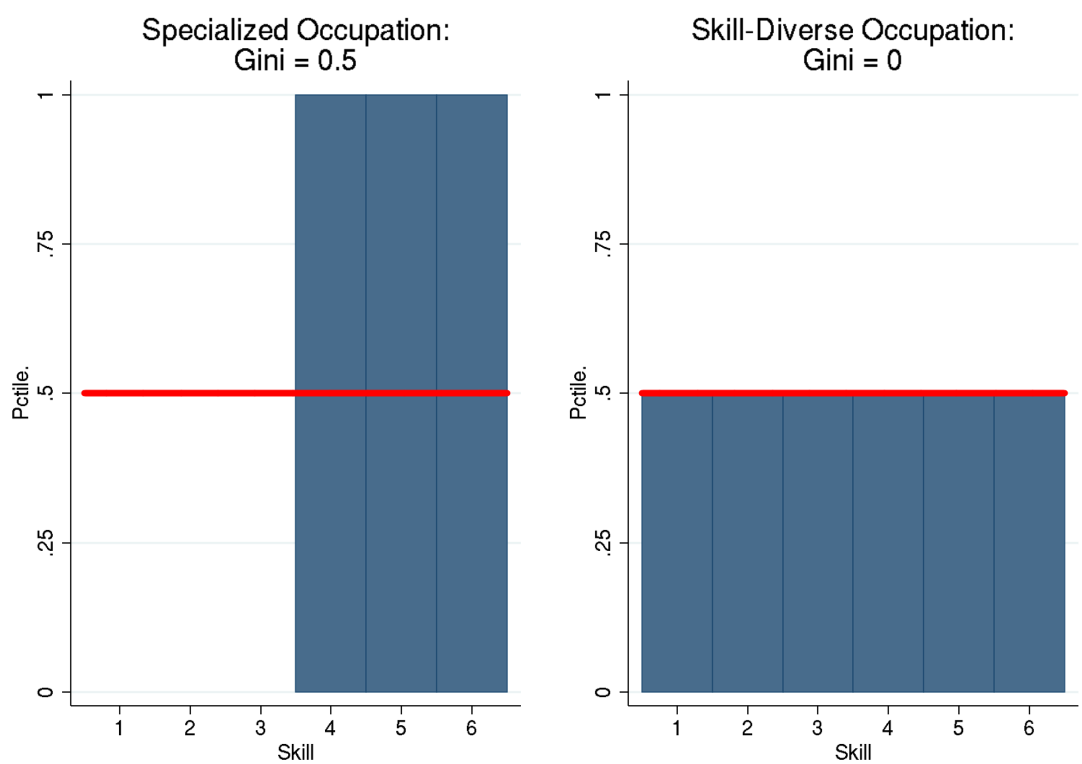

As my primary measures of skill depth and specialization, I construct occupation-level means and Gini coefficients. Although the Gini coefficient is most well known for its use in analyzing income inequality, it is suitable as a measure of dispersion in this case in part because its calculation is invariant to the mean skill percentile. If one

[14] An inherent limitation of these data, and of any similar data, are that the initial list of 35 skills is chosen by O*NET. An implicit and untestable assumption of this analysis is that the skills listed are of similar breadth, an in particular that they are of similar costliness to acquire. My main results focus not on the aggregate level of skill, but on the distribution of a worker's skills, controlling for his/her aggregate level of skill. This may make them less sensitive to definitional considerations. The robustness results included in the Appendix also suggest that results are not especially sensitive to the level of skill aggregation considered, or to the use of the raw scores themselves.

controls for the mean skill percentile, then a high Gini coefficient will be associated with a specialized skillset, while a low Gini coefficient will be associated with a broader skillset. A stylized graphical representation of a highly specialized occupation and a highly unspecialized occupation are shown in Figure 2·4.[15]

In principle, one can construct a mean and/or a Gini coefficient for any level of skill disaggregation, with the obvious limiting case being that the Gini coefficient is mechanically 0 within a set consisting of a single skill. For my primary specifications in this chapter, I construct means and Gini coefficients within two main clusters of skills that I observe in the data; I term these Production skills (8 skills) and Cognitive skills (27 skills). The procedure by which these skill clusters were classified is described below in Section 2.5.4. I also show results based on the construction of a single mean and Gini that incorporates the full set of 35 skills. However, the strength of the division between Production and Cognitive skills observed in the data, as well as the robustness of my main results to alternative specifications, suggests that the Gini coefficient of all skills may be a poor measure of overall specialization within the skillsets that are most relevant to workers in blue-collar and white-collar occupations, respectively.

2.5.2 RAIS

Brazilian matched employer-employee data comes from the *Relação Anual de Informações Sociais* (RAIS) dataset. In Brazil, all firms that are formally registered are required to report information on their employees in each year to the Ministry of Labor. This information is used for the provision of an annual wage supplement, as

[15] The Gini coefficient is always bounded below by zero, regardless of the mean skill percentile. However, because the distribution of skill percentiles is bounded, an increase in the mean skill percentile mechanically reduces the upper bound of the range of feasible Gini coefficients. This is a feature of all measures of dispersion that involve a bounded scale and a finite set of observations (skills). In Table 2.10, I include a robustness check in which I rescale the Gini coefficient by dividing it by its maximum feasible value given its mean. Use of this alternative index does not substantially change the main results of the chapter.

well as a variety of government social programs. The dataset contains both a unique tax identifier for each establishment and a unique identifier for each worker, allowing individuals to be tracked over time as they change employers.[16] The dataset contains standard demographic data on workers, such as education, age, and nationality, and as well as information on workers' job tenure and geography. The dataset also contains data on the industry subsector of establishments at several levels of classification, and the completeness of the data allows me to construct indices regarding the overall size, employment composition, and mean wages of each establishment.

The use of these administrative data provide several advantages over using non-administrative data from the United States. [17] Most immediately, I am able to isolate strictly within-establishment variation in wages by incorporating a full set of establishment fixed effects. Secondly, there is some limited evidence that self-reported occupations systematically inflate their occupation to higher-skilled and higher-paying occupations than the ones whose tasks they perform as measured by administrative data (Fisher and Houseworth, 2012), which may pose problems for inference of occupation-level differences in wage patterns.

On the other hand, using administrative data from another country, and in particular from a developing country, has some disadvantages as well. An important one is that Brazil and the United States code occupations using different systems. In order to map data on the skill content of occupations to Brazilian administrative data, I use existing occupational concordances. The O*NET-SOC 2010 codes used in O*NET are converted to SOC-2000 occupational codes, and then from there to ISCO-

[16]Although individuals may be tracked across employers with this dataset, and although this feature is key to the estimation method used in Section 2.6.4, there are limitations in utilizing the longitudinal component of these data because Brazil is a country with a large informal sector. Since I do not observe informally employed individuals in these data, I cannot rule out the role of selection into formal sector status. Therefore, all results should be interpreted as indicative of wage patterns among formal sector employees only.

[17]As noted in Section 2.5.3, I will make use of the U.S. Current Population Survey (CPS) for some robustness checks.

88 international occupational codes using a publicly available concordance provided by the U.S. Bureau of Labor Statistics.[18] Finally, ISCO-88 codes are matched to 343 Brazilian CBO-94 occupational codes using a cordorance produced by Muendler et al. (2004). In instances where a single CBO-94 occupation is mapped to multiple ISCO-88 occupations, or where a single ISCO-88 occupation is mapped to multiple SOC-2000 occupations, the percentile ranks assigned to each skill for the destination occupation are calculated as the mean of the percentile ranks assigned to that skill in the source occupations, weighted by U.S. occupational employment as reported in the 2015 Occupational Employment Statistics (OES).

Table 2.1: Establishments and Firms By Year In RAIS

	Establishments				*Firms*			
			Mean Size				*Mean Size*	
Year	Total	% New	All	New	Firms	% New	All	New
1995	1,553,659	17.1%	15.3	5.9	1,111,225	16.0%	21.4	5.9
1996	1,613,040	15.8%	14.8	5.3	1,155,703	15.2%	20.6	4.4
1997	1,732,791	17.9%	13.9	4.9	1,263,130	18.4%	19.1	4.4
1998	1,795,798	15.6%	13.6	5.3	1,322,233	15.7%	18.5	4.5
1999	1,863,882	15.2%	13.4	6.0	1,380,148	16.3%	18.1	5.8
2000	1,960,164	15.3%	13.4	5.6	1,461,419	15.4%	17.9	4.8
2001	2,040,808	14.5%	13.3	5.5	1,534,079	15.2%	17.7	4.9
2002	2,142,673	14.3%	20.0	8.9	1,617,713	14.5%	26.5	8.3
2003	2,211,471	13.3%	20.3	10.9	1,673,334	13.4%	26.8	10.9
2004	2,312,940	13.5%	20.8	9.3	1,757,949	13.5%	27.4	8.7
2005	2,412,482	13.3%	21.3	9.0	1,844,068	13.4%	27.9	8.1
2006	2,504,328	12.6%	21.9	9.8	1,928,786	13.0%	28.4	8.7
2007	2,588,073	12.6%	22.7	9.8	2,000,869	13.1%	29.4	8.6
2008	2,712,953	13.5%	22.7	7.6	2,105,641	13.6%	29.2	5.8
2009	2,836,138	13.4%	21.8	6.8	2,213,446	13.9%	27.9	5.0
2010	2,997,673	14.0%	22.5	6.8	2,362,035	14.9%	28.6	5.1
2011	3,144,426	13.9%	22.3	6.6	2,489,955	14.9%	28.2	5.1
2012	3,263,226	12.7%	21.8	6.5	2,603,549	13.2%	27.4	4.5
2013	3,386,424	12.7%	21.1	6.3	2,710,094	13.0%	26.4	4.0
2014	3,210,927	11.8%	21.4	6.2	2,581,144	12.3%	26.7	4.1

From RAIS. New establishments/firms are those that are not previously observed in RAIS since 1986. All employee counts are as of December 31st of that year.

A second disadvantage of using Brazilian data relates to the existence of a large informal sector in Brazil, and to patterns of formalization over time. I cannot rule out that changing patterns of wages or of sorting over time are the result of changing

[18] This concordance appears to be no longer publicly available online, but it is available from the author upon request.

patterns of formalization over time. As Table 2.1 shows, there has been a large increase in the number of establishments reporting to RAIS over time. While newly-formalized establishments are smaller than the overall average, the U-shaped time trend of mean new establishment size suggests that there may be no clear criterion for distinguishing genuinely new establishments from existing ones that are only newly formalized. Additionally, when individuals leave their current employer, a sizeable proportion of individuals leave the RAIS dataset, and it is likely that many of them seek employnment in the informal sector. Selection into formal-sector status may be a concern for those who remain in the dataset. For this reason, I have opted not to include any analyses of wage patterns over time, or any analyses that rely on patterns of employer or occupational change. In future research, I hope to incorporate these sources of variation.

For all worker-level wage regressions, the dependent variable that I use is log annual wages, reported as a multiple of the Brazilian monthly minimum wage. Because all specifications include year fixed effects and I use log wages, results are invariant to this normalization. I focus on the period 1995-2014 in my individual-level regressions for two reasons. The first is that improved data availability allows me to exclude workers who are not in private-sector establishments. The second reason is that prior to 1994, Brazil experienced very high rates of inflation. 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*. There is evidence in the data that the value of the Brazilian minimum wage increased over this time period relative to median formal-sector wages. However, the distribution of wages does not suggest that the minimum wage was strongly binding over this period for my sample of workers used here.[19]

[19]Specifically, in all years, less than 5% of individuals whose wage is reported report that they

I exclude observations with no worker identifier, observations with zero wages, observations in the public and non-profit sectors, and observations in which the worker was not employed at the end of the year. When a worker was reported to hold multiple jobs, I include only the highest-paying job. Additionally, for all worker-level specifications, I restrict the sample to include only males aged 25-54 who are reported to be full-time workers and to have one or more years of tenure.[20]The tenure restriction ensures that the annual earnings measure reflects earnings received by a worker over the course of an entire year.

2.5.3 Current Population Survey (CPS)

A key empirical prediction of this model is that while there may observed wage premia associated with occupation-level skill specialization, these premia will be observed primarily across firms. Within firms, we may actually expect to observe lower wages associated with greater skill specialization, particularly if production inputs are gross complements. In order to distinguish between within-firm and across-firm wage patterns, I make use of matched employer-employee administrative data as described in the previous subsection.

Nonetheless, the use of Brazilian administrative data for this analysis presents potential empirical concerns related to the matching of occupations internationally. In particular, matching occupations in this way is likely to introduce measurement error in my skill content measures. Although I have no particular reason to suspect that any measurement error in skills from this matching process is non-random in any particular way, this concern is at a minimum a potential source of attenuation bias in wage premium estimates. Additionally, we may simply be concerned about the extent to which wage premium patterns in a middle-income country are applicable to

were paid at or close to the minimum wage.

[20]Full-time workers are defined as reported contracted hours of 30-50 hours per week.

a developed economy.

For these reasons, I perform robustness checks using publicly available survey data from the Current Population Survey's Merged Outgoing Rotation Groups Sample (CPS MORG), This dataset contains self-reported information from a random sample of the U.S. civilian noninstitutional population on earnings and occupation as well as education, race, age, and other individual covariates. The CPS also contains data on industry, but no other establishment-level observable characteristics, nor does it contain establishment or firm identifiers.

I construct the CPS robustness check dataset using the 2012-2015 CPS MORG. All regressions include men, ages 25-54, who report that they are employed full-time, and who report an occupation. The dependent variable in all wage regressions is log of weekly earnings. In the CPS MORG files provided publicly by the NBER, occupation is reported using the 2010 Census occupational codes. I map SOC occupations to these codes using a concordance provided by the U.S. Census Bureau (*Industry and Occupation Code Lists & Crosswalks*).

2.5.4 Hierarchical Clustering

As noted in Section 2.4.1, a key empirical concern when measuring occupational specialization is that different occupations may exist in very different labor markets. Skills that are used and valued in some of these markets may be of little or no use in others. For example, O*NET reports that all Economists surveyed reported that skills such as "Equipment Maintenance" and "Installation" were not important to their job. If a measure of specialization incorporates skills that are clearly not in the scope of comparable alternative employement within the same labor market, it will be at best an imprecise measure of specialization. The distinction between traditionally "white-collar" and "blue-collar" occupations may be only one example of different labor markets that might value and employ very different skills.

However, this observation presents a dilemma. It is impractical and problematic to regress wages on all 35 skills individually. Firstly, some skills may be highly correlated within the occupational distribution, which will make any estimates of those skills' association with wages highly imprecise. Secondly, using only occupation-level variation in skill content, it is impossible to construct a comprehensive set of interactions among skills because there are limited degrees of freedom. A better strategy may be to classify skills into broad groups that reflect a division among skills that are most directly used in production activities and skills that are not. But, it is important that such a classification is not performed in an ad hoc manner.

To accomplish this task, I classify skills using a method of agglomerative hierarchical clustering. The hierarchical clustering method provides several benefits over other, similar methods.[21]Most importantly for the purposes of this chapter, it provides a classification of the set of skills that is transparent and readily interpretable.[22] Unlike some alternative measures, it does not require specifying the number of desired clusters in advance. And, by classifying skills into discrete clusters, the clustering algorithm allows me to construct measures of skill specialization within each cluster while also controlling for the mean level of skill within each cluster.

The procedure begins by constructing a 35×35 matrix of pairwise correlations

[21]For an overview of the hierarchical clustering method as compared to alternative methods, see Hastie, Tibshirani, and Friedman (2011). Although, to my knowledge, the hierarchical clustering method has been little-used in Economics to date, it is most similar in effect to the k-means clustering procedure, which has been used in some contexts. Unlike k-means clustering, hierarchical clustering does not require specifying the number of clusters in advance, and it is not susceptible to misclassification in the presence of multiple local minima.

[22]For example, principal component analysis (PCA) and factor analysis are two methods that have a history of use in Economics. Although PCA will by definition capture variation in the skill data more parsimoniously than a clustering algorithm, it does so at a substantial cost in interpretability, because it characterizes component vectors that may place either positive or negative weights on certain skills. In a principal component analysis of the O*NET skill data, the first two components correspond closely to the classifications of Cognitive and Production skills that I report. The correlation between each occupation's score in the first component and its mean Cognitive skill measure is 0.99, while the correlation between each occupation's score in the second component and its mean Production skill measure is 0.85. Together, these two principal components explain approximately 70% of the variation in skills in the O*NET database.

between each of the 35 skills in the O*NET database, where element $[i, j]$ represents the correlation between skills i and j across the set of occupations reported in O*NET. This correlation matrix is then used as a measure of skill dissimilarity.[23] The algorithm then begins by considering each of the 35 skills reported in O*NET as a distinct cluster. From there, the algorithm proceeds iteratively to agglomerate clusters as follows:

1. Identify the two remaining clusters with the highest mean pairwise correlation;

2. Classify the members of these two clusters as belonging to a single cluster;

3. Repeat steps 1 and 2 until only a single cluster remains.

Because the algorithm can be observed in every iteration, the algorithm can be used to separate the set of skills into any desired number of clusters.

Figure 2·5 provides a visual representation of the skill clusters generated by the hierarchical clustering algorithm in the form of a cluster dendrogram (tree diagram). Two key stylized facts emerge from this analysis. Firstly, many skills, such as Reading Comprehension and Writing, are very highly correlated across occupations. This is broadly consistent with the theoretical model presented in this chapter; if individual workers were highly specialized among these skills, then we would expect the correlation between any two skills to be relatively low. These high correlations also suggest that an empirical strategy that includes all 35 skills separately is likely to face severe multicollinearity issues.

Secondly, and more importantly, not all skills are highly correlated with one another. Rather, there exist two distinct groups of skills in the O*NET dataset. The first

[23]The use of alternative distance metrics, such as the cosine angle of separation used by Gathmann and Schönberg (2010), does not alter the classification of skills into the two primary clusters that are analyzed in this chapter. However, the use of the correlation matrix allows for a more straightforward interpretation of the algorithm's results; the value at which any two clusters are merged is the mean pairwise correlation of the elements of each cluster.

Figure 2·5: Results of Agglomerative Hierarchical Clustering Algorithm on O*NET Skill Measures

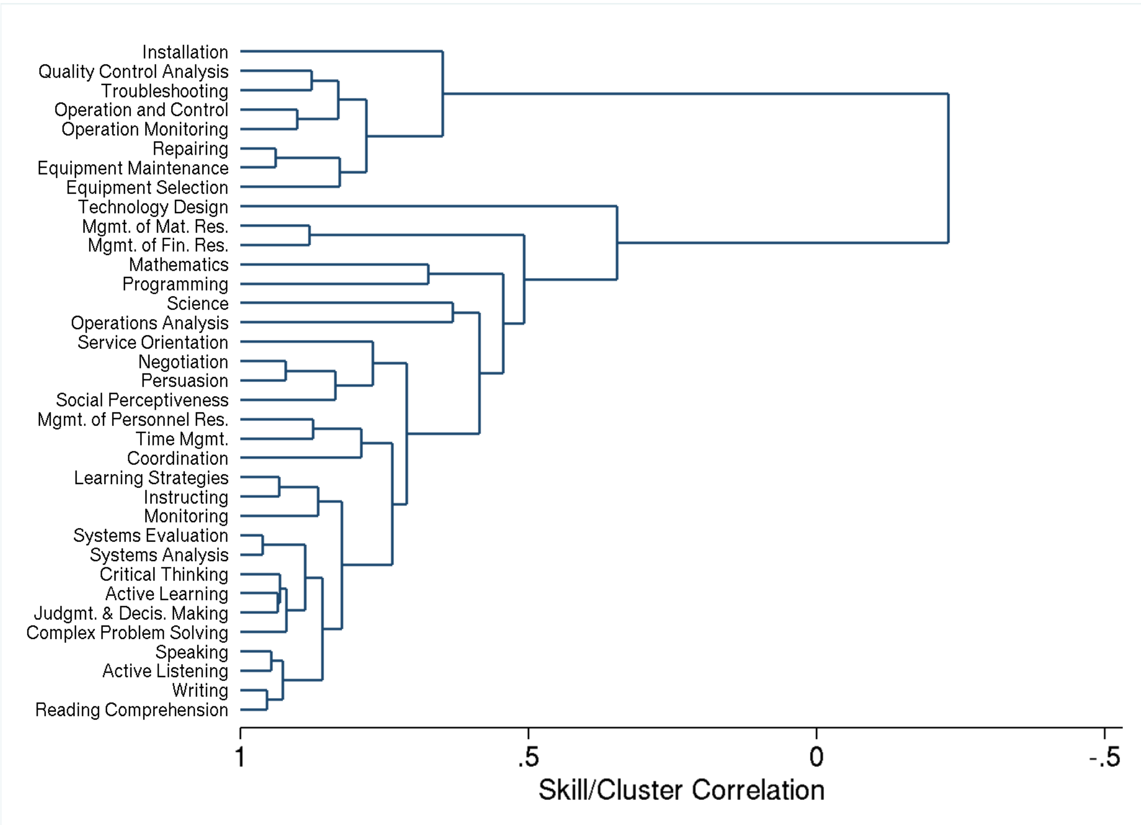

group consists of 8 different skills, which I term Production skills. These skills include Equipment Maintenance, Equipment Selection, Installation, Operation and Control, Operation Monitoring, Quality Control Analysis, Repairing, and Troubleshooting. Each of these skills is broadly associated with production activities, and in particular the descriptions given for each of these skills generally references regular interaction with equipment or machines used in production.

The second cluster of skills, which I term Cognitive skills, includes 27 skills.[24] These skills cover a wide range, and they include both skills that may have a relatively intuitive and explicit definition (Programming, Mathematics) and skills whose definition may be more amorphous (Active Learning, Social Perceptiveness). Because many of these skills are highly correlated across the occupational distribution, I have opted not to disaggregate skills further. However, in principle it may be feasible to do so for future research.

The division of skills into these two primary clusters is highly robust to specification of the clustering procedure. In fact, within the Production skill cluster, 100% of the pairwise correlations are positive, and similarly, 100% of the pairwise correlations in the Cognitive cluster are positive. However, over 80% of the pairwise correlations across clusters are negative; occupations tend to report that they are relatively skilled in one cluster, and relatively less skilled in the other. Returning to Figure 2·3, it is clear that Economists also exhibit this pattern.

Figure 2·6 plots each occupation's mean skill intensity in each of the two groups,

[24] O*NET has also constructed a division of skills into six primary groups (*Skills Search*). To the best of my knowledge, this classification has been constructed manually. All 8 of the skills that I define as Production skills are classified as "Technical" skills by O*NET. Of the 27 skills that I define as Cognitive skills, three are classified by O*NET as Technical skills (Technology Design, Programming, and Operations Analysis), while the remaining skills are separated among "Basic" skills (10 skills), "Complex Problem Solving" skills (1 skill), "Resource Management" skills(4 skills), "Social" skills (6 skills), and "Systems" Skills (3 skills). In the Appendix, I reconstruct the baseline regressions of this chapter using each of those categories as a single Cognitive skill in the calculation of the Gini coefficient to address any concern that my results may arise from the greater number of skills in the Cognitive cluster.

Figure 2·6: Scatterplot of Mean Production Skills vs. Mean Cognitive Skills By Occupation

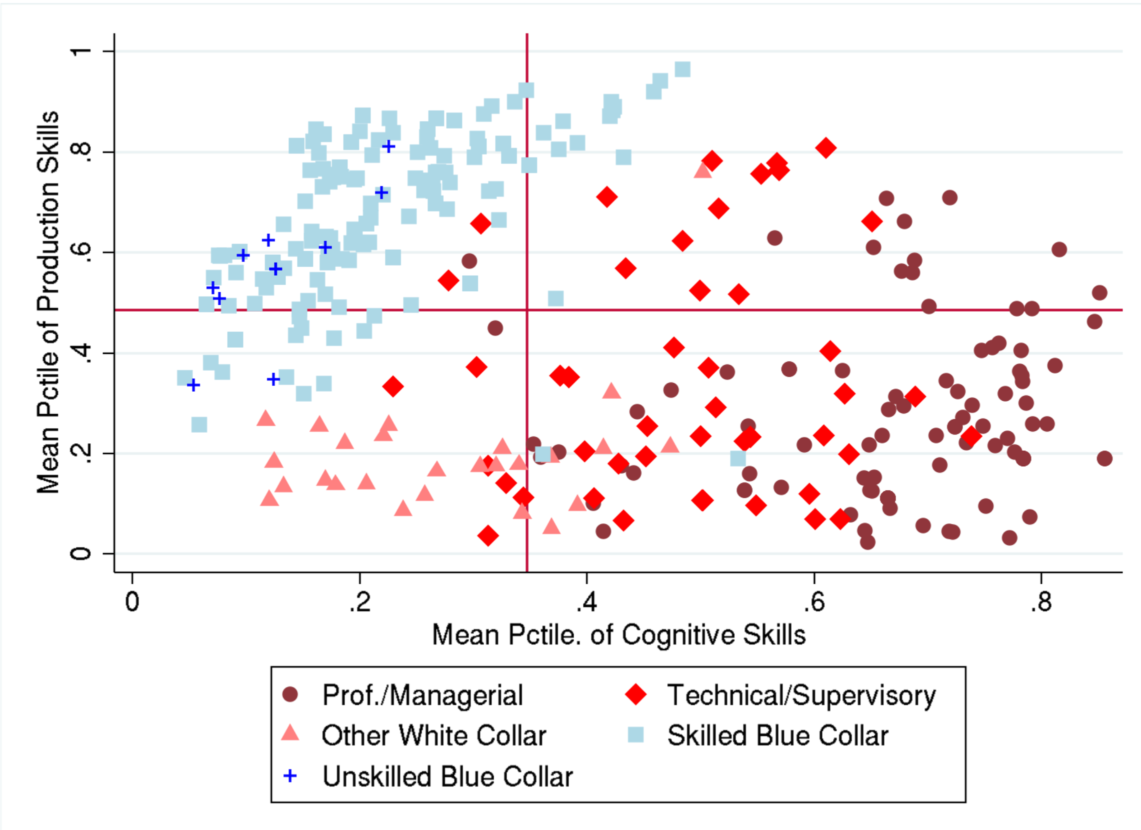

with Cognitive skills on the horizontal axis and Production skills on the vertical axis. Occupations markers denote one of five main occupational groupings generated by to which the occupation belongs (as in Menezes-Filho, Muendler, and Ramey, 2008). Here again there is evidence of a division among occupations; relatively few occupations report similar levels of mean skill intensity for both Production Skills and Cognitive skills. Furthermore, the division of occupations into Production-intensive and Cognitively-intensive maps closely, though imperfectly, to the division of blue-collar and white-collar occupations. For this reason, I perform several robustness checks to my main specifications, separating blue-collar and white-collar occupations to ensure that my primary results related to Cognitive and Production skill content are not being driven by occupational groups with low average levels of skills in those groups.

Figure 2·7 plots each occupation's Gini coefficient of Cognitive Skills against its Gini coefficient of Production skills. Again, there is a clear division among blue-collar and white-collar workers. However, here we can observe another interesting, though perhaps unsurprising pattern; occupations tend to have a higher measured concentration of skills, i.e. to be more specialized, within the class of skills (Cognitive or Production) in which they are less skilled on average. Some portion of this relationship–though not all–is related to the mechanical upper bound that a high mean skill percentile places on the Gini coefficient. In the Appendix, I include these same measures for alternative indices.

2.5.5 Comparison of Specialization Measures with Other Skill Measures

While the notion of labor specialization is not new, to date it has not been an area of focus in the empirical analysis of occupational skill content. Instead, numerous other papers have examined labor market changes along specific dimensions of skill by constructing alternative measures that encapsulate these dimensions. Perhaps the most

Figure 2·7: Scatterplot of Gini Coefficient of Production Skills vs. Cognitive Skills By Occupation

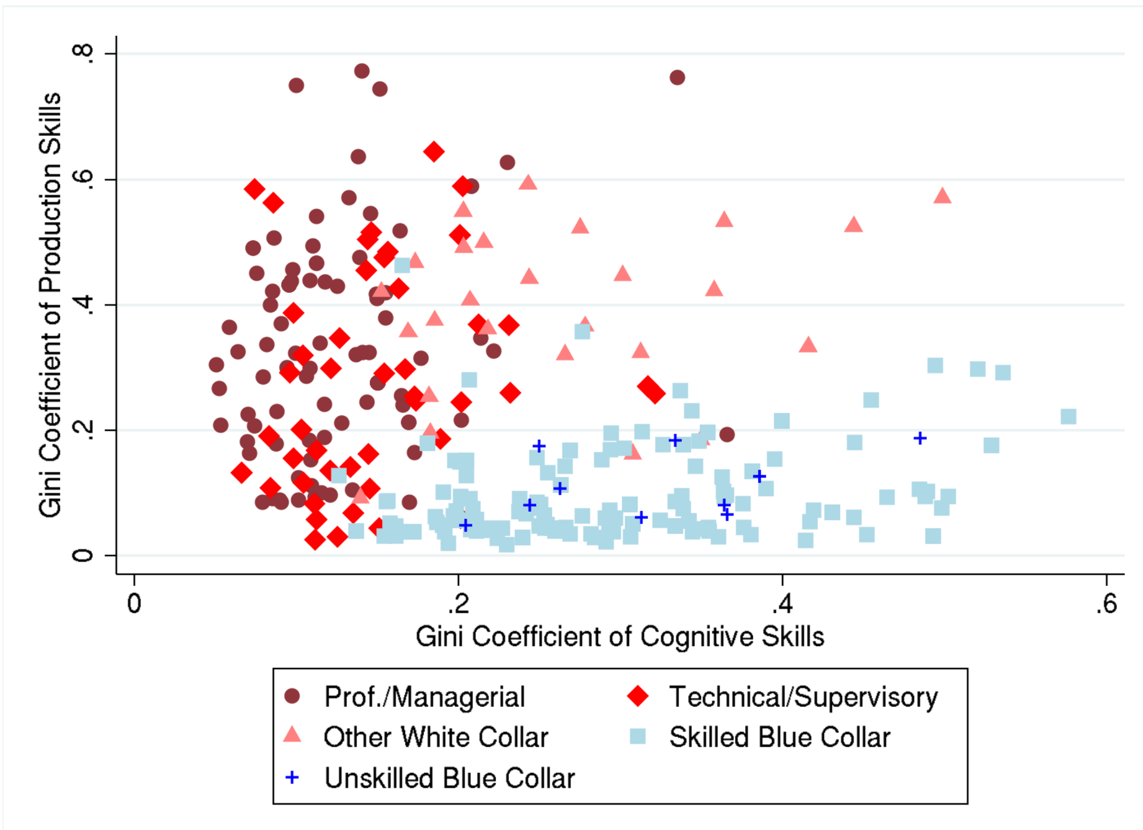

well-known of these are the series of measures constructed and analyzed in several papers by David Autor and coauthors (Autor, Levy, and Murnane, 2003; Autor and Dorn, 2013). In particular, Autor, Levy, and Murnane (2003) use the Dictionary of Occupation Titles (the precursor to O*NET) to examine changes in aggregate occupational structure over time that may be associated with technological change. In particular, since routine tasks are argued to be more suspectible to automation, the authors construct measures of routine and non-routine task content. In Autor and Dorn (2013), the authors introduce an additional, similar measure, which they term the Routine Task Index (RTI). This index is a composition of three separate measures of task content: Routine, Manual, and Abstract tasks.

Table 2.2: Comparison of Autor and Dorn Skill Measures with Specialization measures

	(1) RTI	(2) Routine	(3) Manual	(4) Abstract
Mean of Cognitive Skills	-0.349***	-0.217**	-0.198**	0.904***
	(0.0973)	(0.0903)	(0.0992)	(0.0625)
Mean of Production Skills	0.111	0.540***	0.159	-0.0354
	(0.117)	(0.109)	(0.120)	(0.0754)
Gini of Cognitive Skills	0.0233	-0.112	-0.0950	0.125**
	(0.0929)	(0.0862)	(0.0947)	(0.0597)
Gini of Production Skills	0.373***	0.0834	-0.152	0.0232
	(0.118)	(0.109)	(0.120)	(0.0757)
Observations	438	438	438	438
R-squared	0.156	0.270	0.135	0.661

Each dependent variable is the indicated measure of Autor and Dorn (2013), standardized across Census occupations. Each explanatory variable is a standardized measure of mean skill or specialization from O*NET. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 2.2 provides the results of regressions of standardized values of the four measures (RTI plus its three task components) considered in Autor and Dorn (2013) on standardized values of the mean skill and concentration measures constructed in this chapter. In all specifications, the unweighted regression is conducted on the 438 three-digit Census occupations that I am able to match between the authors' measure and my own. Since all observations are also standardized, the coefficients can be interpreted as the predicted standardized change in each of my measures given a one standard deviation change in the indicated measure from Autor and Dorn, *ceteris paribus*. Columns 1 and 2 of the table shows that both of the measures of "routine" occupational task content are associated with lower levels of Cognitive skill usage in employment, with somewhat higher levels of Production skill, and with more specialization within Production skills. In comparison, Column 3 shows that occupations that are high in "non-routine manual" task content, while also comparatively strong in Production skills, are not associated with specialized Production skillsets. The literature on technological change suggests that routine manual tasks may be most readily automated, and the finding that routine occupations are both comparatively non-cognitive and comparatively skill specialized is consistent with this suggestion. Furthermore, since Autor and Dorn have previously shown that routine occupations tend have higher average wages than non-routine manual occupations, these correlations would be very consistent with an empirical finding that greater specialization in Production skills is associated with higher wages. Finally, Column 4 shows that the Autor-Dorn "abstract" task measure is strongly associated with higher usage of Cognitive skills, and also with higher concentration within the set of those skills that are used.

2.6 Results

2.6.1 Summary of RAIS Data

Before proceeding with the regression analysis, it is useful to understand how and why individuals with more or less specialized skillsets differ. Table 2.3 summarizes these differences within the sample of matched employer-employee data from RAIS. Columns 1 and 3 of this table represent individuals whose occupations are Most Specialized and Least Specialized, respectively, in that the Gini coefficient of their

Table 2.3: Summary Stats for Individuals in Most/Least Skill-Specialized Occupations

	Most Specialized	Avg. Specialization	Least Specialized	Total
Mean Skill Pctile.	0.217	0.347	0.598	0.363
	(0.0575)	(0.110)	(0.115)	(0.149)
Mean Gini of Skills	0.513	0.318	0.156	0.325
	(0.0467)	(0.0673)	(0.0307)	(0.115)
Earnings/MW	2.866	6.282	12.99	6.721
	(2.956)	(8.319)	(13.92)	(9.265)
Age	38.71	38.77	39.45	38.85
	(8.068)	(7.903)	(7.730)	(7.908)
Tenure	2.164	2.331	2.713	2.360
	(1.201)	(1.333)	(1.583)	(1.362)
Contracted Hours	43.44	42.74	42.18	42.77
	(1.854)	(3.105)	(3.038)	(2.959)
Primary/Middle School	0.776	0.488	0.225	0.494
	(0.417)	(0.500)	(0.417)	(0.500)
High School	0.209	0.389	0.397	0.363
	(0.407)	(0.487)	(0.489)	(0.481)
Some College	0.00603	0.0360	0.0770	0.0373
	(0.0775)	(0.186)	(0.267)	(0.189)
College Graduate	0.00730	0.0864	0.301	0.105
	(0.0851)	(0.281)	(0.459)	(0.307)
Median Estab. Size	190	157	326	182
Mean Estab. Size	1227.5	1196.6	1819.4	1290.3
	(2998.5)	(3090.9)	(4312.0)	(3287.4)
Log Mean Estab. Earnings	0.947	1.362	1.742	1.353
	(0.652)	(0.808)	(0.886)	(0.827)
Observations	9,460,483	43,728,192	8,873,158	62,061,833

From a 100% sample of eligible person-years from RAIS, 1995-2014. Mean coefficients are reported, with standard deviations in parentheses. Most specialized occupations are 1+ standard deviations above the sample mean. Least specialized occupations are 1+ standard deviations below the sample mean.

skills as reported in O*NET is at least one standard deviation above (or one standard deviation below) the sample mean. Column 2 represents all other individuals in the sample, and column 4 shows overall averages for the sample.

The first and most immediate observation from the Table is the there is a clear negative relationship between workers' mean skill percentile and the Gini coefficient of their occupation's skills as indicated by O*NET. That is, the most skilled workers also appear to have broader skillsets that are less concentrated in a few skills; that is, they are less specialized. Some of this relationship is likely attributable to the inherent relationship between any measure of dispersion, such as the Gini coefficient, and the mean of a bounded scale with a finite sample size.

The second observation, which is again fundamentally at odds with the baseline model, is that workers with the least specialized skillsets have much higher earnings and higher educational attainment than workers with more specialized ones. They have much higher rates of educational attainment, and they disproportionately work at larger establishments that pay higher wages on average. These last two findings are especially of note; if specialization is valuable, then larger establishments and establishments that pay higher wages should be in theory prefer to hire workers in more specialized occupations. [25] Taken collectively, these results are broadly inconsistent on their face with the notion that workers in specialized occupations will have higher wages. They also underscore the extent to which, in any regression analysis using these data, it is important to account for other observable differences, including in particular differences in the mean level of worker skill.

As previously noted in Sections 2.4.1 and 2.5.4, a significant empirical concern is that skill specialization should ideally be measured against the set of skills that are used within the similar occupations in the same overall labor market. Measuring

[25] In the model presented in this chapter, as in other models of specialization, firm size, technology, and the feasible extent of specialization are directly related. There is also empirical evidence that larger establishments pay higher wages on average (Brown and Medoff, 1989; Oi and Idson, 1999).

Table 2.4: Summary Stats for Individuals in Most/Least Production-Specialized Blue-Collar Occupations

	Most Specialized	Avg. Specialization	Least Specialized	Total
Mean Cognitive Skill Pctile.	0.188	0.206	0.309	0.211
	(0.103)	(0.0874)	(0.106)	(0.0945)
Mean Production Skill Pctile.	0.420	0.697	0.852	0.677
	(0.123)	(0.121)	(0.0582)	(0.154)
Mean Gini of Cognitive Skills	0.342	0.266	0.261	0.274
	(0.0807)	(0.0737)	(0.0965)	(0.0799)
Mean Gini of Production Skills	0.223	0.0846	0.0277	0.0962
	(0.0708)	(0.0387)	(0.00408)	(0.0632)
Earnings/MW	3.723	4.401	5.946	4.426
	(3.622)	(4.872)	(5.465)	(4.813)
Age	37.74	39.26	38.31	39.03
	(7.921)	(7.939)	(7.791)	(7.944)
Tenure	2.182	2.234	2.438	2.242
	(1.177)	(1.239)	(1.360)	(1.242)
Contracted Hours	43.42	43.25	43.07	43.26
	(1.919)	(2.157)	(2.159)	(2.134)
Primary or Middle School	0.689	0.686	0.575	0.679
	(0.463)	(0.464)	(0.494)	(0.467)
High School	0.294	0.291	0.393	0.298
	(0.456)	(0.454)	(0.488)	(0.457)
Some College	0.00824	0.00944	0.0138	0.00959
	(0.0904)	(0.0967)	(0.117)	(0.0975)
College Graduate	0.00883	0.0123	0.0185	0.0123
	(0.0935)	(0.110)	(0.135)	(0.110)
Median Estab. Size	224	160	296	198
Mean Estab. Size	1126.6	1272.3	1854.8	1294.2
	(2436.2)	(3001.7)	(4022.6)	(3026.8)
Log Mean Estab. Earnings	1.090	1.197	1.517	1.206
	(0.655)	(0.742)	(0.767)	(0.740)
Observations	4,532,260	19,415,254	7,387,316	31,334,830

From a 100% sample of eligible person-years from RAIS, 1995-2014. Mean coefficients are reported, with standard deviations in parentheses. Most specialized occupations are 1+ standard deviations above the sample mean. Least specialized occupations are 1+ standard deviations below the sample mean.

Table 2.5: Summary Stats for Individals in Most/Least Cognitively-Specialized White-Collar Occupations

	Most Specialized	Avg. Specialization	Least Specialized	Total
Mean Cognitive Skill Pctile.	0.143	0.441	0.744	0.451
	(0.0470)	(0.160)	(0.0739)	(0.201)
Mean Production Skill Pctile.	0.199	0.264	0.364	0.270
	(0.0577)	(0.216)	(0.152)	(0.203)
Mean Gini of Cognitive Skills	0.376	0.168	0.0790	0.177
	(0.0917)	(0.0430)	(0.0138)	(0.0854)
Mean Gini of Production Skills	0.432	0.322	0.154	0.312
	(0.112)	(0.157)	(0.0988)	(0.162)
Earnings/MW	3.337	8.568	16.56	9.061
	(3.214)	(11.07)	(16.04)	(11.79)
Age	38.27	38.55	39.79	38.68
	(7.900)	(7.882)	(7.650)	(7.867)
Tenure	2.223	2.469	2.758	2.481
	(1.234)	(1.460)	(1.605)	(1.465)
Contracted Hours	43.06	42.18	42.15	42.26
	(2.454)	(3.724)	(2.939)	(3.541)
Primary or Middle School	0.536	0.306	0.130	0.306
	(0.499)	(0.461)	(0.336)	(0.461)
High School	0.415	0.456	0.262	0.428
	(0.493)	(0.498)	(0.440)	(0.495)
Some College	0.0158	0.0677	0.0895	0.0655
	(0.125)	(0.251)	(0.285)	(0.247)
College Graduate	0.0334	0.170	0.517	0.200
	(0.180)	(0.375)	(0.500)	(0.400)
Median Estab. Size	94	143	466	166
Mean Estab. Size	777.2	1268.6	1790.2	1286.5
	(2390.5)	(3345.7)	(5037.0)	(3533.4)
Log Estab. Mean Earnings	1.072	1.498	1.872	1.504
	(0.601)	(0.887)	(0.874)	(0.882)
Observations	3,067,132	23,970,793	3,689,078	30,727,003

From a 100% sample of eligible person-years from RAIS, 1995-2014. Mean coefficients are reported, with standard deviations in parentheses. Most specialized occupations are 1+ standard deviations above the sample mean. Least specialized occupations are 1+ standard deviations below the sample mean.

specialization against a set of skills that are not in general used by similar occupations may make any estimates imprecise. The unexpected patterns observed in Table 2.3 suggest that this particular concern may be relevant in this context. As detailed in Section 2.5.4, patterns of occupational skill usage in the O*NET data support the division of occupational skills into two primary groups, "Production" and "Cognitive" skills, and the relative usage of these skill groups maps closely to the divide between blue-collar and white-collar occupations.

For this reason, in Tables 2.4 and 2.5, I provide summary statistics on two subsets of the RAIS sample, along measures of these two different types of specialization. Table 2.4 divides blue-collar workers by their degree of specialization in Production skills, while Table 2.5 divies white-collar workers by their degree of specialization in Cognitive skills. These tables demonstrate that the key patterns observed in Table 2.3 are observed within each group, and are not driven by the inclusion of both white collar and blue collar workers in the sample. In both tables, there remains a negative relationship between mean skill percentile and the Gini of skills. And, in both tables, as before, less specialized workers earn more than more specialized workers do. Finally, in both tables, less specialized workers work at larger firms which pay higher wages. Noticably, each of these results is considerably more pronounced among white-collar workers than among blue-collar workers. However, each of these results still runs counter to the predictions of the baseline model.

2.6.2 Wage Premia in RAIS

Table 2.6 presents baseline results with all skills incorporated into a single measure of specialization. In Table 2.7, Cognitive and Production skills are incorporated separately against the entire sample. Column 1 of each table shows results with no other covariates included except for year effects, while column 2 incorporates incorporates standard individual-level covariates, including educational group and nationality con-

Table 2.6: Baseline Regressions: All Skills

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
Mean Skill Pctile.	0.376***	0.245***	0.0916***	0.0804***	0.0886***	0.0769***
	(0.0589)	(0.0321)	(0.0308)	(0.0192)	(0.0172)	(0.0135)
Gini of All Skills	-0.0664	0.0211	0.00635	-0.00239	0.0201	0.00519
	(0.0648)	(0.0408)	(0.0243)	(0.0185)	(0.0174)	(0.0132)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.306	0.518	0.575	0.658	0.683	0.819
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

Table 2.7: Baseline Regressions: Cognitive and Production Skills

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
Mean Cognitive Skills	0.395***	0.176***	0.0732**	0.0773***	0.0718***	0.0822***
	(0.0473)	(0.0287)	(0.0322)	(0.0217)	(0.0184)	(0.0155)
Mean Production Skills	0.228**	0.228***	0.0877**	0.0880***	0.0961***	0.0594***
	(0.0912)	(0.0487)	(0.0379)	(0.0235)	(0.0215)	(0.0179)
Gini Cognitive Skills	-0.0691	-0.0563*	-0.00806	0.0125	0.0270	0.0238
	(0.0513)	(0.0305)	(0.0319)	(0.0217)	(0.0198)	(0.0165)
Gini Production Skills	0.145**	0.0967***	0.0586*	0.0585***	0.0607***	0.0417**
	(0.0728)	(0.0363)	(0.0322)	(0.0223)	(0.0217)	(0.0165)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.318	0.525	0.575	0.659	0.684	0.820
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

trols, linear and quadratic terms in age, experience, job and tenure. In column 3, I incorporate 85 detailed two-digit occupational group fixed effects, which control flexibly for any other characteristics of occupational groups that may be associated with higher or lower wages, such as hedonic factors, policies that constrain labor supply, etc. In columns 4 and 5 I add approximately 600 subsector controls based on the Brazilian *Classificação Nacional de Atividades Econômicas* (CNAE) classification[26] , followed by establishment-level controls for the proportions of establishment employment in each of the four primary educational groups and five aggregate occupational groups. All establishment-level controls exclude the individual in question (single-employee establishments are therefore dropped from column 5). Finally, in column 6 I add a full-set of establishment fixed effects, of which there are approximately 1.9 million. These results, as with all individual-level regression results that I present in this chapter, are clustered at level of the 343 individual CBO occupational codes on which my occupational-level measures vary, making them robust to arbitrary serial correlation within individual occupations. Additionally, all results are reported as standardized coefficients. Therefore, all results can be interpreted as the percentage effect on wages of a one standard deviation increase in the covariate of interest.

The results of both Table 2.6 and Table 2.7 show, firstly, that individuals in occupations that are more skilled, as characterized by O*NET, do in fact earn significantly higher wages. This is unsurpising, although it validates the use of O*NET data as a measure of occupation-level differences in skill. However, Table 2.7 also shows that specialization among Production skills is associated with higher earnings, even after controlling for these differences in mean skill. The results are large; Column 3 shows that a one standard deviation increase in the Gini of production skills is associated with nearly 6% higher earnings, even after controlling for mean levels of skill, and

[26]The Brazilian CNAE subsector classification is somewhat comparable in number and scope to the U.S. NAICS classifications.

controlling flexibly for omitted factors that may affect broader occupational groups. Even once a full set of establishment fixed effects are included, Column 6 shows that the wage premium associated with specialization is estimated to be approximately 4%. In contrast, the degree of specialization in Cognitive skills does not appear to be associated with a significant wage premium. And, when these two groups of skill are aggregated together as in Table 2.6, there appear to be no significant relationship between the degree of specialization and wages at all once mean skill and worker characteristics are accounted for.

Additionally, the results of Columns 1 through 6 of Table 2.7 show an interesting pattern. In particular, the wage premium associated with specialization in Production skills appears to be somewhat larger when incorporating variation both across occupational groups and across establishments, as in columns 1 and 2. However, the estimated wage premium shown in Column 3 and that of Column 6 are little different; this suggest that intra-firm wage premium associated with specialization is not substantially different from the inter-firm premium. Among Cognitive skills, without occupational group controls, point estimates would suggest that specialization is actually associated with lower wages. Yet, once establishment characteristics or establishment fixed effects are incorporated, the sign of these estimates change. Because of the limited statistical power of these methods with occupation-level clustering, I cannot statistically reject the hypothesis that both within-firm and across-firm measures of the wage premia associated with specialization are identical. However, such results could be consistent with unspecialized Cognitively-skilled workers being sorted into high-wage establishments, as will be examined in the next section.

One concern with running the baseline regressions on the entire sample is, as with the summary statistics, that we may be concerned that our estimates of the gains or lack of gains from specialization could be driven by individuals who are not, in

Table 2.8: Baseline Regressions: White-Collar and Blue-Collar Workers

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
			White-Collar Workers Only			
Mean Cognitive Skills	0.406***	0.176***	0.0740**	0.0316	0.0405*	0.0515**
	(0.0828)	(0.0386)	(0.0334)	(0.0234)	(0.0225)	(0.0205)
Mean Production Skills	0.296**	0.241***	0.100***	0.0845***	0.0752***	0.0536**
	(0.148)	(0.0699)	(0.0354)	(0.0225)	(0.0235)	(0.0214)
Gini Cognitive Skills	0.0135	-0.0240	0.0327	-0.0259	-0.00623	-0.0199
	(0.108)	(0.0394)	(0.0479)	(0.0285)	(0.0262)	(0.0255)
Gini Production Skills	0.118	0.0752*	0.0400	0.0491**	0.0492**	0.0407***
	(0.105)	(0.0437)	(0.0318)	(0.0202)	(0.0205)	(0.0154)
Observations	29,930,497	29,846,739	29,846,739	29,846,739	29,625,422	29,846,739
Adjusted R-squared	0.308	0.561	0.597	0.674	0.698	0.830
			Blue-Collar Workers Only			
Mean Cognitive Skills	0.106	0.0295	0.00262	0.00817	0.0210	0.0408
	(0.0937)	(0.0706)	(0.0710)	(0.0460)	(0.0435)	(0.0390)
Mean Production Skills	0.499***	0.424***	0.216**	0.254***	0.226***	0.163***
	(0.135)	(0.104)	(0.103)	(0.0567)	(0.0572)	(0.0510)
Gini Cognitive Skills	-0.110**	-0.0771*	-0.0177	0.0359	0.0412*	0.0419*
	(0.0547)	(0.0432)	(0.0421)	(0.0250)	(0.0241)	(0.0217)
Gini Production Skills	0.347**	0.300**	0.300***	0.250***	0.209***	0.143**
	(0.170)	(0.133)	(0.114)	(0.0765)	(0.0776)	(0.0713)
Observations	29,872,564	29,849,681	29,849,681	29,849,681	29,575,222	29,849,681
Adjusted R-squared	0.241	0.382	0.461	0.582	0.614	0.796
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

Table 2.9: Baseline Regressions: Excluding Small Establishments

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
			Estab. Size 10 or More			
Mean Cognitive Skills	0.410***	0.192***	0.0909***	0.0880***	0.0724***	0.0832***
	(0.0495)	(0.0280)	(0.0307)	(0.0194)	(0.0181)	(0.0159)
Mean Production Skills	0.185**	0.197***	0.0797**	0.0866***	0.101***	0.0600***
	(0.0752)	(0.0365)	(0.0359)	(0.0228)	(0.0214)	(0.0180)
Gini Cognitive Skills	-0.0542	-0.0438	0.00676	0.0192	0.0274	0.0230
	(0.0514)	(0.0281)	(0.0313)	(0.0212)	(0.0200)	(0.0169)
Gini Production Skills	0.106	0.0696**	0.0537*	0.0564**	0.0617***	0.0415**
	(0.0666)	(0.0308)	(0.0310)	(0.0221)	(0.0215)	(0.0166)
Observations	53,074,217	52,986,576	52,986,576	52,986,576	52,986,576	52,986,576
Adjusted R-squared	0.345	0.546	0.588	0.660	0.679	0.807
			Estab. Size 50 or More			
Mean Cognitive Skills	0.444***	0.217***	0.0913***	0.0921***	0.0732***	0.0851***
	(0.0580)	(0.0319)	(0.0288)	(0.0193)	(0.0187)	(0.0173)
Mean Production Skills	0.135**	0.161***	0.0766**	0.0861***	0.102***	0.0619***
	(0.0625)	(0.0324)	(0.0338)	(0.0225)	(0.0216)	(0.0187)
Gini Cognitive Skills	-0.0326	-0.0272	0.0136	0.0251	0.0305	0.0244
	(0.0595)	(0.0326)	(0.0311)	(0.0215)	(0.0209)	(0.0182)
Gini Production Skills	0.0468	0.0311	0.0510*	0.0539**	0.0590***	0.0392**
	(0.0686)	(0.0340)	(0.0292)	(0.0226)	(0.0218)	(0.0170)
Observations	40,576,414	40,505,591	40,505,591	40,505,591	40,505,591	40,505,591
Adjusted R-squared	0.373	0.564	0.605	0.669	0.682	0.794
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

general, skilled in similar things. So, in Table 2.8, I show in two panels the results of running the regressions separately on white collar and blue collar workers. The results for blue-collar workers again show clear evidence of the value of specialization in Production. In fact, the estimated wage premium associated with a one standard deviation increase in the Gini of production skills is estimated at 14-30%. However, among white-collar workers, there remains no evidence that specialization in Cognitive skills is valuable.

An additional concern, particularly in looking at the lack of wage premia associated with specialization in Cognitive skills, may be that specialization patterns are particularly driven by small firms. As characterized in the model, individuals in many small firms may be expected to perform a relatively broad set of tasks, and therefore may be unable to developed specialized skills. In larger firms, we might expect that the extent of specialization is less constrained, leading to little value placed on skill specialization. Table 2.9 show regression results excluding any establishments with less than 10 workers in the sample (Panel A), and with less than 50 workers in the sample (Panel B). Note that although a majority of establishments are excluded from each of these subsamples, most individuals remain in both samples, as most formal sector workers in Brazil are employed in larger establishments. Regardless, these results do not suggest that either the estimated wage premium for Production skill specialization or the lack of an estimated wage premium for Cognitive skill specialization are driven primarily by small establishments.

Finally, one may be concerned that the results are driven either by collinearity between the mean percentile of each set of skills and the Gini coefficient, or by the manner in which raw scores from O*NET have been converted to percentiles. So, table 2.10 shows results from two alternative measures. In Panel A, the Gini coefficient for each occupation and each skill group has been rescaled. Because the distribution of skill percentiles is bounded, an increase in the mean skill percentile mechanically reduces the upper bound of the range of feasible Gini coefficients. So, for each occupation and each skill group, the upper bound of feasible Gini coefficients is calculated, and the true Gini coefficient is divided by this measure. In Panel B, I use the standard (un-rescaled) Gini coefficent measure, but I calculate both means and Gini coefficients on the raw O*NET scores, rather than on the data that have been converted to percentiles of the occupational distribution. Again, the basic conclusions of the baseline results are largely reaffirmed.

In the Appendix, I provide regression results based on two alternative index measures of specialization; the Theil index and the standard deviation of reported skills.

Table 2.10: Baseline Regressions: Alternative Skill Measures

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
			Rescaled Gini Measure			
Mean Cognitive	0.444***	0.216***	0.0652**	0.0576***	0.0445***	0.0601***
	(0.0273)	(0.0193)	(0.0255)	(0.0179)	(0.0158)	(0.0137)
Mean Production	0.134**	0.166***	0.0639**	0.0641***	0.0679***	0.0409***
	(0.0541)	(0.0292)	(0.0301)	(0.0149)	(0.0131)	(0.0123)
Rescaled Gini	-0.0323	-0.0260	-0.00999	-0.00473	0.00353	0.00285
Cognitive	(0.0341)	(0.0193)	(0.0195)	(0.0130)	(0.0124)	(0.0105)
Rescaled Gini	0.0662*	0.0457**	0.0458**	0.0457***	0.0430***	0.0311***
Production	(0.0352)	(0.0190)	(0.0209)	(0.0130)	(0.0127)	(0.0103)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.316	0.524	0.575	0.659	0.684	0.820
			*Using Raw O*NET Scores*			
Mean Cognitive	0.465***	0.192***	0.101**	0.0155	-0.00493	0.0112
(Raw Measure)	(0.0719)	(0.0428)	(0.0514)	(0.0354)	(0.0325)	(0.0216)
Mean Production	0.206***	0.207***	0.133**	0.112***	0.119***	0.0762***
(Raw Measure)	(0.0590)	(0.0355)	(0.0532)	(0.0304)	(0.0298)	(0.0223)
Gini Cognitive	0.0271	-0.0216	0.0423	-0.0416	-0.0490**	-0.0495***
(Raw Measure)	(0.0836)	(0.0499)	(0.0454)	(0.0270)	(0.0245)	(0.0173)
Gini Production	0.120	0.0874*	0.0802*	0.0795***	0.0834***	0.0641***
(Raw Measure)	(0.0860)	(0.0454)	(0.0473)	(0.0306)	(0.0305)	(0.0221)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.322	0.525	0.575	0.659	0.684	0.820
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

Table 2.11: CPS Regressions: All Skills

	(1) Log Wkly. Earn.	(2) Log Wkly. Earn.	(3) Log Wkly. Earn.	(4) Log Wkly. Earn.	(5) Log Wkly. Earn.
Mean Skill Pctile.	0.305***	0.178***	0.120***	0.107***	0.106***
	(0.0338)	(0.0237)	(0.0198)	(0.0146)	(0.0130)
Gini of All Skills	0.0348	0.0119	0.00939	0.00402	0.00412
	(0.0322)	(0.0229)	(0.0180)	(0.0135)	(0.0121)
Observations	173,126	173,126	173,126	173,126	173,126
Adjusted R-squared	0.188	0.300	0.321	0.338	0.345
Individual Controls	No	Yes	Yes	Yes	Yes
Occ. Group Controls	No	No	Yes	Yes	Yes
Industry Controls	None	None	None	Aggregate	Detailed
Num. Clusters	467	467	467	467	467

Robust standard errors in parentheses are clustered by Census occupation. From the unweighted CPS Merged Outgoing Rotation Groups 2012-2015. Individual Controls include quadratic terms in age and experience, race indicators, and educational group indicators. All specifications include year controls. *** p<0.01, ** p<0.05, * p<0.1

The Theil index is a popular form of generalized entropy index that, like the Gini coefficient, is frequently used to measure inequality. The standard deviation measures variation in the depth of skills, and it has a monotonic relationship with the Gini coefficient with a fixed mean on a bounded scale such as the one used here. These results suggest similar magnitudes for wage premia as the Gini coefficient measure of specialization. They also provide some stronger suggestive evidence that within establishments, specialization in Cognitive skills is associated with higher wages event though it is not associated with higher wages across establishments.

2.6.3 Wage Premia in CPS

As noted in Section 2.5.3, even though the use of RAIS data from Brazil allows me to look at the wage premia associated with specialization both within and across firms, the use of these data comes with additional limitations and empirical concerns. Therefore, in this subsection, I replicate the key baseline regressions of the previous subsection using the CPS Merged Outgoing Rotation Groups (MORG) from 2012-2015.

Tables 2.11 and 2.12 provide results for the specialization wage premia across all skills and within the two dimensions of Cognitive and Production skills, respectively.

Table 2.12: CPS Regressions: Cognitive and Production Skills

	(1) Log Wkly. Earn.	(2) Log Wkly. Earn.	(3) Log Wkly. Earn.	(4) Log Wkly. Earn.	(5) Log Wkly. Earn.
Mean Cognitive Skills	0.223***	0.118***	0.0761***	0.0856***	0.0851***
	(0.0358)	(0.0255)	(0.0207)	(0.0158)	(0.0140)
Mean Production Skills	0.0155	0.0621**	0.0117	0.00981	0.00301
	(0.0279)	(0.0267)	(0.0230)	(0.0181)	(0.0169)
Gini Cognitive Skills	-0.0724**	-0.0636***	-0.0584***	-0.0377***	-0.0392***
	(0.0281)	(0.0195)	(0.0170)	(0.0128)	(0.0111)
Gini Production Skills	0.00659	0.0188	0.0209	0.0144	0.0124
	(0.0337)	(0.0294)	(0.0183)	(0.0148)	(0.0131)
Observations	173,126	173,126	173,126	173,126	173,126
Adjusted R-squared	0.204	0.302	0.322	0.339	0.346
Individual Controls	No	Yes	Yes	Yes	Yes
Occ. Group Controls	No	No	Yes	Yes	Yes
Industry Controls	None	None	None	Aggregate	Detailed
Num. Clusters	467	467	467	467	467

Robust standard errors in parentheses are clustered by Census occupation. From the unweighted CPS Merged Outgoing Rotation Groups 2012-2015. Individual Controls include quadratic terms in age and experience, race indicators, and educational group indicators. All specifications include year controls. *** p<0.01, ** p<0.05, * p<0.1

These results correspond to Tables 2.6 and 2.7. As before, when looking at all skills simultaneously, there is no evidence that specialization is associated with either higher or lower wages once an occupation's mean level of reported skill is taken into account. However, once Production and Cognitive skills are separated, there is evidence of wage differences associated with specialization in each. The results of Table 2.12, however, are not entirely congruent with the results found when looking at RAIS as in Table 2.7. Specifically, there is no clear evidence of a wage premium associated with specialization within Production skills in these results. And, there appears to be a wage penalty associated with specialization within Cognitive skills, a finding which is only consistent with the RAIS regressions before the inclusion of broader occupational group fixed effects.

Table 2.13 breaks out the sample into blue-collar and white-collar occupations, as in Table 2.8, and sheds insight into the discrepancy from the baseline regressions. Just as was observed in the Brazilian context, among blue-collar occupations, there are large observed wage premia associated with specialization among Production skills. This is coupled, however, with a finding that specialization within Cognitive skills is associated with a wage penalty for these workers. Among white-collar workers, while

Table 2.13: CPS Regressions: White-Collar and Blue-Collar Workers

	(1) Log Wkly. Earn.	(2) Log Wkly. Earn.	(3) Log Wkly. Earn.	(4) Log Wkly. Earn.	(5) Log Wkly. Earn.
			White-Collar Workers Only		
Mean Cognitive Skills	0.265***	0.148***	0.128***	0.133***	0.123***
	(0.0543)	(0.0379)	(0.0294)	(0.0220)	(0.0189)
Mean Production Skills	-0.0420	0.0297	-0.0245	-0.0125	-0.0168
	(0.0308)	(0.0332)	(0.0290)	(0.0238)	(0.0220)
Gini Cognitive Skills	-0.0817	-0.0604*	-0.000651	0.0156	0.00143
	(0.0525)	(0.0356)	(0.0270)	(0.0185)	(0.0171)
Gini Production Skills	-0.0170	0.00730	-0.0103	-0.00604	-0.00584
	(0.0323)	(0.0306)	(0.0195)	(0.0170)	(0.0147)
Observations	110,718	110,718	110,718	110,718	110,718
Adjusted R-squared	0.167	0.289	0.313	0.327	0.334
Num. Clusters	301	301	301	301	301
			Blue-Collar Workers Only		
Mean Cognitive Skills	0.0708*	0.0428	0.0118	0.0517**	0.0523**
	(0.0424)	(0.0379)	(0.0390)	(0.0247)	(0.0240)
Mean Production Skills	0.135***	0.0928**	0.121***	0.0730**	0.0711**
	(0.0488)	(0.0404)	(0.0413)	(0.0309)	(0.0289)
Gini Cognitive Skills	-0.118***	-0.101***	-0.104***	-0.0661***	-0.0561***
	(0.0246)	(0.0203)	(0.0209)	(0.0137)	(0.0126)
Gini Production Skills	0.158***	0.0949*	0.140***	0.0775**	0.0775**
	(0.0598)	(0.0501)	(0.0528)	(0.0388)	(0.0348)
Observations	62,408	62,408	62,408	62,408	62,408
Adjusted R-squared	0.097	0.163	0.173	0.201	0.212
Num. Clusters	166	166	166	166	166
Individual Controls	No	Yes	Yes	Yes	Yes
Occ. Group Controls	No	No	Yes	Yes	Yes
Industry Controls	None	None	None	Aggregate	Detailed

Robust standard errors in parentheses are clustered by Census occupation. From the unweighted CPS Merged Outgoing Rotation Groups 2012-2015. Individual Controls include quadratic terms in age and experience, race indicators, and educational group indicators. All specifications include year controls. *** p<0.01, ** p<0.05, * p<0.1

Cognitive skills appear to be valued, there is no clear evidence of wage differences associated with specialization within either group of skills. Overall, these finding are fairly consistent with those obtained via RAIS.

In summary, while there are some inconsistencies between the regression results produced using the CPS and those produced using the RAIS data, these data continue to strongly support the notion that for blue-collar workers in particular, there is a sizeable wage premium associated with being specialized within one's Production skills. However, there remains weak evidence for any wage premium associated with specialization among white-collar workers who use primarily Cognitive skills.

2.6.4 Sorting and Occupational Switchers

The baseline results presented in the previous suggestion provide considerable evidence that specialization in one's set of Productions skills is valuable. This supports for a key prediction of the theoretical model of specialization based on underlying ability. However, the lack of clear evidence in support of the value of specialization among Cognitive skills raises questions as well. Although a lack of a wage premium associated with specialization of Cognitive skills could, for example, be the result of measurement error, or a lack of variation in skillsets within detailed occupational categories. However, the pattern of the results is also suggestive in particular that less specialized workers may be sorted into high-wage firms. Such a finding is at odds with the theoretical model of Section 2.3, which clearly suggests that more specialized workers will be located in more productive and higher-paying firms, because of the fundamental trade-off between the increasing returns to scale implied by specialization and the existence of coordination costs or other convex costs that limit the scope of that specialization.

To investigate this further, in this subsection, I examine patterns of worker sorting directly. As described in Section 2.4.2, I do so by constructing a two-step procedure based on the overlapping high-dimensional fixed effect methods of (Abowd, Kramarz, and Margolis, 1999). In the first step, I regress workers log wages on a full set of establishment and individual fixed effects along with individual covariates to account for life-cycle wage patterns. Then, I regress the estimated establishment fixed effects on the measures of mean skill and of skill specialization to determine the extent to which specialized workers are sorted into higher-wage firms, net of other observable characteristics.

Figure 2·8: Scatterplot of Mean AKM Establishment and Person FEs by Occupation

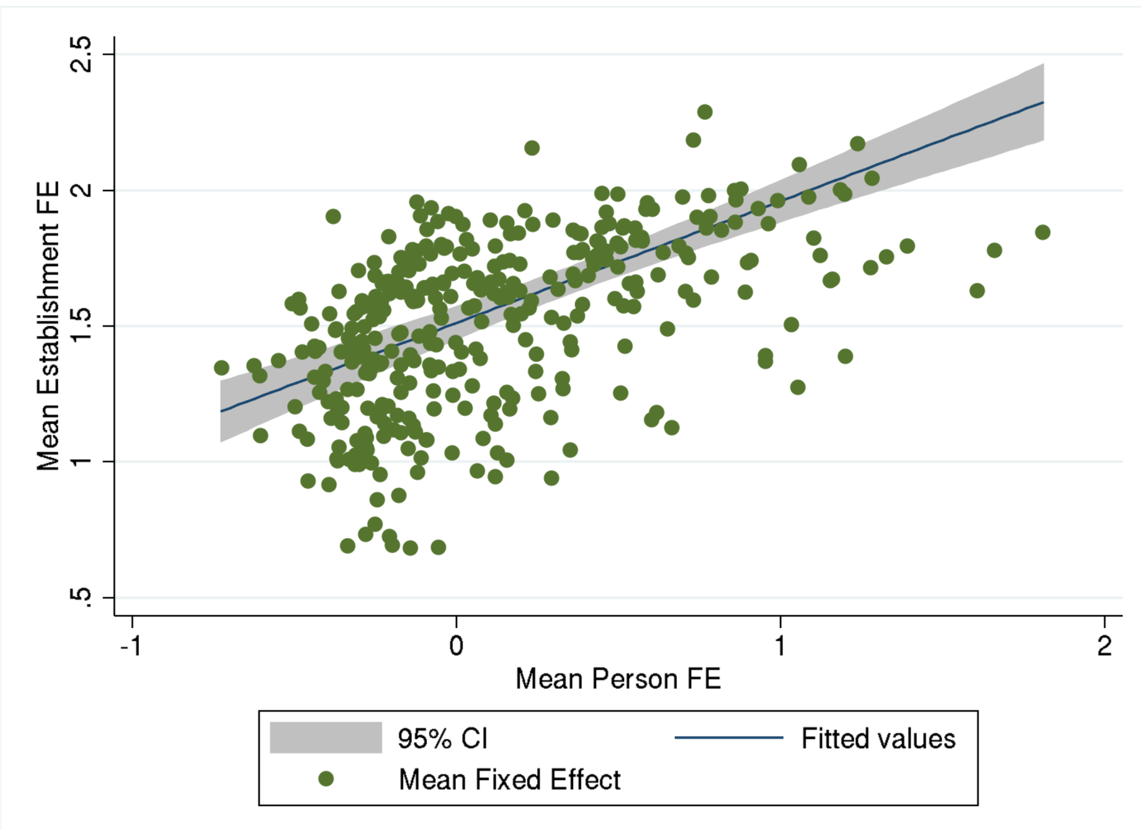

The initial AKM paper, along with many follow-up studies, have sought to analyze the extent of sorting between high-wage workers and high-wage firms in general, rather than along particular particular dimensions of workers' observable characteristics.

That is, they consider the correlation between estimated establishment effects and worker effects across the entire sample. This pattern of sorting can also be observed occupation-by-occupation. Figure 2·8 shows a scatterplot of the mean establishment FE and the mean worker FE identified by the AKM procedure for each occupation. As can be seen from the plot, there is evidence of positive assortative matching between workers and firms. The overall correlation between worker and establishment fixed effects is approximately 0.134. This is somewhat lower than the correlations found, for example, by Card et. al in West Germany over the same period (Card, Heining, and Kline, 2013), and slightly lower than the correlations of approximately 0.16 found between observable worker characteristic effects and establishment fixed effects in the Brazilian data for 1990 and 1997 in Menezes-Filho, Muendler, and Ramey (2008).

Table 2.14: Regression of AKM Establishment Fixed Effects on Skill Measures

	(1) Estab. FE	(2) Estab. FE	(3) Estab. FE	(4) Estab. FE
Mean Cognitive Skills	0.0748***	0.00327	-0.0338	-0.0547**
	(0.0279)	(0.0501)	(0.0344)	(0.0266)
Mean Production Skills	0.168***	0.239***	0.188***	0.0472
	(0.0583)	(0.0591)	(0.0422)	(0.0300)
Gini Cognitive Skills	-0.0776**	-0.0738**	-0.0628**	-0.0546**
	(0.0364)	(0.0347)	(0.0278)	(0.0228)
Gini Production Skills	0.117***	0.130***	0.0848***	0.0316*
	(0.0439)	(0.0399)	(0.0268)	(0.0178)
Observations	56,499,238	56,499,238	56,433,558	56,433,558
Adjusted R-squared	0.096	0.104	0.219	0.314
Occupation Controls	None	Aggregate	Aggregate	Two-Digit
Individual Controls	No	No	Yes	Yes

Notes: Robust standard errors are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specification include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure as well as education level and nationality indicators. *** p<0.01, ** p<0.05, * p<0.1

Table 2.14 presents the results from the second stage of the procedure outlined above. As can be seen from the table, workers who have specialized Production skillsets are more likely to be found at high-wage establishments, which are presumed to firms with access to better production technologies. While the inclusion of rich individual covariates and two-digit occupational indicators attenuates this finding somewhat, these findings are very consistent with the predictions of the standard model, and with the findings of the previous subsection. However, the same pattern does not hold among Cognitive skills. Rather, high-wage establishments appear to employ individuals in occupations that possess broad Cognitive skillsets. Again, this finding is robust to the inclusion of controls based on differences in observable individual characteristics, to the mean levels of skill in each skill cluster, and to the inclusion of broader occupational indicators.

Table 2.15: Skill Measure Wage Premia Identified from Occupation Switchers

	(1) Log Earnings	(2) Log Earnings
Mean Cognitive Skills	0.0448***	0.0441***
	(0.00452)	(0.00451)
Mean Production Skills	0.00570	0.00587
	(0.00474)	(0.00475)
Gini Cognitive Skills	-0.000141	-0.000243
	(0.00581)	(0.00588)
Gini Production Skills	-0.00394	-0.00371
	(0.00428)	(0.00434)
Observations	56,434,639	56,367,869
Adjusted R-squared	0.931	0.931
Individual Controls	No	Yes

Notes: Robust standard errors are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specification include individual and establishment fixed effects, education group by year fixed effects and quadratic and cubic terms in age - 40. Individual controls include linear and quadratic terms in potential experience, and job tenure. *** p<0.01, ** p<0.05, * p<0.1

Finally, in order to understand the mechanisms by which specialization and wages are related further, I run a single-step AKM regression in which I include establishment and worker fixed effects simultaneously, along with the occupational content measures as above. Rather than measuring worker sorting, this measures the extent to which the worker and establishment fixed effects are able to capture the sources of wage patterns associated with skill specialization in full. Table 2.15 shows the results of these regressions, both with only a limited set of life-cycle and educational controls included as in the standard AKM specification, and with additional tenure and experience controls included in Column 2. The results show that, as predicted, occupation switchers' wages do not appear to be increased or reduced when they switch from an occupation that is less skill specialized to one that is more skill specialized, and vice versa. Rather, these findings are consistent with the sources of wage premia from specialization being either firm-level technology or individual-level ability.

2.7 Conclusion

When Adam Smith wrote his famous chapter regarding a pin factory, it seems evident that he was thinking of the gains from specialization in traditional manufacturing and other similar activities. And, the evidence presented here suggests that specialization of skills in production is indeed valuable. A one standard deviation increase in an occupation's Gini coefficient of Production skills is associated with 4-6% higher earnings, even after controlling for observable individual-level differences, controlling for variation across occupational categories, and after controlling flexibly for establishment-level differences. The evidence also suggests that the earnings premium from being in a production skill-specialized occupation is even larger for blue-collar workers. These findings are consistent with the predictions of the simple model of occupational choice which I also present in this chapter.

But, the evidence presented here also suggests that our existing models of specialization may be too limited to account for the role that cognitive skills play in production. Not only is there an absence of clear evidence that specialization in cognitive skills is compensated, but patterns of worker sorting suggest that workers in occupations with broader cognitive skillsets are more likely to be located in higher-paying firms. Even if cognitive skills are highly complementary in production, which might lead to within-firm sorting of high-ability workers to broad cognitive skillsets, the Becker and Murphy model suggests that higher-wage firms should employ more specialized workers, but they do not.

Perhaps this empirical finding is unsurprising. Many types of labor that most people would consider to be highly skilled are not particularly specialized in their skillsets. Managers and CEOs who oversee broad swathes of production activity may immediately come to mind as requiring broad skills. However, many other types of labor are similarly highly skilled, yet incompletely specialized. Consider, for example, economists. A typical Economics department may have researchers in different fields, but it does not have one person who takes all the first-order conditions, one person who does all of the writing, or one person who gives all of the presentations. Indeed, as of 2012, the median number of authors on a published Economics paper was only 2.3 (Card and DellaVigna, 2013). Furthermore, according to the U.S. O*NET database, Economists have very broad skillsets; of 35 skills about which they were surveyed, Economists reported on average that 18 were "Important" or "Very Important" to the performance of their job. The economists surveyed by O*NET report that their job requires levels of skill that are above the 80th percentile of the occupational distribution in 13 of 35 skills measured (*O*NET Resource Center - O*NET 21.0 Production Database*). None of these observations map neatly to a model in which gains from specialization induce firms to hire workers that specialize to the maximum feasible extent.

How, then, should we think of the role of specialization among workers who use primarily cognitive skills? Perhaps broad cognitive skillsets among supervisory workers reduce coordination costs. Or, perhaps broad cognitive skillsets allow workers to solve difficult problems or to discover firm-specific innovations. The theoretical literature has presented a range of models that consider knowledge hierarchies, outside innovators, and other mechanisms that could have relevance for the observations of this chapter (Garicano and Rossi-Hansberg, 2006; Dessein and Santos, 2006; Legros, Newman, and Proto, 2013). But, absent empirical measures of specialization of the type developed in this chapter, there has to date been very little scope to test these models.

In interpreting these results, it is also important to have a careful understanding of their limitations. The O*NET data used in this chapter cannot, for example, be used to track the skill content of individuals, only occupations. Although I am able to include broader occupational group fixed effects, I cannot rule out the role of hedonic differences or other correlated supply-side factors that might influence overall wage patterns. Because the wage premia from specialization are intimately tied with worker sorting across firms, an ideal test of the model's predictions would require exogenous variation in total factor productivity that is not factor-augmenting, and this seems unlikely to find. Finally, although the Brazilian data used in this chapter are detailed administrative data which confer many benefits, the existence of a large informal sector in Brazil poses problems for any attempt to analyze changes across firms over time, and measurement error associated with mapping occupational codes may be a source of considerable attenuation bias.

Nonetheless, taken collectively, these results suggest that there is much more work to be done to understand the role of occupational skills within firms. Although one-dimensional models of human capital may be tractable and very useful in many contexts, they also have significant limitations. Additional research using models that consider human capital as a multidimensional object may yet provide considerable insights into the complex relationship between skills, specialization, and productivity.

Chapter 3

The Impact of Having a Job at Migration on Settlement Decisions: Ethnic Enclaves as Job Search Networks

3.1 Introduction

Ethnic enclaves are a persistent phenomenon in America and around the world. They persist over long periods of time, in spite of freedom of movement. They also persist in spite of observational evidence which suggests that residents of ethnic enclaves have worse labor market outcomes than non-residents. There is no guarantee that this observational evidence reflects any causal relationship between enclaves and immigrant outcomes, but it does suggest that enclaves are properly seen as a sorting phenomenon: Low-skilled new immigrants cluster in enclaves, while highly-skilled new immigrants more frequently choose to locate in areas where their ethnic group does not constitute a large proportion of the population. Since contemporary debates in immigration policy have focused heavily on issues relating to immigrant selection and immigrant networks, there may be great value in understanding the causes of this sorting within destination countries.

The persistence of enclaves suggests that new immigrants are responding to socioeconomic incentives when they choose where to locate. Yet, it remains difficult to separate among the plausible causes of enclave formation. Theoretical justifications for enclaves include discrimination, human capital externalities, and private benefits from avoiding assimilation (Edin, Fredriksson, and Åslund, 2003; Lazear, 1999; Cutler, Glaeser, and Vigdor, 2008). However, one of the most compelling hypotheses is that immigrants sort themselves to take advantage of network effects in the labor market. According to this line of reasoning, immigrants locate in enclaves so that they can use local networks to learn about and apply for jobs. This hypothesis is supported by several papers that suggest the existence of immigrant job search networks, particularly among refugees whose initial locations are randomly assigned Edin, Fredriksson, and Åslund (2003), Munshi (2003), and Beaman (2011). Yet, little empirical research exists to support the notion that individuals actually choose where to live with potential job search benefits in mind. Similarly, existing research cannot tell us whether the seeming patterns of negative selection in enclave residence are related to the fact that lower-skilled individuals may benefit disproportionately from any job search networks that exist.

In this chapter, I seek to determine whether immigrants' locational choices respond to potential job search networks by exploiting a simple discontinuity: immigrants who already have a job have little immediate need for the benefits of such a network. That is, I examine whether immigrants, and highly-skilled immigrants in particular, who have a job offer at the time of migration make different locational decisions than immigrants who do not already have a job offer. If new immigrants who arrive without a job offer are taking advantage of ethnic enclave networks for job search, then they will be more likely to locate in areas where the density of other individuals in the network is high. Meanwhile, immigrants who arrive with job offers may benefit less from these networks, and so these immigrants will be less likely to locate in enclave areas. I demonstrate these predictions with a simple theoretical model of job search that may occur both before and after migration.

For my empirical analysis, I use geographic data from the New Immigrant Survey

(Jasso et al., 2006), matched to aggregate data from the U.S. Census at the ZIP code level. The New Immigrant Survey (NIS) provides a cross-sectional survey of approximately 8,500 new legal immigrants to the United States, and it is designed to be a representative sample of all new legal immigrants. In addition to asking about employment at the time of migration, it also provides a detailed retrospective survey of pre-migration characteristics that allow me to control for selection on a broad range of observable attributes. Importantly, the NIS includes data on pre-migration salary, hours, occupation and industry that may control for selection on these attributes without the reverse causality issues that would be present if post-migration controls were used in their place.

My empirical analysis shows that immigrants who migrate with a job offer are significantly less likely to locate in neighborhoods with a high proportion of individuals in their language group. Even within a given metropolitan area, my analysis suggests that immigrants who migrate with a job offer locate in less ethnically-dense neighborhoods than those who migrate without a job offer. This result is robust to the inclusion of a variety of covariates, and also to several different methods of defining ethnic groups based on languages spoken at home. In keeping with the predictions of my theoretical model, this suggests that immigrants who arrive without job offers make a trade off between the job-search benefits of enclave membership and commuting costs.

While there do appear to be labor market benefits of locating in an ethnic enclave area, the benefits of this decision need not be homogeneous. Indeed, my model allows for rich selection behavior in locational decisions. Based on plausible assumptions about the nature of commuting costs and the job-matching function, I show that the job search benefits to higher-skilled immigrants of enclave residence may be lower than the job search benefits to lower-skilled immigrants. This would imply that upon receiving a job offer, high-skilled immigrants' locational decisions are impacted less than low-skilled immigrants' locational decisions. To test this hypothesis, I examine whether the impact of a pre-migration job offer is heterogeneous across the distribution of observable skills. However, my analysis fails to find significant evidence of heterogeneous effects, either on education or on other observable characteristics.

Also in keeping with the broader literature on selection in migration, my model suggests that immigrants who arrive with a job offer may be observably different from immigrants who do not arrive with a job offer. I find strong empirical evidence of this assertion. Immigrants who arrive with a job offer are more highly educated than immigrants who do not arrive with one, and they are considerably more highly paid both before and after migration. These differences suggest that selection on unobservables may be a concern. Any unobservable characteristics that are be correlated with the decision to locate in an ethnic enclave may bias regression results if they are also correlated with the probability of receiving a pre-migration job offer. As a robustness check, I perform an analysis using a kernel-based propensity score matching method. This analysis may be more robust to any potential impact of unobservables, as long as the correlation between unobservable characteristics and the probability of receiving a job offer is locally close to 0. The propensity score analysis is suggestive of a treatment effect; the magnitude of the estimated average treatment effect is similar to the results from OLS and probit regression, though the reduced power of this method means that these results are only marginally significant.

This chapter proceeds in sections as follows. Section 3.2 provides an overview of literature related to migration, and in particular to the role of ethnic enclaves. Next, Section 3.3 introduces a simple theoretical model that supports my hypotheses. This model illustrates that there is an inherent trade-off for new, unemployed immigrants in choosing where to locate that does not exist for immigrants who arrive with a job offer. It also provides a framework for incorporating ethnic enclaves into the broader literature on the selection of immigrants. Then, section 3.4 provides an empirical framework that allows me to test a key prediction of the theoretical model, that new immigrants who have jobs at the time of migration will be less likely to locate in ethnic enclaves that immigrants who do not have jobs at the time of migration. Section 3.5 describes the New Immigrant Survey and Census data that I use in this chapter, as well as the additional information that I intend to use once I have been granted access to the Restricted Use version of the Survey. Section 3.6 provides and discusses my main results, and section 3.7 provides a range of robustness checks, include a propensity score analysis. Finally, Section 3.8 concludes the chapter.

3.2 Overview of the Literature

Migration is an issue of considerable policy interest, and so there exists a large migration literature. Most of the literature focuses broadly on the issue of selection into migration, or on the impact of migration on native workers throughout the skill distribution (Borjas, 1987; Card, 2001; Borjas and Friedberg, 2009; Peri and Sparber, 2009). Many of these studies attempt to instrument for endogenous labor market sorting, typically under the assumption that migration patterns are persistent over time (Card, 2001). However, for this very reason, most of this literature can say little about the nature of this endogenous sorting, or about what mechanisms might cause migrants to sort in a persistent fashion. Apart from broad policies that encourage or discourage migration, there exist many related policy questions that may benefit from an understanding of these mechanisms. For example, to fully understand the impacts of U.S. policies that encourage family reunification, urban planning policies, or welfare programs, it is important to understand the nature of migrant social networks, particularly as they relate to ethnic enclaves.

There is evidence that network effects do matter for job search among the general population. The best such evidence comes from Bayer, Ross, and Topa (2008). These authors found that an urban resident is more likely to work with a resident of his/her own city block than on neighboring blocks. Since travelling a single city block to a job presumably imposes only very small commuting costs, this is supposed to identify potential network effects. Not only does the study find evidence of such a pattern, but it finds that this effect is stronger when residents have similar socio-economic characteristics. This supports the notion that ethnic enclaves form to take advantage of network effects, and it additionally suggests that enclaves may matter at a very localized level.

Among the immigrant population, Goel and Lang (2016) has proposed a theoretical model in which immigrants take advantage of their ethnic networks to reduce job search costs. Looking at a survey of new immigrants to Canada, they test the predictions of this model empirically by looking at whether new immigrants claim to have a friend or relative in the country at the time of immigration. They also use geographic measures of immigrant density to define ethnic enclaves. This model and results suggest that residents of ethnic enclaves may benefit from the more rapid arrival of job offers, even when the distribution of potential job offers to be received does not change. These predictions are consistent with my hypothesis. However, this theoretical model is based on a population of homogeneous immigrants, and so it is unable to make predictions about effects across different levels of worker skill.

There is also evidence that network effects matter in other decisions that migrants make. For example, Bertrand, Luttmer, and Mullainathan (2000) uses Census data to look at whether network effects matter in the application for welfare benefits. As in this chapter, they group migrants by common language, and they exploit variation in the density of migrants to conclude that network effects cause members of language groups with high existing welfare use to be more likely to apply for welfare benefits themselves. Furtado and Theodoropoulos (2012) uses a similar method to look at the take-up of disability programs. Similarly, research by Patel and Vella (2012) shows that the occupational choices of immigrant groups tend to be persistent over time. The authors show not only that new immigrants tend to choose the same occupations as previous migrants, but also that they earn a wage premium by doing so. This suggests that networks may be used as a way to receive better job offers than would otherwise be available.

While studies such those of Bertrand, Luttmer, and Mullainathan and Patel and Vella are able to employ careful instrumental variables strategies to achieve clean identification of network effects, their methods are not well-suited to considering initial locational decisions. Their methods can identify network effects only from within-metropolitan area variation in welfare use. In contrast, new immigrants choose where to locate not only within metropolitan areas, but across them as well, and there may be selection in these decisions. It is for this reason that I have adopted a selection on observables strategy, with careful control for pre-migration characteristics.

Another literature related to ethnic enclaves looks not at the causes of their formation, but at the outcomes of their residents. These studies frequently use the initial placement of refugee immigrants, as assigned by government policy, to address the issue of immigrant sorting. The work of Edin, Fredriksson, and Åslund (2003) used this method to look at the impact of living in an ethnic enclave on refugee immigrant earnings in Sweden. They found not only that immigrants benefit from initial placement in an ethnic enclave, but they found strong evidence of immigrant sorting after initial placement, particularly among low-skilled refugees. Work by Beaman (2011) on refugees assigned around the U.S. provides additional insight by suggesting that network effects may be heterogeneous. Refugees who were assigned to locations with a large social network of very recent immigrants showed lower incomes, while refugees who were assigned to locations with many longstanding network members appear to benefit from their assignment. These heterogeneous treatment effects support the labor market networks hypothesis, and they also suggest that the persistence of ethnic enclaves benefits new immigrants. However, these studies cannot show whether network effects are also important for non-refugees, or when locations are endogenously chosen.

Finally, looking solely at low-skilled agricultural migrants from Mexico, research by Munshi (2003) finds that migrants with larger networks are more likely to take higher-paying non-agricultural jobs. Because Munshi looks at repeat migrants, he is able to include individual fixed effects, which control effectively for selection on unobservables to cleanly identify network effects. However, it is unclear from this analysis whether such effects might be similar for more highly-skilled migrants, who typically come from countries other than Mexico, and who may be searching for very different types of jobs. It is also unclear whether effects may be similar in largely metropolitan ethnic enclaves.

Each of these existing studies seems to suggest that job search networks may play a significant role in location decisions. However, none of these studies directly addresses the situation of having a job offer at the time of migration. More importantly, none of these studies is able to address some key questions that may be relevant for policymaking. For example: Are the effects of arriving with a job heterogeneous across the skill distribution or across ethnic groups? Furthermore, do migrants with jobs locate in different metropolitan areas altogether, or in different parts of the same metropolitan areas? Could highly localized policies such as urban planning policies that impact the formation of dense enclaves have an impact on immigrant outcomes? For these questions, the detailed retrospective survey data that I use has considerable advantages.

3.3 Theoretical Model

In this section, I present a model that encapsulates the basic decisions made by potential legal immigrants, and that encompasses both selection in the decision to migrate and selection in the decision to locate in an ethic enclave. This model is a variant on the Roy (1951) model of selection, as applied to migration by Borjas (1987) and others. This model enriches the Borjas model by encompassing two separate problems; the pre-migration decision of the potential migrant, and the decision upon migration of where a migrant should locate. As in the seminal Roy model, the decisions of individuals are a function of individual country-specific shocks, which may be correlated.

Let us first begin with the problem of the potential migrant in his/her native country. In each period, each risk-neutral potential migrant i earns a wage $\mu_0 + \epsilon_{0i}$, where μ_0 is the average wage for a potential migrant in this country, and ϵ_{0i} is a time-invariant, country-specific individual wage shock that is distributed with mean 0 and finite variance σ_0^2.[1] All potential immigrants are assumed to be employed in their native countries. However, upon migration, an immigrant may be either employed or unemployed, and he/she chooses a location in the new country in which to reside. Locations vary by their ethnic density $D \in [0, D_{max}]$, i.e. the fraction of residents in that location that are immigrants. An individual immigrant is said to live in an

[1] In the Borjas model, these country-specific shocks are generally interpreted to be the residual wage differential net of the effects of observable characteristics, so that μ_0 is the mean wage for individuals conditional on observable characteristics. To simplify notation and to provide the model with a more intuitive interpretation, I assume that μ_0 is the unconditional mean, and therefore that ϵ_{0i} incorporates the effects of education and other observable characteristics, as well as unobserved heterogeneity in potential migrants. This also implies that the correlation between ϵ_{0i} and ϵ_{1i} is likely to be positive and relatively close to 1. In practice, ϵ_{0i} and ϵ_{1i} may be assumed to take an error components form, and the correlation of residual wage differences net of observables may be positive, negative, or zero.

ethnic enclave if the ethnic density of his/her location is greater than some level $\omega\bar{D}$, where $\bar{D}$ is the average density of that ethnic group in the new country. Once they have settled, immigrants cannot change their location.[2]

An employed immigrant earns a net wage in each period $\mu_1 + \epsilon_{1i} - \gamma(D)$, where μ_1 is the average wage for a migrant, ϵ_{1i} is time-invariant wage shock in the receiving (new) country, also mean-0 with finite variance σ_1^2, and $\gamma(D)$ is a commuting cost that is assumed to be increasing in D, the ethnic density of the location in which the immigrant resides.[3] Assume for simplicity and for later use that $\gamma(0) = 0$ and that $\gamma(\cdot)$ is twice continuously differentiable. Earnings of unemployed immigrants in each period are normalized to 0, and they do not depend on D. Immigrants and potential immigrants have complete information in this model, and there is no uncertainty, so each migrant knows the mean wage in each country, as well as his/her individual wage shocks ϵ_{0i} and ϵ_{1i}. Individuals in their native country may choose to migrate to the new country, and upon migration they choose D, the ethnic density of the location in which they choose to live.

In this model, both immigrants and non-immigrants may find receive job offers in the receiving country. However, in each period an immigrant who has already migrated to the receiving country is more likely to find a job than an immigrant who has not yet migrated. In this model, job offers arrive by a Poisson process, where the per-period arrival probability depends on that individual's country-specific shocks and, for unemployed immigrants, the ethnic density of their location. Thus, the probability that a non-immigrant finds a job is $\lambda_0(\epsilon_0, \epsilon_1)$, because it may depend on country-specific shocks in both his/her native country and in the new country.

[2] The basic predictions of this model will hold as long as relocation is costly and/or it is not instantaneous.

[3] This assumption simply imposes the requirement that locating in an ethnic enclave is costly for employed individuals. If the jobs performed by employed immigrants are located outside the ethnic enclaves themselves, then this assumption will be met.

The probability that an unemployed immigrant finds a job is $\lambda_1(\epsilon_1, D)$, because it may depend on his/her country-specific shock in the new country as well as the ethnic density of the location in which the immigrant has chosen to reside. If locating in an ethnic enclave helps a new immigrant to find a job, then $\lambda_1(\cdot)$ will be increasing in D.[4] I assume also that both $\lambda_0(\cdot)$ and $\lambda_1(\cdot)$ are increasing in ϵ_1, and that $\lambda_1(\cdot)$ is twice continuously differentiable.

Migration for immigrants who are not currently employed costs a fraction c of that immigrant's utility in his home county, while migration for immigrants who have a job offer at the time of migration is assumed to be costless.[5] Additionally, I will assume for simplicity that when a potential migrant receives a job offer, that he or she migrates with certainty. Formally, this requires that the supports of ϵ_{0i} and ϵ_{1i} are bounded above and below, respectively, and it requires some restrictions on the relative values of μ_0, μ_1, ϵ_{0i} and ϵ_{1i} which will also depend on the functions $\lambda_0(\cdot)$, $\lambda_1(\cdot)$, and $\gamma(\cdot)$. I assume that these restrictions are satisfied.

Define V_{0i} as the value function for a potential migrant i in his/her native country, V_{1Ui} as the value function for an individual i who has migrated but is unemployed, and V_{1Mi} as the value function for an individual i who has migrated and is employed. Then, the following equations characterize the model:

$$V_{0i} = \mu_0 + \epsilon_{0i} + \lambda_0(\epsilon_{0i}, \epsilon_{1i})\beta V'_{1Mi} + (1 - \lambda_0(\epsilon_{0i}, \epsilon_{1i}))\beta V'_{0i} \tag{3.1}$$

$$V_{1Ui} = 0 + \lambda_1(\epsilon_{1i}, D)\beta V'_{1Mi} + (1 - \lambda_1(\epsilon_{1i}, D))\beta V'_{1Ui} \tag{3.2}$$

$$V_{1Mi} = \mu_1 + \epsilon_{1i} - \gamma(D) + \beta V'_{1Mi} \tag{3.3}$$

[4] My model does not explicitly model either the decisions of firms or the informational content of social network ties, and therefore it does not incorporate the type of asymmetric information problem that would imply that $\frac{\partial\lambda(\cdot)}{\partial D} > 0$. However, this result can be easily obtained from a model with a structure similar to that of Montgomery (1991).

[5] The assumption that migration is costless for individuals with a job at the time of migration simplifies the notation of the problem, but it does not change the intuition of the results. It is also consistent with a world in which employers pay for the migration of the foreign workers they hire.

Since ϵ_{1i} is time-invariant, and since there is no job separation in this model, the value function for employed immigrants can be rewritten as:

$$V_{1Mi} = \frac{\mu_1 + \epsilon_{1i} - \gamma(D)}{1-\beta} \tag{3.4}$$

From (3.4), it is easy to see that immigrants who arrive with a job offer will seek to minimize their commuting costs $\gamma(D)$, and since $\gamma(\cdot)$ is assumed to be increasing in D, that all immigrants who arrive with job offers will choose $D = 0$. However, immigrants who arrive without job offers will choose D based on both the likelihood that they will find a job, and on the commuting costs that they will incur upon becoming employed.

Substituting the expression in (3.4) for V_{1Mi} into the equation for unemployed immigrants, we get that:

$$V_{1Ui} = \frac{\frac{\beta}{1-\beta}\lambda_1(\epsilon_{1i}, D)(\mu_1 + \epsilon_{1i} - \gamma(D))}{1-(1-\lambda_1(\epsilon_{1i}, D))\beta} \tag{3.5}$$

Similar substitution into the equation for non-immigrants implies that:

$$V_{0i} = \frac{\mu_0 + \epsilon_{0i}}{1-(1-\lambda_0(\epsilon_{0i}, \epsilon_{1i}))\beta} + \frac{\frac{\beta}{1-\beta}\lambda_0(\epsilon_{0i}, \epsilon_{1i})(\mu_1 + \epsilon_{1i} - \gamma(D))}{1-(1-\lambda_0(\epsilon_{0i}, \epsilon_{1i}))\beta} \tag{3.6}$$

A non-immigrant chooses to migrate without a job offer if the value of doing so exceeds the value of remaining in his/her native country, net of his/her migration costs $c \cdot V_{0i}$. Then, from the above, this model implies that migration will occur whenever:

$$\frac{\frac{\beta}{1-\beta}\lambda_1(\epsilon_{1i}, D)(\mu_1 + \epsilon_{1i} - \gamma(D))}{1-(1-\lambda_1(\epsilon_{1i}, D))\beta} - \frac{(1+c)(\mu_0 + \epsilon_{0i})}{1-(1-\lambda_0(\epsilon_{0i}, \epsilon_{1i}))\beta} - \frac{(1+c)\frac{\beta}{1-\beta}\lambda_0(\epsilon_{0i}, \epsilon_{1i})(\mu_1 + \epsilon_{1i} - \gamma(D))}{1-(1-\lambda_0(\epsilon_{0i}, \epsilon_{1i}))\beta} \geq 0 \tag{3.7}$$

Empirically, one cannot observe the decisions of individuals who have not chosen to migrate, because it is not in general possible to observe or identify the value of ϵ_{1i}

for anyone who has not migrated.[6] However, we do know that this constraint will be satisfied for any immigrant who is unemployed. Therefore, I will refer to (3.7) as the "Migration Constraint."

At the time of migration, a new immigrant who has migrated without a job offer chooses where to locate. Unlike a new immigrant who migrates with a job offer, this immigrant is not merely choosing to minimize commuting costs. Instead, his/her choice of D reflects an inherent trade-off between the probability of finding employment, and the cost that will be incurred to commute to that employment once it has been found. That is, the new immigrant solves the following problem:

$$\max_{D} \frac{\frac{\beta}{1-\beta}\lambda_1(\epsilon_1, D)(\mu_1 + \epsilon_1 - \gamma(D))}{1 - (1 - \lambda_1(\epsilon_1, D))\beta}, \text{ subject to} \tag{3.8}$$

$$\text{the Migration Constraint (3.7)}$$

$$0 \leq D \leq D_{max}$$

We can take the first order condition to solve for the optimal choice of D. I will abuse some notation by dropping all subscripts and arguments (ex. $\lambda \equiv \lambda_1(\epsilon_1, D)$). I will also take advantage of the conditions that $\beta \in (0, 1)$ and that $\lambda \geq 0$. Then, with some algebra, taking the first order condition yields the following:

$$\frac{\mu + \epsilon - \gamma}{\gamma'} = \frac{\lambda}{\lambda_D} \cdot \frac{1 - \beta + \beta\lambda}{1 - \beta} \tag{3.9}$$

where λ_D is the partial derivative of λ with respect to D.

From the above equations, the following features of the model are clear:

- *Locational decisions are not impacted* directly *by pre-migration characteristics.* This is true by construction. However, it is important to recognize that in

[6] Heckman and Honoré (1990) demonstrate that under the traditional assumption made in the Roy model that wage shocks are distributed log-normally, and assuming that the correlation of ϵ_0 and ϵ_1 is known, it is possible to identify the distribution of ϵ_0 for non-migrants. However, the log-normal distribution is not bounded above, which is a necessary assumption for this model.

general, since $Corr(\epsilon_0, \epsilon_1) \neq 0$, pre-migration characteristics will be correlated to observable locational outcomes. Additionally, pre-migration characteristics impact the likelihood of receiving a pre-migration job offer. Therefore, any pre-migration characteristics that are correlated to post-migration characteristics may still present endogeneity issues in empirical analysis.

- *All of the possible selection behaviors admissible in the Borjas (1987) model are also admissible in this model.* A key insight of Borjas selection model is that depending on the correlation of country-specific shocks, we may observe that immigrants are positively or negatively selected from the populations of their native countries, or that there may be "refugee" selection of individuals who are at the bottom of the native country's wage distribution but the top of the new country's wage distribution. All of these cases are also admissible in this model.[7]

- *Conditional on migration, the average wage of an immigrant's native country has no impact on his/her locational decision.* This follows naturally from the problem shown above, and from the assumption that ϵ_0 is mean-0. Taken literally, this implies that empirically, any observed correlation between the mean wage of a country and the locational decisions of immigrants who did not arrive with a job offer must result from differences in the distribution of ϵ_1 across countries. However, if the model is relaxed to allow for non-commuting benefits of ethnic density (see below), then it is possible that differences in the extent of non-commuting benefits across ethnic groups or across pre-migration income levels may lead to an observable correlation between average pre-migration wages and average locational decisions.

[7] In particular, the special case in which $\lambda_0(\epsilon_0, \epsilon_1) = 0 \; \forall \; \epsilon_0, \epsilon_1$, $\lambda_1(\epsilon_1, D) = 0 \; \forall \; \epsilon_1, D$, and $\gamma(D) = 0$ collapses neatly into the standard Borjas (1987) model form.

- *Immigrants who arrive with job offers will choose $D = 0$ regardless of their characteristics.* As discussed above, this is a relatively strong assumption of the model made by construction. If the function $\gamma(D)$ is taken more generally to specify "net costs of living in an area of ethnic density," and if the assumption of monotonicity is relaxed, then immigrants who arrive with job offers may choose $D > 0$. This model also does not account for the locational choices of firms, which likely make commuting costs non-monotonic in density below some threshold. However, as long as the non-commuting benefits of ethnic density are not increasing in ϵ_1 quickly relative to the commuting costs of living in an area of high ethnic density, we should observe that immigrants who arrive with job offers are less likely to live in an ethnic enclave than immigrants who do not arrive with job offers.

- *Under some conditions, an increase in ϵ_1 will lead to a decrease in the equilibrium ethnic density chosen for immigrants who migrate without a job offer.*

In this chapter, I will focus my attention on the hypothesis that immigrants who arrive with job offers are less likely to locate in an ethnic enclave than immigrants who do not, as well as the hypothesis that effects may be heterogeneous in observable skill.

3.4 Empirical Design

The theoretical model in Section 3.3 demonstrates that the if ethnic enclaves serve as job search networks for new immigrants, then the existence of these networks can be observed by observing the locational decisions of new immigrants. A key testable prediction of the model is that immigrants who migrate without a job offer are more likely to locate in an ethnic enclave than immigrants who arrive with a job offer. Therefore, I seek to test for the existence of job search networks by testing this prediction.

Consider an individual i of ethnic group k living in geographic area j. Then, I seek to estimate α_1 in a model of the form:

$$E_{ijk} = \Phi\left(\alpha_0 + \alpha_1 Offer_i + \alpha_2 X_i + \delta_k + \eta_j + \nu_i\right) \tag{3.10}$$

Where X_i are individual characteristics, E_{ijk} is a binary measure of whether migrant i of group k in area j settles in an ethnic enclave, $Offer_i$ is a measure of whether the immigrant had a job offer at the time of his/her migration, δ_k and η_j are group and area fixed effects, and ν_i is a stochastic error term. Ethnic enclaves are defined as areas of substantially greater immigrant density than the average density for that ethnic group.

More formally, define ethnic density as follows:

$$Density_{jk} = \frac{\frac{\text{Number of people from group } k \text{ in area } j}{\text{Total population in area } j}}{\frac{\text{Number of people from group } k}{\text{Total population in country}}} \tag{3.11}$$

Then, the discrete dependent variable measure of ethnic density is based on whether the ethnic density exceeds a multiple ω of the national average ethnic density for that group:

$$E_{jk} = \mathbf{1}\left(Density_{jk} \geq \omega \bar{D}_k\right) \tag{3.12}$$

with $\omega = 2$ in the baseline estimates. Each of these variables is constructed using Census geographic data. A more detailed description of the construction of these variables can be found in the Appendix in Table C.1 and Table C.2.

While a binary definition of ethnic enclaves has some intuitive advantages, there is some evidence to suggest that using a binary dependent variable specification ignores considerable heterogeneity in outcomes. Specifically, there is great heterogeneity in the extent to which different ethnic groups appear to sort themselves into enclaves (Toussaint-Comeau, 2008), which implies that the set of locations that meet a binary threshold E_{jk} based on a fraction of the national average density may be a more or less selected subset of locations for different ethnic groups. To address this concern, I also construct a continuous measure of ethnic density using the inverse hyperbolic sine transformation.[8]Define:

$$D_{jk} \quad = \quad \ln\left(Density_{jk} + \left(Density_{jk}^2 + 1\right)^{1/2}\right) \tag{3.13}$$

Then, I run OLS regressions, fitting the model:

$$D_{ijk} = \alpha_0 + \alpha_1 Offer_i + \alpha_2 X_i + \delta_k + \eta_j + \nu_i \tag{3.14}$$

As shown in Section 3.6, the results of the OLS and probit regressions are broadly similar. Additionally, as shown in Table 3.4, marginal effects are fairly consistent across a wide range of binary density thresholds.

3.5 Data

3.5.1 The New Immigrant Survey

The New Immigrant Survey (NIS), administered by by the Princeton University Office of Population Research, is designed to be a nationally-representative multi-cohort longitudinal study of new legal immigrants to the United States. The first cohort of immigrants were sampled using administrative data on new migrants from May to

[8]MacKinnon and Magee (1990) show that the inverse hyperbolic sine transformation has similar properties to the logarithmic transformation, except that it is defined at 0. However, my main results are not substantially affected by using the log of ethnic density instead, as Bertrand, Luttmer, and Mullainathan (2000) and others have done.

November of 2003, and a baseline survey was conducted between June 2003 and June 2004. In the initial cohort, 12,500 adults were sampled, along with 1,250 children. Of these, there are 8,573 completed interviews, for a response rate of 68.6%. Since the sample is designed to be nationally representative, it includes a sample from each of the top 85 MSAs, each of the top 38 counties, and a random sample of 10 other MSAs and 15 other counties. A follow-up survey of the first wave of immigrants has been conducted and recently released for research use, but data from this follow-up survey is not used in this analysis.[9]

Each adult immigrant is asked a series of questions about each job that he/she reports having. This includes questions about about occupation, salary, benefits, hours worked, etc. However, most critically for this analysis, I code the variable $Offer_i$ based on whether immigrant i responds yes to the question "Had you been offered this job before coming to the United States to live?" for any of his/her reported jobs. The survey also asks questions such as "Did you get this job with the help of a relative?" and "Do any of your relatives work for this business?" I use these measures as additional covariates in some regression specifications, as they are observable measures of potentially endogenous pre-existing networks. In the Appendix, Table C.1 provides some information about the questions that I use to code variables from the NIS.

A key feature of the NIS is that it provides an exceedingly broad retrospective survey. For example, in addition to data on current employment status, the NIS asks for detailed data on such things as pre-migration industries and occupations, pre-migration wages, prior migration history, education, language skills, and social networks. While the Public Use version of the NIS contains limited geographic information, the NIS collected data on immigrants' location at the time of the interview at the ZIP code level, and I match this data to U.S. Census data at the ZIP code

[9]See http://nis.princeton.edu/project.html and http://nis.princeton.edu/overview.html for further information.

level to calculate local ethnic density and identify ethnic enclaves.

The key empirical challenge of attempting to do this type of analysis with a representative sample of the population is that, in general, new immigrants are neither exogenously offered jobs, nor are they exogenously located in ethnic enclaves. Yet, any number of characteristics may be correlated with both of these measures, and these are a potential source of omitted variable bias. For observable characteristics, the detailed retrospective data contained in the NIS provides clear advantages, as it allows me to control for such factors as pre-migration wages. The effects of unobservable characteristics, however, remain of concern. Like other theoretical models of migration behavior, (see, e.g. Borjas, 1987; Borjas and Friedberg, 2009), my theoretical model shows that selection on unobservable characteristics is likely to occur. Yet, the instrumental variables methods in the immigration literature either cannot account for endogenous sorting across metropolitan areas at the time of migration (Bertrand, Luttmer, and Mullainathan, 2000), or their validity is limited to a small subset of the immigrant population (Munshi, 2003).

The selection on observables identification strategy of this chapter, then, rests on the breadth of the NIS dataset to account for variables that could otherwise be sources of omitted variable bias. In exchange for this limitation, this chapter provides potentially broader external validity and more precise geographic variation than previous research. As shown in Section 3.6, using NIS covariates I am able to control for pre-existing family networks, occupational/industry sorting, limited measures of cultural attachment, and differences in leisure preferences that yield differences in pre-migration hours or salaries. I have also undetertaken an analysis using the method of propensity score matching, which may be somewhat more robust to selection on unobservables (Angrist and Pischke, 2008). Further discussion of the potential benefits of this method, as well as results, follow in Section 3.7.

3.5.2 Data on Ethnic Density

Although the NIS sample is designed to be nationally representative, it is relatively small and not suitable for calculating the ethnic density of geographic areas. Therefore, in this chapter, I use Census Bureau data from the SF-3 sample of the 2000 Decennial Census to measure the ethnic density of each ZIP code. Specifically, I consider ethnic groups based on language spoken at home, which is available in both the NIS and in Census data. I then define both a continuous measure D_{jk} and a binary measure E_{jk} for each location-language pair based as described in Section 3.4. Table C.2 provides some summary information about the manner in which I define ethnic group membership and density, while Table C.3 provides a list of language groups included in the model, along with within-group counts.

3.6 Baseline Results

3.6.1 Initial Analysis

I begin my analysis of the new immigrants by looking at overall characteristics of the immigrants in the NIS, and in particular by looking for potential differences in the characteristics of immigrants who arrive with a job offer versus immigrants who do not arrive with such an offer. Table 3.1 presents summary statistics on a variety of covariates used in my regression specifications, including means and sample standard errors. As the table shows, immigrants who arrive with a job offer do in fact differ on a number of key characteristics. Most noticeably, they are substantially more likely to be highly educated; nearly 80% of immigrants who arrive with a job offer have at least a Bachelors degree, while only about 40% of immigrants who arrive without an offer have that much education.

Immigrants who arrive with jobs are also somewhat less likely to have gotten their current job with the help of a relative. This may suggest that immigrants who do not arrive with a job offer may be more reliant on informal network ties even before migration than immigrants who arrive with an offer, and therefore their decision to locate in an enclave may be the result of existing network ties rather than the potential of an enclave to provide them with new ties. I include this variable, as well as the variable on whether a relative works at the employer, as covariates in my model. The inclusion of these covariates may not make my model robust to the impact of pre-existing network ties connecting non-relatives. However, since the majority of visas granted to new legal immigrants in the United States are awarded based on some sort of family-related sponsorship, and yet my regressions do not find evidence that family relationships impact the locational decision, I argue that the impact of these networks on the likelihood of having a pre-migration offer is likely to be small.

Table 3.1 also shows summary information on income for the portion of the sample that reports income. All incomes have been converted by the NIS to 2003 U.S. dollars, using PPP where applicable. However, this leads to some individuals reporting very high or very low incomes. Therefore, individuals with incomes below \$100 and above \$1,000,000 have been removed from the sample (and they have been removed whenever income is a covariate elsewhere in my analysis as well).[10] Individuals who arrive with a job offer do have incomes that are considerably higher than those of individuals who arrive without a job offer, both pre-migration and post-migration. However, notably, the sample suggests that individuals who arrive with a job offer see, on average, their incomes increase by approximately 55% upon migration. The increase in income for individuals who arrive without a job offer, however, is not statistically significant.

In the Appendix, I provide additional tables showing additional frequency statistics on who receives a pre-migration job offer and who does not. In general, the likelihood of receiving a pre-migration offer is quite heterogeneous across occupations

[10] All reported incomes have been converted to annual incomes based on reported usual hour/week and usual weeks/year worked. A small number of immigrants report that they are paid daily or per-unit; these individuals have been dropped from the analysis of income.

Table 3.1: Summary Statistics for New Immigrants

		Offered Job Prior To Move			Differences	
	All	Yes	No	Missing	Yes - No	Yes - (No ∪ Miss)
General Characteristics						
Less Than HS	0.281	0.066	0.241	0.353	-0.176	-0.232
	(0.005)	(0.010)	(0.007)	(0.008)	(0.018)	(0.019)
High School	0.227	0.093	0.244	0.232	-0.151	-0.145
	(0.005)	(0.012)	(0.007)	(0.007)	(0.018)	(0.017)
Some College	0.094	0.053	0.109	0.086	-0.056	-0.044
	(0.003)	(0.009)	(0.005)	(0.004)	(0.013)	(0.012)
Bachelors Degree	0.243	0.485	0.234	0.215	0.251	0.260
	(0.005)	(0.020)	(0.007)	(0.006)	(0.019)	(0.018)
Graduate School	0.152	0.302	0.170	0.113	0.133	0.162
	(0.004)	(0.018)	(0.006)	(0.005)	(0.017)	(0.015)
Years of Schooling	12.698	15.997	13.325	11.581	2.672	3.558
	(0.056)	(0.149)	(0.077)	(0.084)	(0.202)	(0.211)
Years of Schooling in U.S.	0.795	0.361	1.157	0.501	-0.797	-0.469
	(0.024)	(0.051)	(0.041)	(0.030)	(0.105)	(0.092)
Female	0.518	0.376	0.405	0.649	-0.029	-0.153
	(0.005)	(0.019)	(0.008)	(0.008)	(0.021)	(0.021)
Year Born	1964	1967	1967	1960	-1	3
	(0.146)	(0.347)	(0.158)	(0.252)	(0.418)	(0.561)
Helped By Relative To Get Job	0.167	0.139	0.174	0.093	-0.035	-0.032
	(0.005)	(0.014)	(0.006)	(0.024)	(0.016)	(0.016)
Relative Works For Company	0.115	0.096	0.114	0.144	-0.018	-0.022
	(0.004)	(0.012)	(0.005)	(0.015)	(0.014)	(0.014)
Offered Job Prior To Move	0.138					
	(0.005)					
N	8,573	625	3,906	4,042		
Salary Information (Where Available)						
First Post-Migration Job	22,879	50,365	19,182	17,230	27,618	31,849
	(674)	(2268)	(707)	(1722)	(2137)	(1571)
	[3,085]	[399]	[2,155]	[531]		
Last Pre-Migration Job	20,418	32,486	18,264	19,485	14,222	13,639
	(831)	(3264)	(1066)	(1297)	(2686)	(2592)
	[3,040]	[350]	[1,404]	[1,286]		
Current	29,147	59,461	24,763	18,498	25,911	35,202
	(756)	(3,894)	(600)	(3,983)	(1,736)	(1,290)
	[3,553]	[470]	[2,966]	[117]		

Notes: Standard errors in parentheses. Counts for salary data in brackets. All salaries are reported in 2003 U.S. dollars; conversions are made in PPP terms where applicable. Salary averages and counts exclude individuals with reported annual incomes less than \$100 or greater than \$1,000,000.

and industries.

3.6.2 Primary Regression Analysis

Table 3.2: Primary regression results: Assignment based on density of first non-English language

	Binary Enclave Indicator (Probit)					
	(1)	(2)	(3)	(4)	(5)	(6)
Offer at Migration	-0.268**	-0.256**	-0.254**	-0.378***	-0.336***	-0.349***
	(0.127)	(0.123)	(0.125)	(0.139)	(0.129)	(0.127)
Years School		-0.0177**	-0.0181**	-0.0123	-0.0169	0.00264
		(0.00868)	(0.00808)	(0.0195)	(0.0213)	(0.0232)
Age		-0.000722	-0.000809	-0.00373	-0.00462	-0.00703**
		(0.00265)	(0.00256)	(0.00378)	(0.00299)	(0.00333)
Female		-0.0525	-0.0537	-0.0904	-0.107	-0.130
		(0.0428)	(0.0428)	(0.0832)	(0.107)	(0.109)
Married		-0.0194	-0.0200	-0.0965	-0.0793	-0.0766
		(0.0755)	(0.0717)	(0.136)	(0.149)	(0.162)
Separated		0.0838	0.0831	-0.0352	-0.120	-0.126
		(0.0620)	(0.0603)	(0.269)	(0.320)	(0.342)
Has Children		0.0421	0.0410	0.0295	0.0340	0.0600
		(0.0478)	(0.0479)	(0.0738)	(0.0870)	(0.0861)
Got Job Relative			0.00456	-0.00891	-0.0101	-0.0936
			(0.0913)	(0.110)	(0.0998)	(0.125)
Relative at Emp.			-0.0690	-0.0397	-0.0181	-0.0276
			(0.0578)	(0.153)	(0.137)	(0.128)
Log Pre-Migration Salary					-0.0269	-0.0232
					(0.0266)	(0.0309)
Pre-Migration Hours					0.00156	0.000852
					(0.00194)	(0.00228)
Poor English						0.497***
						(0.0918)
Religious						0.184*
						(0.0970)
Poor Health						-0.101
						(0.437)
Occ. & Industry	No	No	No	No	Yes	Yes
Pre-Migration Income Sample	No	No	No	Yes	Yes	Yes
Observations	3,389	3,372	3,372	1,219	1,219	1,219
Avg. Marginal Effect	-0.0889	-0.0847	-0.0839	-0.124	-0.105	-0.108
	(0.0427)	(0.0413)	(0.0420)	(0.0451)	(0.0406)	(0.0393)

Notes: Standard errors in parentheses, clustered at MSA × language-group. All regressions include a full set of MSA and language-group fixed effects. Log salary is measured in U.S. 2003 dollars.
*** p<0.01, ** p<0.05, * p<0.1

Table 3.2 and Table 3.3 show my initial regression results using the basic specifications shown in Equations 3.10 and 3.14, respectively Table 3.2 shows the results of probit regression of my binary measure E_{jk} on whether one received a pre-migration offer, as well as a variety of covariates. For reference, the estimated average marginal effect of the indicator $Offer_i$ is also shown at the bottom of the table. Table 3.3 shows

Table 3.3: Primary regression results (OLS Specification): Assignment based on density of first non-English language

	Continuous Enclave Indicator (OLS)					
	(1)	(2)	(3)	(4)	(5)	(6)
Offer at Migration	-0.241**	-0.236**	-0.235**	-0.276***	-0.204***	-0.211***
	(0.0995)	(0.0938)	(0.0957)	(0.0793)	(0.0622)	(0.0644)
Years School		-0.0168***	-0.0171***	-0.0162	-0.0206*	-0.00851
		(0.00391)	(0.00373)	(0.0107)	(0.0124)	(0.0134)
Age		8.66e-05	3.69e-05	0.00182	0.00219	0.000595
		(0.00174)	(0.00169)	(0.00412)	(0.00340)	(0.00288)
Female		-0.00809	-0.00892	-0.0103	-0.0440	-0.0533
				(0.0370)	(0.0576)	(0.0567)
Married		0.00645	0.00602	0.00748	0.0155	0.0196
		(0.0579)	(0.0574)	(0.103)	(0.115)	(0.118)
Separated		0.0980	0.0970	0.201	0.142	0.144
		(0.102)	(0.101)	(0.205)	(0.245)	(0.255)
Has Children		0.0784***	0.0776***	0.0587	0.0615	0.0725
		(0.0243)	(0.0248)	(0.0593)	(0.0672)	(0.0664)
Got Job Relative			0.00436	0.0321	0.0404	-0.0105
			(0.0617)	(0.0674)	(0.0747)	(0.0776)
Relative at Emp.			-0.0478	-0.0833	-0.0971	-0.103
			(0.0346)	(0.120)	(0.116)	(0.113)
Log Pre-Migration Salary					-0.0408***	-0.0355**
					(0.0134)	(0.0139)
Pre-Migration Hours					-0.00233*	-0.00282**
					(0.00123)	(0.00124)
Poor English						0.327***
						(0.0865)
Religious						0.0684
						(0.0863)
Poor Health						0.0300
						(0.218)
Occ. & Industry	No	No	No	No	Yes	Yes
Pre-Migration Income Sample	No	No	No	Yes	Yes	Yes
Observations	3,507	3,487	3,487	1,311	1,311	1,311
R-squared	0.337	0.343	0.343	0.392	0.417	0.427

Notes: Standard errors in parentheses, clustered at MSA × language-group. All regressions include a full set of MSA and language-group fixed effects. Log salary is measured in U.S. 2003 dollars.
*** p<0.01, ** p<0.05, * p<0.1

OLS regression results on my continuous measure of enclave density D_{jk}, with similar covariate specifications. For each of these regressions, language group k is assigned using the first non-English language reported to be spoken at home, though the use of the other measures described in Table C.2 gives me very similar results. All of these regressions include a full set of language group and geographic area fixed effects, and they are two-way clustered at the MSA by language-group level to be robust to correlated shocks within language groups within an MSA. The estimated effect of having a job offer at the time of migration is consistently negative and significant, implying that immigrants who have received a job offer are in fact less likely to locate in an ethnic enclave. Additionally, the estimated magnitude of this effect varies little when one controls for standard demographic covariates and for measures of family-related job networks. Notably, if one controls for pre-migration salary, hours, occupation and industry as shown in columns 5 and 6, the magnitude of the effect does not diminish at all, and as shown in column 4, this does not appear to be driven by the limited sample for which pre-migration employment information is available. This suggests that the apparent effect of a job offer is unlikely to be driven solely by selection on the skill levels of those that receive a job offer. These regressions suggest that a new immigrant who arrives with a job offer is roughly 8.9% to 10.8% less likely to locate in an ethnic enclave than an immigrant who arrives without a job offer.

These regression results also show that there is a strong negative relationship between one's level of schooling and residence in an enclave, even after controlling for job offers and other covariates. Additionally, individuals who have children appear to be significantly more likely to locate in an enclave.[11] Individuals who report that they speak English or understand English at a poor level are considerably more likely to locate in enclaves. Finally, individuals who report relatively low pre-migration

[11] This binary measure indicates whether the respondent reports that they live with one or more of their own children under the age of 18.

salaries and low pre-migration average hours of work are significantly more likely to locate in an enclave. These coefficients are all consistent with other explanations that have been proposed for why immigrants locate in enclaves. In particular, they suggest that the ability to assimilate, leisure preferences, or other preferences may all play a substantial role in new immigrants' locational decisions. Yet, the significance and robustness of the effect of a job offer suggests that job market network effects are important as well.

Finally, it is noteworthy in my results in Tables 3.2 and 3.3 that my measures of family-based networks, whether one got a job from his/her relative and whether one's relative works for the company, do not appear to be significantly related to the decision to locate in an ethnic enclave once other covariates and the reception of a pre-migration job offer have been controlled for. Although I am unable to control for pre-existing non-familial networks as a source of selection on who receives a job offer, this suggests that unobserved networks may be unlikely to a major source of bias in these estimates.

The probit regression results shown in Table 3.2 suggest sizeable marginal effects of having a job offer on the probability of locating in an ethnic enclave, where an enclave is defined as having a population of that group that is twice the national average proportion. In Table 3.4, I show the marginal effects from probit regressions based on a variety of alternative thresholds. The results show that the marginal effects of a job offer are fairly similar even for comparatively restrictive definitions of what constitutes an enclave area. Additionally, the reported dependent variable means make it clear that immigrants are in fact highly geographically concentrated in the U.S.; more than one third of immigrants in the sample live in ZIP codes that have at least five times the national average proportion of their language group, and more than a quarter are in ZIP codes that have ten times the national average proportion

Table 3.4: Marginal Effects in Probit Regressions with Different Enclave Thresholds

	(1)	(2)	(3)	(4)	(5)
			Language Density $\geq 1\times$ *National Avg.*		
Marginal Effect	-0.0606	-0.0573	-0.0570	-0.0734	-0.0753
	(0.0256)	(0.0239)	(0.0246)	(0.0214)	(0.0222)
Mean Dep. Var.	0.790	0.790	0.790	0.785	0.785
			Language Density $\geq 2\times$ *National Avg.*		
Marginal Effect	-0.0889	-0.0847	-0.0839	-0.105	-0.108
	(0.0427)	(0.0413)	(0.0420)	(0.0406)	(0.0393)
Mean Dep. Var.	0.624	0.623	0.623	0.605	0.605
			Language Density $\geq 3\times$ *National Avg.*		
Marginal Effect	-0.0852	-0.0833	-0.0819	-0.0763	-0.0803
	(0.0338)	(0.0321)	(0.0331)	(0.0313)	(0.0298)
Mean Dep. Var.	0.502	0.503	0.503	0.480	0.480
			Language Density $\geq 4\times$ *National Avg.*		
Marginal Effect	-0.0707	-0.0699	-0.0698	-0.0502	-0.0525
	(0.0356)	(0.0343)	(0.0346)	(0.0377)	(0.0360)
Mean Dep. Var.	0.418	0.419	0.419	0.408	0.408
			Language Density $\geq 5\times$ *National Avg.*		
Marginal Effect	-0.0665	-0.0641	-0.0645	-0.0385	-0.0413
	(0.0353)	(0.0345)	(0.0351)	(0.0310)	(0.0304)
Mean Dep. Var.	0.341	0.342	0.342	0.358	0.358
			Language Density $\geq 10\times$ *National Avg.*		
Marginal Effect	-0.0455	-0.0428	-0.0433	-0.0486	-0.0527
	(0.0338)	(0.0280)	(0.0276)	(0.0353)	(0.0344)
Mean Dep. Var.	0.269	0.269	0.269	0.285	0.285
Demographics	No	Yes	Yes	Yes	Yes
Family/Relatives	No	No	Yes	Yes	Yes
Occ. & Industry	No	No	No	Yes	Yes
Pre-Migration Salary/Hours	No	No	No	Yes	Yes
Social Indicators	No	No	No	No	Yes

Notes: Robust standard errors in parentheses. All regressions include a full set of state/Census division and language-group fixed effects. Demographic controls include age, gender, years of schooling, and marital/parental status indicators. Family controls include indicators for whether an individual got their job from a relative or has a relative at their employer. Log Salary is measured in U.S. 2003 dollars. Social indicators include self-reported poor English language ability, poor health, and religiosity.
*** p<0.01, ** p<0.05, * p<0.1

of their group.

Table 3.5: Primary regression (OLS Specification) Based on Within-MSA Variation Only

	Continuous Language Density (OLS)				
	(1)	(2)	(3)	(4)	(5)
Offer at Migration	-0.110*	-0.113*	-0.111*	-0.101	-0.106
	(0.0627)	(0.0618)	(0.0616)	(0.108)	(0.109)
Demographics	No	Yes	Yes	Yes	Yes
Family/Relatives	No	No	Yes	Yes	Yes
Occ. & Industry	No	No	No	Yes	Yes
Pre-Migration Salary/Hours	No	No	No	Yes	Yes
Social Indicators	No	No	No	No	Yes
Observations	3,507	3,487	3,487	1,311	1,311
R-squared	0.557	0.561	0.561	0.629	0.637

Notes: Standard errors in parentheses, clustered at MSA $\times$ language-group (one-way clustering). All regressions include a full set of MSA, language-group, and MSA $\times$ language-group fixed effects. Demographic controls include age, gender, years of schooling, and marital/parental status indicators. Family controls include indicators for whether an individual got their job from a relative or has a relative at their employer. Log Salary is measured in U.S. 2003 dollars. Social indicators include self-reported poor English language ability, poor health, and religiosity.
*** p<0.01, ** p<0.05, * p<0.1

Although the theoretical model shown in Section 3.3 describes the locational choice using a single ethnic density measure, in the real world, locational decisions are slightly more complicated. In practice, one can think of an immigrant's decision of where to locate as having two main components: "In which metropolitan area should I locate?" and "Where within this metropolitan area should I locate?" These choices are made simultaneously. However, if individuals choose where to migrate based on labor market networks, then one should be able to empirically identify that individuals with a job offer locate in less ethnically-dense neighborhoods, even within the same metropolitan area. Table 3.5 repeats the baseline OLS regressions, but instead of including a full set of language group fixed effects plus MSA effects, this analysis includes a fixed effect for each language-group $\times$ MSA combination. Thus, these regressions are identified from only the within-MSA variation in which neighborhoods individuals choose. These results are clustered one-way at the MSA $\times$ language-group level.

As Table 3.5 shows, even looking within metropolitan areas, individuals who arrived in the U.S. with a job offer locate in less ethnically dense areas. However, the magnitude of these estimates is also smaller than the corresponding baseline estimates in Table 3.3, because some portion of the effect of a job offer is the result of those who arrive with a job offer choosing to live in different metropolitan areas altogether from those who arrive without an offer. The coefficients on all other covariates are also broadly similar to the baseline regressions. Overall, this provides additional supportive evidence that labor market effects are a significant factor in immigrants' locational decisions.

Table 3.6: Regression with Heterogeneous Effects Based on Schooling

	(1)	(2)	(3)	(4)	(5)
Panel A: Binary Enclave Indicator (Probit)					
Offer at Migration	-0.275**	-0.279**	-0.277**	-0.413***	-0.417***
	(0.116)	(0.113)	(0.115)	(0.135)	(0.143)
Offer × School (Demeaned)	0.0125	0.0121	0.0123	0.0386	0.0343
	(0.0197)	(0.0198)	(0.0189)	(0.0338)	(0.0341)
Years School (Lang. Group Demeaned)	-0.0186**	-0.0187**	-0.0191***	-0.0235	-0.00343
	(0.00750)	(0.00778)	(0.00738)	(0.0261)	(0.0281)
Observations	3,381	3,372	3,372	1,219	1,219
Avg. Marginal Effect	-0.0870	-0.0884	-0.0876	-0.110	-0.112
	(0.0403)	(0.0395)	(0.0401)	(0.0391)	(0.0387)
Panel B: Continuous Language Density (OLS)					
Offer at Migration	-0.224**	-0.233**	-0.232**	-0.234**	-0.234**
	(0.105)	(0.102)	(0.103)	(0.0950)	(0.0963)
Offer × School (Demeaned)	-0.000946	-0.00138	-0.00138	0.0146	0.0115
	(0.0189)	(0.0188)	(0.0181)	(0.0261)	(0.0264)
Years School (Lang. Group Demeaned)	-0.0171***	-0.0167***	-0.0170***	-0.0230*	-0.0104
	(0.00363)	(0.00382)	(0.00384)	(0.0124)	(0.0136)
Observations	3,498	3,487	3,487	1,311	1,311
R-squared	0.342	0.343	0.343	0.418	0.427
Demographics	No	Yes	Yes	Yes	Yes
Family/Relatives	No	No	Yes	Yes	Yes
Occ. & Industry	No	No	No	Yes	Yes
Pre-Migration Salary/Hours	No	No	No	Yes	Yes
Social Indicators	No	No	No	No	Yes

Notes: Standard errors in parentheses, clustered at MSA × language-group. All regressions include a full set of MSA and language-group fixed effects. Demographic controls include age, gender, years of schooling, and marital/parental status indicators. Family controls include indicators for whether an individual got their job from a relative or has a relative at their employer. Log Salary is measured in U.S. 2003 dollars. Social indicators include self-reported poor English language ability, poor health, and religiosity.Years of Schooling is actual years of schooling less the mean for each individual's language group.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Next, we will consider the possibility of heterogeneous effects in migration. The theoretical model suggests that individuals who have a high observable type in the receiving country will be less likely to locate in an ethnic enclave because they receive job offers at a relatively higher rate regardless of their location. While the baseline regression results show that more educated individuals are less likely to locate in an enclave, they do not allow for the possibility of heterogeneous effects. Table 3.6 repeats the baseline specification using both the binary enclave indicator and the continuous measure of ethnic density, but it includes an interaction between the reception of a pre-migration job offer and an individual's years of schooling.[12] As discussed in Section 3.3, under some plausible assumptions, the effect of a job offer on the location decision will be smaller for individuals with more education, because high-skilled immigrants who arrive without a job offer will receive an offer more quickly on average than low-skilled immigrants who arrive without a job offer. However, my empirical results do not appear to support this prediction. The interaction effect of an additional year of schooling is not statistically significant, which suggests that heterogeneity is not present.[13] Yet, when this interaction effect is included, my estimate of the average marginal effect of having an offer on an immigrant with an average level of education is virtually unchanged; an immigrant who receives a job offer before migrating is roughly 8.7% to 11.2% less likely to locate in an ethnic enclave.

[12]Here, each individual's years of schooling is subtracted from the mean for his/her language group. Results are similar whether one uses this measure, a "raw" measure of years of schooling, or a set of discrete indicators for different levels of educational attainment.

[13]In other regressions which are not shown here, I have tested for heterogeneous effects using interactions with English-language ability, as well using interactions with the full set household characteristics included in the baseline regressions. I find no significant evidence of heterogeneous effects along any of these dimensions.

3.7 Robustness Checks

The results shown in Tables 3.2 and 3.3 are highly suggestive, but there are a number of potential concerns with the interpretation of these results. For one, immigrants who arrive to the United States legally and who already have a job offer at the time of migration are likely to have arrived through an employer-sponsored visa. The process by which one receives this visa is different than the process for other visa classifications because employers typically have to show that an immigrant provides particular skills that are in demand and not that are otherwise available to the firm. In fact, the majority of immigrants who receive new legal resident status receive it due to family-sponsored preferences or because they are the direct relatives of existing U.S. citizens. (*2011 Yearbook of Immigration Statistics* 2012) To the extent that these employer-sponsored individuals may be different from non-sponsored individuals in unobservable ways (such as unobserved differences in underlying preferences), they may bias our estimates of labor market effects.

Table 3.7: Frequencies of Offer At Migration vs. Employee Sponsorship in the NIS

Had at Migration	*Employee Sponsorship*				
	No		Yes		
	#	% of Row	#	% of Row	Total
No	2,668	68.3%	1,238	31.7%	3,906
Yes	89	14.2%	536	85.8%	625
Total	2,757	60.8%	1,774	39.2%	4,531

Source: New Immigrant Survey.

One potential solution to this issue would be to run regressions on a sample consisting of only those individuals who were not employer-sponsored. Then, the effect of a job offer would reflect the impact of a job offer only for those who are not observably. However, as Table 3.7 shows, the NIS is too small to have much power in conducting such an analysis. Over 85% of individuals who arrived in the U.S. with a job offer are sponsored by their employers, and only 89 individuals in the sample had an offer at migration but were not employer sponsored. Accordingly, the baseline results are insignificant when run on this subsample.

Table 3.8: Regression with Interaction Based on Employer Sponsorship (Probit)

	(1)	(2)	(3)	(4)	(5)
	Binary Enclave Indicator (Probit)				
Employer Sponsored	-0.0819	-0.0722	-0.114*	-0.295***	-0.258**
	(0.0519)	(0.0514)	(0.0683)	(0.109)	(0.123)
Years School		-0.0162**	-0.0176**	-0.0177	0.00111
		(0.00757)	(0.00866)	(0.0217)	(0.0233)
Age		-0.00139	-0.000818	-0.00458*	-0.00673**
		(0.000906)	(0.00267)	(0.00264)	(0.00303)
Female		-0.0599	-0.0612	-0.131	-0.154
		(0.0555)	(0.0439)	(0.114)	(0.117)
Married		-0.0195	-0.00684	-0.0932	-0.0894
		(0.0448)	(0.0705)	(0.145)	(0.158)
Separated		0.0191	0.119	-0.105	-0.106
		(0.0392)	(0.0788)	(0.321)	(0.345)
Has Children		0.0465***	0.0333	0.0312	0.0526
		(0.0164)	(0.0474)	(0.0903)	(0.0907)
Got Job Relative			0.00254	-0.0425	-0.129
			(0.0849)	(0.0968)	(0.123)
Relative at Emp.			-0.0855*	-0.0490	-0.0492
			(0.0505)	(0.135)	(0.125)
Log Pre-Migration Salary				-0.0276	-0.0234
				(0.0250)	(0.0292)
Pre-Migration Hours				0.000606	-0.000100
				(0.00199)	(0.00239)
Occ. & Industry	No	No	No	Yes	Yes
Demographics	No	Yes	Yes	Yes	Yes
Family/Relatives	No	No	Yes	Yes	Yes
Pre-Migration Salary/Hours	No	No	No	Yes	Yes
Social Indicators	No	No	No	No	Yes
Observations	6,388	6,349	3,413	1,231	1,231
Avg. Marginal Effect	-0.0264	-0.0232	-0.0369	-0.0911	-0.0784
	(0.0168)	(0.0167)	(0.0227)	(0.0334)	(0.0375)

Notes: Standard errors in parentheses, clustered at state/Census division × language-group. All regressions include a full set of state/Census division and language-group fixed effects. Demographic controls include age, gender, years of schooling, and marital/parental status indicators. Family controls include indicators for whether an individual got their job from a relative or has a relative at their employer. Log Salary is measured in U.S. 2003 dollars. Social indicators include self-reported poor English language ability, poor health, and religiosity.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

An alternative strategy, which I show in Table 3.8, is to re-run the baseline regressions with an indicator for employer sponsorship instead of an indicator for receiving a job offer. These results are very similar to the baseline regression results. In particular, the effects on other covariates change little relative to the baseline analysis, which we might expect if employer sponsorship is highly correlated with these other covariates. Thus, while it remains possible that employer-sponsored individuals are systematically different from non-sponsored individuals, there is no strong evidence to suggest that these individuals are a source of bias in the baseline estimates.

3.7.1 Propensity Score Analysis

As discussed above, the types of individuals who receive job offers are on average more highly educated and with higher incomes than individuals who do not receive offers. Large differences in average observable covariates may complicate the usual concerns about endogeneity in regression results. For example, when the true effects of observable covariates on an outcome are nonlinear, then the endogeneity bias from omitted variables that are correlated with these covariates will be magnified.

The method of propensity score matching is a close relative of traditional regression analysis that may be beneficial in situations where these concerns exist. More properly, OLS can be understood as a matching estimator with a particular weighting scheme (Angrist and Pischke, 2008). In essence, propensity score matching creates a unique set of control observations for each treatment observation (treatment being defined in this case as the receipt of a job offer before migration). This puts the highest weight in analysis on individuals with the highest probability of treatment. Ordinary Least Squares, in contrast, puts the highest weight on individuals with the highest variance in regressors.

Some analysis, especially in the literature on job training programs, has suggested that the method of propensity score matching may be beneficial in situations where the broader population differs from the treatment group (Dehejia and Wahba, 2002). In the job training case, endogeneity bias arises from the fact that individuals seeks job training are disproportionately likely to have received an unobservable negative income shock in the recent past. In my model, the potential for endogeneity arises from the possibility that pre-migration immigrant networks are simply less strong for the types of immigrants who tend to receive pre-migration job offers, or that immigrants who receive pre-migration job offers have systematically different preferences from those who do not. Thus, in this subsection I undertake a preliminary application of propensity score matching estimators to identify the average treatment effect of receiving a job offer prior to migration on the decision to locate in an ethnic enclave.

A key requirement of propensity score matching methods is that there be common support over the distributions of covariates. If the assumption of common support is violated, then the set of suitable control observations for each treatment observation will be small or nonexistent, causing these observations to be dropped from the analysis. Figure 3·1 and Figure 3·2 provide kernel density estimates of the pre-migration salaries and total current salaries of individuals who received a pre-migration offer and those who did not. Figure 3·3 provide a dual bar chart showing the density of the distribution of years of schooling for individuals who received a pre-migration offer and those who did not. As expected from the results in Table 3.1, the distributions of these variables are not identical, and individuals who received a pre-migration offer have more schooling and higher salaries on average. However, it appears from these graphs that the assumption of common support is generally met.

Figures 3·4 and 3·5 show histograms of the propensity scores produced by my initial probit regressions of receiving a pre-migration job offer or having an employer sponsored visa on the covariates reported in Table 3.2, as well as the square of the continuous variables years of school, age, and pre-migration income.[14] Again, it does appear that the assumption of common support is met, but the distribution

[14] The inclusion of these polynomial terms is suggested by Angrist and Pischke (2008) and employed by Dehejia and Wahba (2002).

Figure 3·1: Kernel Density Estimation of Pre-Migration Salary

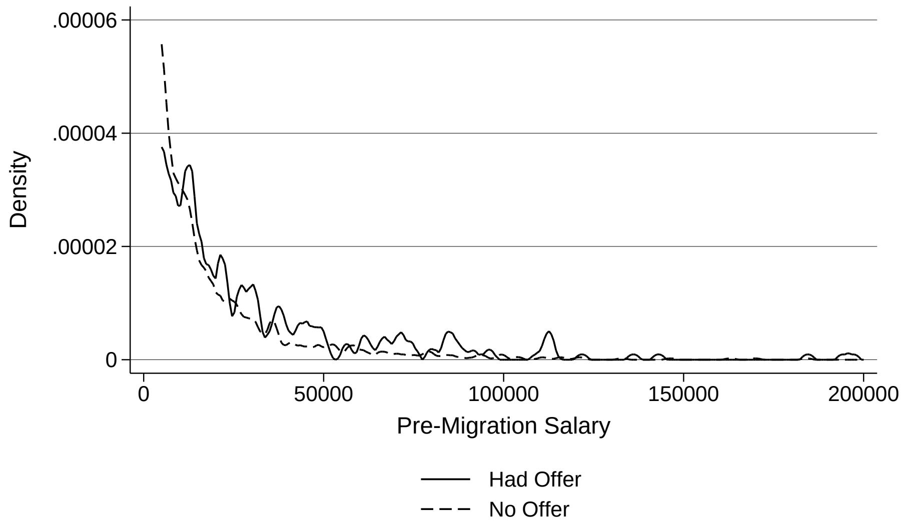

Kernel density plot using an Epanechnikov kernel with bandwidth of $1,000. Excludes individuals earning less than $5,000 or more than $200,000.

Figure 3·2: Kernel Density Estimation of Current Salary

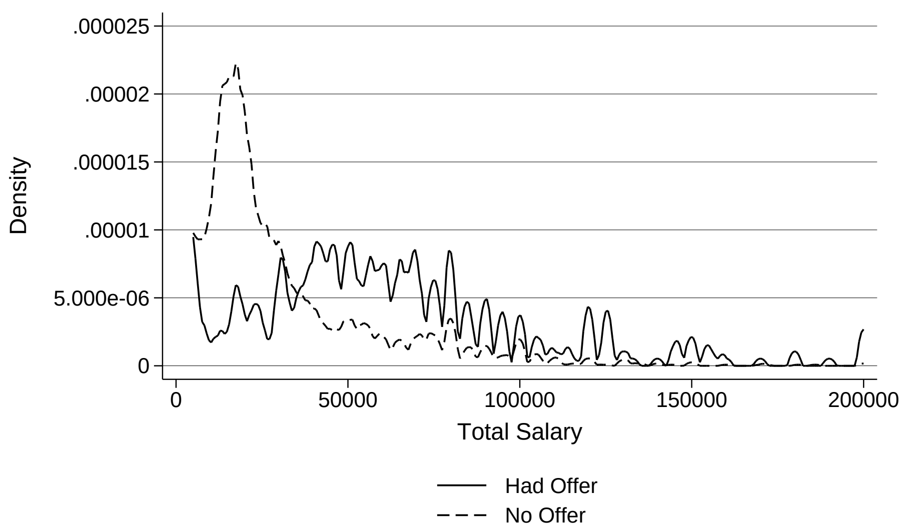

Kernel density plot using an Epanechnikov kernel with bandwidth of $1,000. Excludes individuals earning less than $5,000 or more than $200,000.

Figure 3·3: Distribution of Years of Schooling

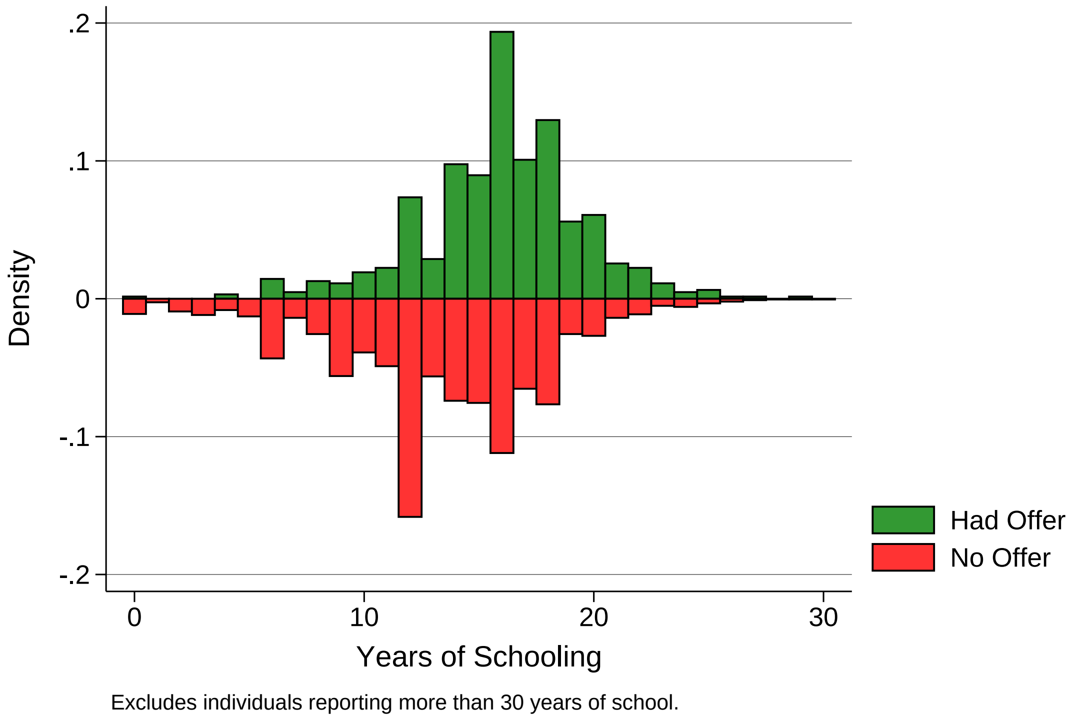

of propensity scores for individuals who did not receive a job offer is comparatively skewed to the less in both cases and is relatively less dense at high propensity scores. While this does suggest that the method of propensity score matching may lead to improved estimates in this case, it also suggests that the power of propensity score estimates may be relatively low.

Table 3.9 shows the results of propensity score matching methodology, using an Epanechnikov kernel matching method with default bandwidth assumptions. Because the standard errors produced by propensity score estimates make strong assumptions about homoskedasticity, I provide bootstrapped standard errors for these propensity score estimates. In general, my preliminary propensity score estimates are reasonably similar in magnitude to the estimates of average marginal effects in my initial probit

Figure 3·4: Comparison of Propensity Scores for Pre-Migration Job Offer

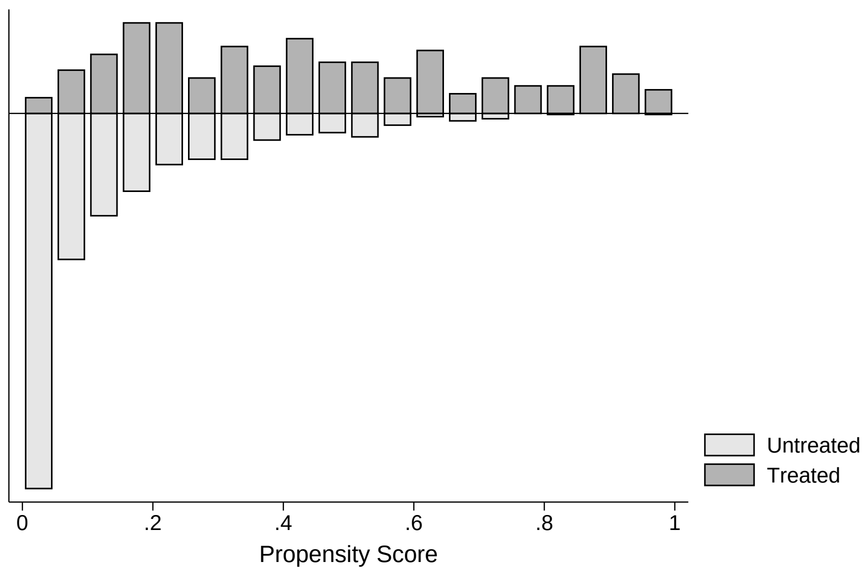

Figure 3·5: Comparison of Propensity Scores for Employer Sponsorship

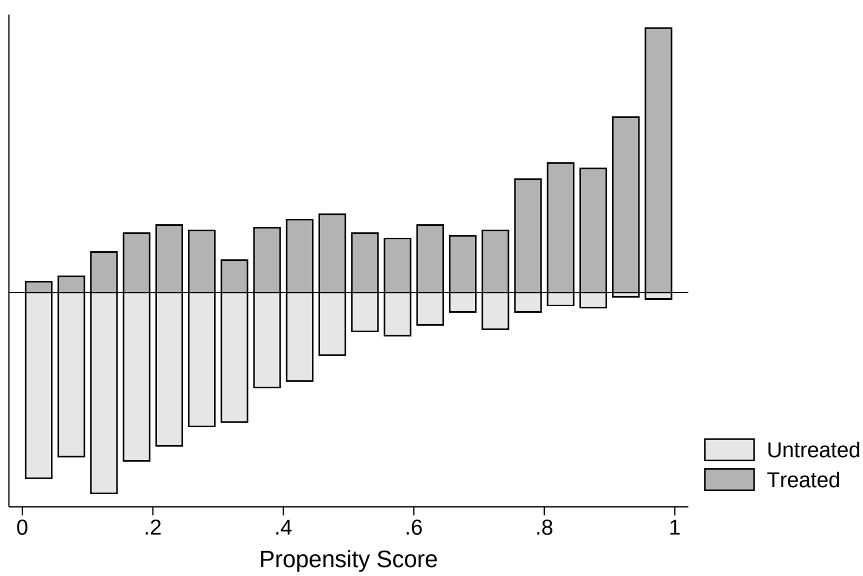

Table 3.9: Estimated Avg. Treatment Effects Using Propensity Score Analysis

	(1)	(2)	(3)	(4)
	Offer At Migration		Employer Sponsored	
Outcome Measure:	Binary	Continuous	Binary	Continuous
Avg. Treatment Effect	-0.0251	-0.209	-0.112***	-0.202**
	(0.0746)	(0.145)	(0.0363)	(0.0877)
Observations	1,150	1,102	1,343	1,261

Notes: Bootstrapped standard errors in parentheses. Estimates of the average treatment effect of having a job offer at migration or being employer sponsored are computed using the Epanechnikov kernel with default bandwidth. Propensity scores are computed via probit regression including all covariates presented in column (5) of the primary regression specification table, as well as the square of each continuous variable. *** p<0.01, ** p<0.05, * p<0.1

and OLS regressions, and in my regressions based with an employer sponsorship interaction. However, the standard errors of these estimates are larger, and so these results are in general no longer statistically significant at standard levels of confidence. Once I have access to more disaggregated geographic data, I expect that further analysis will provide me with additional insights.

Finally, I consider as a robustness check the suggestion of Angrist and Pischke (2008) that the initial probit scores used in propensity score matching may be used to limit the sample considered by more traditional regression methods. They suggest that when selection is a concern, that trimming from the sample individuals with very high or very low propensity scores may improve OLS estimates, and they show this phenomenon using an example from the literature on job training programs. Accordingly, Table 3.10 replicates my primary regression specifications, but excluding those individuals with a propensity score less than 0.1 or greater than 0.9, similar to what Angrist and Pischke did. The estimates that I produce using this method are generally quite similar to those produced using the baseline standard probit and OLS regressions. Even though the sample size is much smaller in these sample-trimmed regressions, they continue to show that individuals who received a job offer

Table 3.10: Regressions With Trimmed Sample Excluding Individuals with Low and High Propensity Scores

	(1)	(2)	(3)	(4)	(5)
	Panel A: Binary Enclave Indicator (Probit)				
Offer at Migration	-0.309**	-0.317**	-0.317**	-0.373***	-0.382***
	(0.151)	(0.146)	(0.144)	(0.129)	(0.120)
Observations	589	589	589	574	574
Avg. Marginal Effect	-0.0940	-0.0951	-0.0950	-0.100	-0.0999
	(0.0439)	(0.0414)	(0.0405)	(0.0327)	(0.0292)
	Panel B: Continuous Language Density (OLS)				
Offer at Migration	-0.236***	-0.235***	-0.233***	-0.154***	-0.159**
	(0.0886)	(0.0804)	(0.0819)	(0.0586)	(0.0629)
Observations	601	601	601	601	601
R-squared	0.391	0.401	0.403	0.460	0.472
Demographics	No	Yes	Yes	Yes	Yes
Family/Relatives	No	No	Yes	Yes	Yes
Pre-Migration Salary/Hours	No	No	No	Yes	Yes
Social Indicators	No	No	No	No	Yes

Notes: Robust standard errors in parentheses. All regressions include a full set of state/Census division and language-group fixed effects. Demographic controls include age, gender, years of schooling, and marital/parental status indicators. Family controls include indicators for whether an individual got their job from a relative or has a relative at their employer. Log Salary is measured in U.S. 2003 dollars. Social indicators include self-reported poor English language ability, poor health, and religiosity. Excludes individuals with a propensity score from probit estimation of less than 0.1 or greater than 0.9, using the same propensity scoring method as in Table 3.9.
*** p<0.01, ** p<0.05, * p<0.1

are roughly 10% less likely to locate in an ethnic enclave, and that this result is statistically significant.

3.8 Conclusion

In order to improve our understanding of ethnic enclaves, we need to understand why they form, and how and why some individuals elect to locate in them upon migration, while others do not. However, isolating the factors that influence the locational decision of new immigrants poses multiple empirical challenges. Living in an ethnic enclave may provide new immigrants not only with improved job search prospects, but also a wide range of social and economic benefits.[15] Furthermore, it is reasonable to expect that many of these benefits may accrue to immigrants regardless of whether they are employed or unemployed, and it is also plausible that some of these benefits may be correlated with the the likelihood of receiving a job offer in the first place. Yet, new immigrants are not randomly assigned to jobs, nor are they typically randomly assigned treatment along any other dimensions that we might expect to strongly influence their locational decisions. But, as I show in this chapter, individuals who arrive in the U.S. with a job offer already in hand have different characteristics from those who arrive without an offer. Absent cleaner sources of exogenous variation, the best we may hope for is to control carefully for these differences.

Still, by taking advantage of the detailed data and the representative sample of the New Immigrant Survey, this chapter shows that immigrants who arrive with a job offer are considerably less likely to locate in an ethnic enclave area, even after carefully controlling for a wide range of observable pre-migration characteristics. These results are also robust to numerous alternative specifications, and they are relatively free of concerns related to reverse causality. In keeping with the theoretical predictions of

[15] See, e.g., Bertrand, Luttmer, and Mullainathan (2000) and Patel and Vella (2012).

my model, this provides strong evidence that immigrants do locate in ethnic enclaves for reasons related to job search, and that this affects locational decisions both within and across metropolitan areas.

Perhaps unsurprisingly, regression analysis also provides suggestive evidence that several other factors influence the locational decision of migrants. Individuals who report that they speak English poorly or who have relatively low levels of education are more likely to locate in enclaves. So are individuals who had relatively low salaries or who worked relatively few hours pre-migration, given their occupation and industry. These findings support the notions that immigrants sort themselves into enclaves in part based on their personal preferences, their productive capacity, and/or their ability to assimilate.

Finally, this chapter also suggests additional avenues for research. In particular, I somewhat surprisingly find no significant evidence of heterogeneous effects of a job offer on locational choices. That is, the reduced likelihood of enclave residence associated with a job offer is the same whether an individual is highly educated or has little education. While I cannot rule out the role of sample size limitations, this suggests that the probability of job offers to immigrants may have little relationship to their own education. It also suggests that the patterns of negative selection into enclave residence are unlikely to be driven by the existence of job search networks alone. Both of these observations warrant further inquiry.

Contemporary immigration policy debates, particularly in the United States, remain heavily focused on issues related to the selection of immigrants. Proposed policy changes, such as restrictions on family reunification visas and moves toward "points-based" systems that reward pre-existing employment, might not only alter the set of immigrants who migrate. They may also have a particular influence on the demand for and the quality of the job search networks that immigrants exploit. And, as this chapter suggests, these policy changes may alter the locational decisions of migrants as well, with ramifications for the geographic distribution of immigrant groups that could persist far into the future. Isolating the mechanisms associated with locational choice remains a significant empirical challenge, but with additional research, there may be much more to learn about the role of policy and networks both in immigrants' economic outcomes and in broader patterns of regional economic activity.

Appendix A

Chapter 1 Appendices

A.1 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 1·4, 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.[1]

Therefore, my final measure of labor market formality is:

$$Formality_{s,o} = \frac{N_{formal,s,o}}{N_{s,o}} \tag{A.1}$$

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.

A.2 Additional Figures and Results

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

Figure A·1: Unadjusted Age Wage Profile in New Worker Sample, All Years

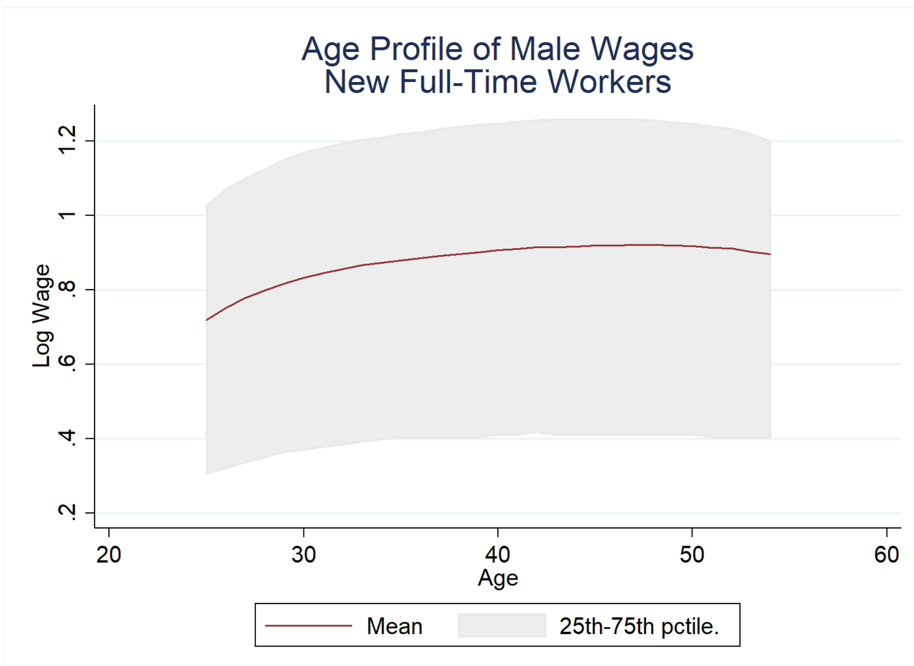

Notes: From the RAIS new worker sample, 1995-2014.

Figure A·2: Tenure Wage Profiles, Indexed by New Local Worker Average Occupational Wage

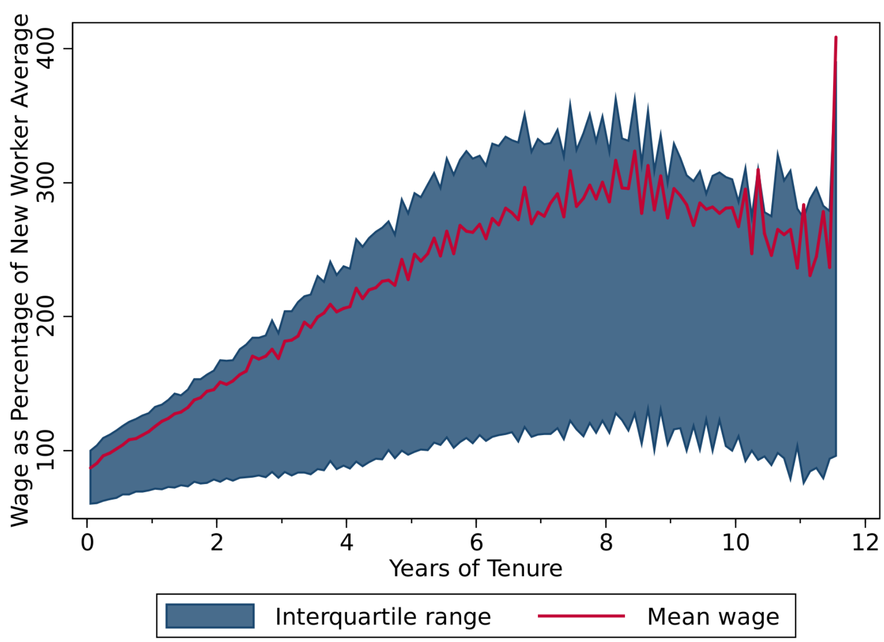

Notes: From the RAIS work history sample, 1995-2014. Mean wage and interquartile range are indexed, where 100 is the mean wage for new (< 1 year of tenure) workers of the same occupation in the same Brazilian micro-region in that year.

Figure A·3: Probability of Voluntary and Non-voluntary Separation by Years of Tenure

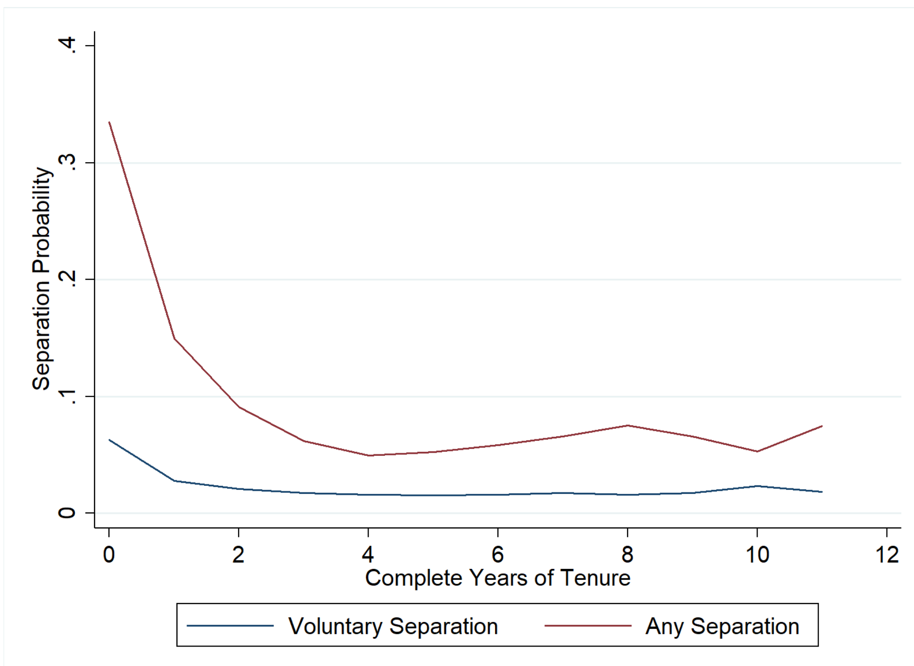

Notes: From the RAIS work history sample, 1995-2014. Voluntary and Non-Voluntary separation are as categorized in Table A.1

Table A.1: 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 A.2: 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 1.2. Standard errors in parentheses are clustered by establishment. *** p<0.01, ** p<0.05, * p<0.1

Table A.3: 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 1.2. Standard errors in parentheses are clustered by establishment.

Table A.4: 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) IV: Other Occs./Estabs.
Panel A: 1995-2001				
Occ. Growth	0.00326***	0.00844	0.0435***	0.124*
	(0.000612)	(0.00542)	(0.0147)	(0.0696)
Observations	2,098,000	1,944,169	434,638	625,932
Adjusted R-squared	0.894			
K-P F Stat		439.3	39.99	10.75
Panel B: 2002-2008				
Occ. Growth	0.00448***	0.0183***	0.0162	0.0670
	(0.000654)	(0.00316)	(0.00993)	(0.0424)
Observations	3,167,182	2,979,514	769,224	1,005,043
Adjusted R-squared	0.894			
K-P F Stat		210	810.6	13.50
Panel C: 2008-2014				
Occ. Growth	0.00596***	0.0185***	0.00397	0.0634***
	(0.000592)	(0.00164)	(0.00461)	(0.0136)
Observations	4,408,500	4,201,445	1,177,606	1,470,986
Adjusted R-squared	0.885			
K-P F Stat		264.9	660.6	49.51
Estab. FEs	Yes	Yes	Yes	Yes
Individual FEs	Yes	Yes	Yes	Yes
Occ-Micro-Year FEs	Yes	Yes	Yes	Yes

Notes: All regression specifications shown are as in the baseline regressions of Table 1.2, but with the panel period restricted to the seven-year interval shown. Standard errors reported in parentheses are two-way clustered by firm and micro-region. *** p<0.01, ** p<0.05, * p<0.1

Table A.5: New Worker Wage Regressions Excluding Fixed Effects

	(1) OLS	(2) IV: Other Occs.	(3) IV: Other Estabs.	(4) IV: Other Occs./Estabs.
Panel A: Excluding Individual FEs				
Occ. Growth	0.00215***	0.0237***	0.00265	0.0758**
	(0.000549)	(0.00255)	(0.00419)	(0.0323)
Observations	11,067,007	10,505,513	3,041,058	3,860,114
Adjusted R-squared	0.734			
K-P F Stat		223.2	662.6	67.04
Panel B: Excluding Establishment FEs				
Occ. Growth	0.00166	0.00901***	-0.00576	-0.0213
	(0.00107)	(0.00333)	(0.00507)	(0.0232)
Observations	9,895,100	9,329,615	2,470,999	3,209,726
Adjusted R-squared	0.829			
K-P F Stat		379.2	736.9	165
Panel C: Excluding Local Labor Market FEs				
Occ. Growth	0.00851***	0.0143***	0.0204***	0.0643**
	(0.000516)	(0.00306)	(0.00375)	(0.0289)
Observations	9,875,692	9,324,578	2,475,797	3,218,512
Adjusted R-squared	0.858			
K-P F Stat		373.2	513.2	70.03

Notes: All regression specifications shown are as in the baseline regressions of Table 1.2, but with one category of fixed effects excluded from the regression specifications in each panel. Standard errors reported in parentheses are two-way clustered by firm and micro-region. *** p<0.01, ** p<0.05, * p<0.1

Appendix B

Chapter 2 Appendices

B.1 Proof of Proposition 2

Assume that there exists a profit maximizing equilibrium in which there exists an interval of tasks $S = [\underline{s}, \bar{s}]$ over which two employees, i and i', are each involved in production.

By profit maximization, it also follows that the marginal product of each worker's time must be the same at each point in the interval. That is:

$$\frac{\partial Y(s)}{T_i(s)} = \frac{\partial Y(s)}{T_{i'}(s)}$$

for all s in the overlapping interval. Since the ordering of tasks in production is arbitrary, the assumption that S is compact is of no consequence here.

Additionally, assume without loss of generality that employee i is engaged in task-specific human capital investment along this interval. Then, from equal marginal product of time, it follows that there exists an $\epsilon > 0$ such that by partitioning S at the point $\underline{s} + \epsilon$, allocating all production performed by i' to i in $[\underline{s}, \underline{s} + \epsilon]$, and allocating all production performed by i to i' in $[\underline{s} + \epsilon, \bar{s}]$, the overall level of production is maintained, and both employees' time constraints are satisfied.

However, from (2.7), it follows that the marginal product of time of worker i will be increased by this reallocation, implying that i can produce more output. And, since the marginal product of time of worker i' cannot decrease as a result of this reollocation, the overlapping initial allocation cannot be an optimum.

B.2 Proof of Proposition 3

Suppose that employee i produces intermediate goods s and s'. Then it follows from Proposition 2 and from profit maximization that:

$$\frac{\partial Y_i(s)}{\partial T_i(s)} = \frac{\partial Y_i(s')}{\partial T_i(s')}$$

which from (2.7) is shown to be true only if $T_i(s) = T_i(s')$. The constraint that $\eta < \frac{1}{1+\theta}$ is a sufficient condition to ensure that this equal time allocation maximizes profit, otherwise tasks may be sufficiently substitutable that there exists no equilibrium in which workers perform a finite measure of tasks.

To see the second result, observe that since all workers who specialize are identical (and since firms cannot wage discriminate based on information about worker-specific preferences that they cannot observe), any within-firm equilibrium must involve all workers receiving equal compensation, and therefore by cost minimization they must have identical marginal product. Since, by assumption, firms must allocate workers over the entire unit interval of production tasks, this can only be true if $|s_i| = \frac{1}{n_f^*}$.

B.3 Additional Tables

Table B.1: Baseline Regressions: Using O*NET Categorization

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
Mean Cognitive Skills	0.429***	0.203***	0.0774**	0.0768***	0.0662***	0.0818***
	(0.0337)	(0.0206)	(0.0301)	(0.0204)	(0.0186)	(0.0152)
Mean Production Skills	0.172*	0.198***	0.0831**	0.0838***	0.0943***	0.0563***
	(0.0906)	(0.0491)	(0.0392)	(0.0227)	(0.0204)	(0.0172)
Gini Cognitive Skills	-0.0263	-0.0273	0.0119	0.00965	0.0155	0.0241**
	(0.0455)	(0.0286)	(0.0234)	(0.0166)	(0.0166)	(0.0122)
Gini Production Skills	0.122*	0.0821**	0.0563*	0.0577***	0.0617***	0.0415***
	(0.0715)	(0.0365)	(0.0319)	(0.0215)	(0.0209)	(0.0160)
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. Mean skill in each group is measured as the average of skill within 6 broad O*NET categories, weighted equally. Gini of Cognitive skill is measured as the Gini of skill across the same broad O*NET categories. Gini of Production Skills is in Table 2.7 because all Production skills are within the "Technical" category of O*NET skills. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

Table B.2: Baseline Regressions: Alternative Skill Measures

	(1)	(2)	(3)	(4)	(5)	(6)
	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings	Log Earnings
			Theil Index			
Mean Cognitive Skills	0.411***	0.188***	0.0636**	0.0580***	0.0586***	0.0670***
	(0.0423)	(0.0238)	(0.0323)	(0.0208)	(0.0188)	(0.0151)
Mean Production Skills	0.173**	0.192***	0.0632**	0.0708***	0.0776***	0.0454***
	(0.0703)	(0.0388)	(0.0320)	(0.0161)	(0.0159)	(0.0134)
Theil Cognitive Skills	-0.0510	-0.0365	0.0157	0.0406**	0.0469***	0.0400***
	(0.0361)	(0.0227)	(0.0247)	(0.0171)	(0.0169)	(0.0139)
Theil Production Skills	0.0966*	0.0616**	0.0291	0.0348***	0.0361***	0.0242***
	(0.0526)	(0.0265)	(0.0192)	(0.0121)	(0.0113)	(0.00833)
Theil Between Groups	-0.00396	-0.0125	-0.0275	-0.0376**	-0.0261	-0.0289**
	(0.0350)	(0.0242)	(0.0237)	(0.0169)	(0.0163)	(0.0139)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.317	0.524	0.575	0.659	0.684	0.820
			Standard Deviation of Skills			
Mean Cognitive Skills	0.459***	0.222***	0.0604*	0.0485**	0.0311	0.0478**
	(0.0355)	(0.0214)	(0.0345)	(0.0233)	(0.0223)	(0.0189)
Mean Production Skills	0.105**	0.148***	0.0656***	0.0516***	0.0546***	0.0272**
	(0.0507)	(0.0260)	(0.0234)	(0.0157)	(0.0136)	(0.0122)
SD Cognitive Skills	-0.0216	-0.00614	0.0178	0.0239*	0.0285**	0.0241**
	(0.0389)	(0.0234)	(0.0198)	(0.0140)	(0.0133)	(0.0112)
SD Production Skills	0.0177	0.0216	0.0463***	0.0278***	0.0248***	0.0142*
	(0.0300)	(0.0167)	(0.0155)	(0.0104)	(0.00954)	(0.00777)
Observations	59,803,061	59,696,420	59,696,420	59,696,420	59,200,644	59,696,420
Adjusted R-squared	0.312	0.522	0.576	0.659	0.684	0.820
Individual Controls	No	Yes	Yes	Yes	Yes	Yes
Occupational Controls	None	None	Two-digit	Two-digit	Two-digit	Two-digit
Employer Controls	None	None	None	Subsector	Sub.+Controls	Estab. FEs

Robust standard errors in parentheses are clustered by CBO occupation. From a 100% sample of eligible worker-year observations from RAIS. All specifications include year controls. Individual controls include linear and quadratic terms in age, potential experience, and job tenure, four educational indicators, and nationality indicators. Unreported establishment controls include log establishment size and proportions of workers in each of five occupation categories eight age groups and four educational categories excluding the individual. *** p<0.01, ** p<0.05, * p<0.1

Appendix C

Chapter 3 Appendices

Table C.1: Description of Key Variables in the New Immigrant Survey

Variable	NIS Questions	Response Coding
Had At Migration	"Had you been offered this job before coming to the United States to live?"	1 if yes for any current job
Got Job Relative	"Did you get this job with the help of a relative of yours or your husband/wife?"	1 if yes for any current job
Relative At Emp.	"Do any of your, or your husband/wife's relatives work for this business"	1 if yes for any current job

Table C.2: Description of Enclave Measures Using U.S. Census and NIS Data

Key Definitions	
List of Languages (NIS)	"What languages do you currently speak at home?"
Binary threshold	$2\times$ National Avg. % Spoken
Continuous measure	$\ln\left(\frac{\text{Local \% Spoken}}{\text{Nat'l Avg. \% Spoken}} + \left(\left(\frac{\text{Local \% Spoken}}{\text{Nat'l Avg. \% Spoken}}\right)^2 + 1\right)^{1/2}\right)$

Methods of Individual Group Assigment	
Measure used	Method of group assignment
Results shown	
First Lang.	The first non-English language reported to be spoken by each immigrant (results reported).
Robustness checks (available on request)	
Lang. Density	The reported language with the highest local percentage density relative to its national average
Abs. Lang. Density	The reported language with the highest local percentage density

Notes: All enclaves are defined at the ZIP code level using data from the 2000 Census, SF-3 Sample, available at http://factfinder2.census.gov/.

Table C.3: New Immigrant Survey Language Groups

Group ID	Language	Number of Speakers as First Language
2	Spanish	2,312
3	Portuguese	51
4	Russian	284
5	French	83
6	Arabic	226
7	Chinese	501
8	Tagalog	257
9	Vietnamese	181
10	Korean	126
11	Bengali	57
12	Hindi	107
13	Polish	173
14	Amharic	162
15	Creole	111
16	Other European	410
17	Other Non-European	586
18	Other Spoken in Philippines	75
19	Other Spoken in India	541
20	Other	2

Notes: Counts are as provided by NIS in response to the question "What languages do you currently speak at home?". English is excluded from analysis.

Table C.4: Key Variables By Education Group

	Offered Job Prior To Move	Helped By Relative To Get Job	Relative Works For Company
Less Than HS	0.038	0.267	0.186
	(0.006)	(0.014)	(0.012)
	[923]	[940]	[1,060]
High School	0.096	0.272	0.140
	(0.018)	(0.026)	(0.019)
	[281]	[287]	[321]
Some College	0.072	0.157	0.093
	(0.012)	(0.017)	(0.013)
	[447]	[458]	[486]
Bachelors Degree	0.249	0.105	0.077
	(0.012)	(0.009)	(0.007)
	[1,215]	[1,263]	[1,359]
Graduate School	0.222	0.051	0.053
	(0.014)	(0.007)	(0.007)
	[851]	[895]	[944]
Unknown/No Answer	0.048	0.242	0.157
	(0.007)	(0.015)	(0.012)
	[814]	[839]	[903]

Notes: Standard errors in parentheses. Counts in brackets.

Table C.5: Frequency of Pre-Migration Offers by Occupational Group

	Received Pre-Migration Offer			
Occupation	Yes	No	Missing	Total
OFFICE AND ADMIN. SUPPORT	4.0%	12.3%	13.5%	12.1%
SALES AND RELATED	7.1%	11.4%	13.1%	11.8%
EXEC., ADMIN. AND MANAGERIAL	11.5%	9.3%	10.8%	10.2%
TEACHERS	4.7%	7.0%	8.9%	7.6%
SETTER, OPERATORS, AND TENDERS	3.3%	7.5%	7.1%	7.0%
MATH. AND COMPUTER SCIENTISTS	16.9%	6.8%	2.0%	5.6%
HEALTH DIAG. AND TREAT. PRACTIT.	18.8%	3.5%	4.7%	5.4%
TRANSP. AND MAT. MOVING	1.3%	4.5%	3.9%	4.0%
MANAGEMENT RELATED	2.9%	3.3%	4.6%	3.8%
CONSTRUCTION TR. AND EXTRACTION	0.9%	4.8%	3.2%	3.7%
ENGINEERS, ARCHITECTS, SURVEYORS	6.4%	3.1%	2.7%	3.2%
FARMING, FISHING, FORESTRY	0.4%	3.6%	3.4%	3.2%
FOOD PREPARATIONS AND SERVING	4.9%	3.3%	2.3%	3.0%
INSTALLATION, MAINTENANCE, AND REPAIR	1.1%	3.5%	2.7%	3.0%
Total (including omitted)	100.0%	100.0%	100.0%	100.0%
Percentage Reporting Pre-Migration Occupation	72.2%	60.5%	55.9%	59.2%

Table C.6: Frequency of Pre-Migration Offers by Industry Group

	Received Pre-Migration Offer			
Occupation	Yes	No	Missing	Total
EDUC., HEALTH, SOCIAL SVCS.	25.9%	14.0%	17.9%	16.8%
MANUFACTURING	12.6%	13.2%	13.0%	13.1%
PROFESSIONAL AND RELATED SVCS.	22.4%	11.6%	9.1%	11.4%
RETAIL TRADE	4.9%	10.4%	10.5%	10.0%
ENTERT., ACCOM., AND FOOD SVCS.	6.9%	7.7%	5.8%	6.8%
AGRICULTURE, FORESTRY, FISHERIES	0.7%	6.8%	7.3%	6.5%
OTHER SERVICES	5.1%	5.6%	6.1%	5.8%
PUBLIC ADMINISTRATION	1.6%	4.9%	6.5%	5.3%
CONSTRUCTION	2.0%	6.4%	4.7%	5.3%
FINANCE, INSURANCE, REAL ESTATE	4.7%	4.5%	4.5%	4.5%
TRANSPORT. AND WAREHOUSING	2.0%	4.2%	4.3%	4.1%
WHOLESALE TRADE	3.5%	3.2%	4.0%	3.6%
INFO. AND COMMUNICATION	4.2%	2.7%	2.3%	2.7%
UNCODABLE	2.4%	2.7%	2.6%	2.6%
UTILITIES	0.2%	0.8%	0.5%	0.6%
MINING	0.4%	0.5%	0.5%	0.5%
ARMED FORCES	0.4%	0.5%	0.3%	0.4%
NOT IN LABOR FORCE	0.0%	0.0%	0.1%	0.1%
Total	100.0%	100.0%	100.0%	100.0%
Percentage Reporting Pre-Migration Industry	72.2%	60.5%	55.9%	59.2%

Table C.7: Frequency of Pre-Migration Offers by Social Indicators

	Received Pre-Migration Offer			
	Yes	No	Missing	Total
English is Poor	0.198	0.415	0.585	0.480
	(0.016)	(0.008)	(0.008)	(0.005)
Poor Health	0.010	0.025	0.051	0.036
	(0.004)	(0.002)	(0.003)	(0.002)
Regular Religious Attendance	0.598	0.620	0.603	0.611
	(0.020)	(0.008)	(0.008)	(0.005)

Bibliography

2011 Yearbook of Immigration Statistics (Sept. 2012). Tech. rep. Office of Immigration Statistics, U.S. Department of Homeland Security. URL: http://www.dhs.gov/sites/default/files/publications/immigration-statistics/yearbook/2011/ois_yb_2011.pdf.

Abowd, John M., Robert H. Creecy, and Francis Kramarz (Apr. 2002). *Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data*. Tech. rep. TP-2002-06. U.S. Census Bureau.

Abowd, John M., Francis Kramarz, and David N. Margolis (Mar. 1999). "High Wage Workers and High Wage Firms". In: *Econometrica* 67.2, pp. 251–333. ISSN: 1468-0262. DOI: 10.1111/1468-0262.00020. (Visited on 10/01/2015).

Acemoglu, Daron and Jorn-Steffen Pischke (1998). "Why Do Firms Train? Theory and Evidence". In: *The Quarterly Journal of Economics* 113.1, pp. 79–119. ISSN: 0033-5533.

Angrist, Joshua D. and Jörn-Steffen Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.

Autor, David H. (Jan. 2013). *The "Task Approach" to Labor Markets: An Overview*. Working Paper 18711. National Bureau of Economic Research. (Visited on 11/11/2014).

Autor, David H. and David Dorn (Aug. 2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market". In: *American Economic Review* 103.5, pp. 1553–1597. ISSN: 0002-8282. DOI: 10.1257/aer.103.5.1553. (Visited on 02/26/2018).

Autor, David H., David Dorn, and Gordon H. Hanson (Oct. 2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States". In: *American Economic Review* 103.6, pp. 2121–2168. ISSN: 0002-8282. DOI: 10.1257/aer.103.6.2121. (Visited on 09/21/2017).

Autor, David H., Frank Levy, and Richard J. Murnane (Nov. 2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333. ISSN: 0033-5533. (Visited on 09/14/2015).

Barth, Erling and Harald Dale-Olsen (Oct. 2009). "Monopsonistic Discrimination, Worker Turnover, and the Gender Wage Gap". In: *Labour Economics* 16.5, pp. 589–597. ISSN: 0927-5371. DOI: 10.1016/j.labeco.2009.02.004.

Barth, Erling et al. (Feb. 2016). “It’s Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States”. In: *Journal of Labor Economics* 34.S2, S67–S97. ISSN: 0734-306X. DOI: 10.1086/684045.

Baumgardner, James R. (1988). “The Division of Labor, Local Markets, and Worker Organization”. In: *Journal of Political Economy* 96.3, pp. 509–527. ISSN: 0022-3808. URL: http://www.jstor.org/stable/1830357 (visited on 09/12/2016).

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa (Dec. 2008). “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes”. In: *Journal of Political Economy* 116.6, pp. 1150–1196. ISSN: 0022-3808, 1537-534X. DOI: 10.1086/595975. URL: http://www.jstor.org/stable/10.1086/595975 (visited on 03/07/2012).

Beaman, L.A. (2011). “Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the US”. In: *The Review of Economic Studies*.

Becker, Gary S. (1962). “Investment in Human Capital: A Theoretical Analysis”. In: *Journal of Political Economy* 70.5, pp. 9–49. ISSN: 0022-3808. (Visited on 09/26/2016).

Becker, Gary S. and Kevin M. Murphy (Nov. 1992). “The Division of Labor, Coordination Costs, and Knowledge”. In: *The Quarterly Journal of Economics* 107.4, pp. 1137–1160. ISSN: 0033-5533, 1531-4650. DOI: 10.2307/2118383. URL: http://qje.oxfordjournals.org/content/107/4/1137 (visited on 09/08/2016).

Bellon, Matthieu (2016). *The Characteristics of Worker Flows by Firm Growth: Empirical Evidence from a Matched Firm-Worker Dataset from France*. Working Paper. (Visited on 09/22/2017).

Bertrand, Marianne, Erzo F. P. Luttmer, and Sendhil Mullainathan (2000). “Network Effects and Welfare Cultures”. In: *The Quarterly Journal of Economics* 115.3, pp. 1019–1055. ISSN: 00335533. URL: http://www.jstor.org/stable/2586902.

Bhaskar, V. and Ted To (Apr. 1999). “Minimum Wages for Ronald McDonald Monopsonies: A Theory of Monopsonistic Competition”. In: *The Economic Journal* 109.455, pp. 190–203. ISSN: 1468-0297. DOI: 10.1111/1468-0297.00427.

Bleakley, Hoyt and Jeffrey Lin (2012). “Thick-Market Effects and Churning in the Labor Market: Evidence from U.S. Cities”. In: *Journal of urban economics* 72.2-3, pp. 87–103. ISSN: 0094-1190. DOI: 10.1016/j.jue.2012.04.003.

Booth, Alison L. and Pamela Katic (Sept. 2011). “Estimating the Wage Elasticity of Labour Supply to a Firm: What Evidence Is There for Monopsony?*”. In: *Economic Record* 87.278, pp. 359–369. ISSN: 1475-4932. DOI: 10.1111/j.1475-4932.2011.00728.x.

Borjas, George J. (1987). “Self-Selection and the Earnings of Immigrants”. In: *The American Economic Review* 77.4, pp. 531–553. ISSN: 0002-8282. URL: http://www.jstor.org/stable/1814529 (visited on 09/12/2011).

Borjas, George J. and Rachel M. Friedberg (2009). “Recent Trends in the Earnings of New Immigrants to the United States”. In: *National Bureau of Economic Research Working Paper Series* No. 15406. URL: http://www.nber.org/papers/w15406.

Brown, Charles C. and James L. Medoff (1989). *The employer size-wage effect*. National Bureau of Economic Research Cambridge, Mass., USA. URL: http://www.nber.org/papers/w2870 (visited on 09/28/2016).

Brummund, Peter (2011). *Variation in Monopsonistic Behavior across Establishments: Evidence from the Indonesian Labor Market*. Working Paper. (Visited on 09/22/2017).

Burdett, Kenneth and Dale T. Mortensen (May 1998). “Wage Differentials, Employer Size, and Unemployment”. In: *International Economic Review* 39.2, pp. 257–273. ISSN: 0020-6598. DOI: 10.2307/2527292. (Visited on 09/13/2011).

Bureau, U.S. Census. *Industry and Occupation Code Lists & Crosswalks*. https://www.census.gov/topics/employment/industry-occupation/guidance/code-lists.html.

Card, David (2001). “Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration”. In: *Journal of Labor Economics* 19.1, pp. 22–64. (Visited on 09/21/2017).

Card, David and Stefano DellaVigna (2013). “Nine Facts about Top Journals in Economics”. In: *Journal of Economic Literature* 51.1, pp. 144–161. ISSN: 0022-0515. URL: http://www.jstor.org/stable/23644706 (visited on 09/12/2016).

Card, David, Jörg Heining, and Patrick Kline (2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality*”. In: *The Quarterly Journal of Economics* 128.3, pp. 967–1015. DOI: 10.1093/qje/qjt006. eprint: /oup/backfile/content_public/journal/qje/128/3/10.1093_qje_qjt006/4/qjt006.pdf. URL: http://dx.doi.org/10.1093/qje/qjt006.

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

Card, David et al. (Dec. 2017). “Firms and Labor Market Inequality: Evidence and Some Theory”. In: *Journal of Labor Economics* 36.S1, S13–S70. ISSN: 0734-306X. DOI: 10.1086/694153.

Correia, Sergio (2016). “A Feasible Estimator for Linear Models with Multi-Way Fixed Effects”. In: *Duke University Preliminary Version. URL: www. scorreia. com/research/hdfe. pdf*. (Visited on 09/21/2017).

Costinot, Arnaud (Apr. 2009). “On the Origins of Comparative Advantage”. In: *Journal of International Economics* 77.2, pp. 255–264. ISSN: 0022-1996. DOI: 10.1016/j.jinteco.2009.01.007. (Visited on 04/14/2016).

Cutler, D.M., E.L. Glaeser, and J.L. Vigdor (2008). “Is the melting pot still hot? Explaining the resurgence of immigrant segregation”. In: *The Review of Economics and Statistics* 90.3, pp. 478–497.

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

Dehejia, Rajeev H. and Sadek Wahba (Feb. 2002). "Propensity Score-Matching Methods for Nonexperimental Causal Studies". In: *Review of Economics and Statistics* 84.1, pp. 151–161. ISSN: 0034-6535. DOI: 10.1162/003465302317331982. URL: http://dx.doi.org/10.1162/003465302317331982 (visited on 10/01/2013).

Deming, David J. (Aug. 2015). *The Growing Importance of Social Skills in the Labor Market*. Working Paper 21473. National Bureau of Economic Research. URL: http://www.nber.org/papers/w21473 (visited on 09/14/2015).

Depew, Briggs and Todd Sorensen (2013). "The Elasticity of Labor Supply to the Firm over the Business Cycle". In: *Labour Economics* 24.C, pp. 196–204. ISSN: 0927-5371.

Dessein, Wouter and Tano Santos (2006). "Adaptive Organizations". In: *Journal of Political Economy* 114.5, pp. 956–995. ISSN: 0022-3808. DOI: 10.1086/508031. URL: http://www.jstor.org.ezproxy.bu.edu/stable/10.1086/508031 (visited on 09/08/2016).

Dix-Carneiro, Rafael and Brian K. Kovak (2017). *Margins of Labor Market Adjustment to Trade*. Working Paper 23595. National Bureau of Economic Research. (Visited on 09/21/2017).

Edin, Per-Anders, Peter Fredriksson, and Olof Åslund (Feb. 2003). "Ethnic Enclaves and the Economic Success of Immigrants: Evidence from a Natural Experiment". In: *The Quarterly Journal of Economics* 118.1, pp. 329–357. ISSN: 0033-5533. URL: http://www.jstor.org/stable/25053906 (visited on 01/23/2012).

Falch, Torberg (2010). "The Elasticity of Labor Supply at the Establishment Level". In: *Journal of Labor Economics* 28.2, pp. 237–266. ISSN: 0734-306X. DOI: 10.1086/649905.

Fisher, Jonathan D. and Christina Houseworth (Sept. 2012). *Occupation Inflation in the Current Population Survey*. SSRN Scholarly Paper ID 2205198. Rochester, NY: Social Science Research Network. URL: http://papers.ssrn.com/abstract=2205198 (visited on 09/29/2016).

Fishman, Arthur and Avi Simhon (June 2002). "The Division of Labor, Inequality and Growth". In: *Journal of Economic Growth* 7.2, pp. 117–136. ISSN: 1381-4338, 1573-7020. DOI: 10.1023/A:1015672012193. (Visited on 11/19/2014).

Furtado, Delia and Nikolaos Theodoropoulos (2012). *Immigrant networks and the take-up of disability programs: Evidence from US census data*. Tech. rep. 6649. Discussion Paper series, Forschungsinstitut zur Zukunft der Arbeit. URL: http://www.econstor.eu/handle/10419/62380 (visited on 10/15/2014).

Gan, Li and Qi Li (May 2016). "Efficiency of Thin and Thick Markets". In: *Journal of Econometrics* 192.1, pp. 40–54. ISSN: 03044076. DOI: 10.1016/j.jeconom.2015.10.012. (Visited on 09/24/2017).

Garicano, Luis and Esteban Rossi-Hansberg (Nov. 2006). "Organization and Inequality in a Knowledge Economy". In: *The Quarterly Journal of Economics* 121.4, pp. 1383–1435. ISSN: 0033-5533, 1531-4650. DOI: 10.1093/qje/121.4.1383. URL: http://qje.oxfordjournals.org.ezproxy.bu.edu/content/121/4/1383 (visited on 09/08/2016).

Gathmann, Christina and Uta Schönberg (Jan. 2010). "How General Is Human Capital? A Task-Based Approach". In: *Journal of Labor Economics* 28.1, pp. 1–49. DOI: 10.1086/648416. (Visited on 04/10/2014).

Goel, D. and K. Lang (2016). *Social Ties and the Job Search of Recent Immigrants*. Working Paper 9942. IZA. URL: http://ftp.iza.org/dp9942.pdf.

Guvenen, Faith et al. (Sept. 2015). "Multidimensional Skill Mismatch". In: *Working Paper Series* No. 2015-022. URL: http://research.stlouisfed.org/wp/more/2015-022/ (visited on 09/09/2016).

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (Apr. 2011). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. 2nd ed. 2009. Corr. 7th printing 2013 edition. New York, NY: Springer. ISBN: 978-0-387-84857-0.

Heckman, James J. and Bo E. Honoré (1990). "The Empirical Content of the Roy Model". In: *Econometrica* 58.5, pp. 1121–1149. ISSN: 0012-9682. DOI: 10.2307/2938303. URL: http://www.jstor.org/stable/2938303 (visited on 09/12/2011).

Helpman, Elhanan et al. (2017). "Trade and Inequality: From Theory to Estimation". In: *The Review of Economic Studies* 84.1, pp. 357–405.

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

Hirsch, Boris, Thorsten Schank, and Claus Schnabel (2010). "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany". In: *Journal of Labor Economics* 28.2, pp. 291–330.

Hotchkiss, Julie L. and Myriam Quispe-Agnoli (Dec. 2009). *Employer Monopsony Power in the Labor Market for Undocumented Workers*. SSRN Scholarly Paper ID 1459045. Rochester, NY: Social Science Research Network. (Visited on 09/21/2017).

Jasso, Guillermina et al. (Mar. 2006). *The New Immigrant Survey 2003 Round 1 (NIS-2003-1) Public Release Data*. Tech. rep. URL: http://nis.princeton.edu. (visited on 04/10/2012).

Kambourov, Gueorgui and Iourii Manovskii (Feb. 2009). "Occupational Specificity of Human Capital*". In: *International Economic Review* 50.1, pp. 63–115. ISSN: 1468-2354. DOI: 10.1111/j.1468-2354.2008.00524.x. URL: http://onlinelibrary.wiley.com.ezproxy.bu.edu/doi/10.1111/j.1468-2354.2008.00524.x/abstract (visited on 11/11/2014).

Kremer, Michael and Eric Maskin (Aug. 1996). *Wage Inequality and Segregation by Skill*. SSRN Scholarly Paper ID 3631. Rochester, NY: Social Science Research Network. (Visited on 09/21/2017).

Kuhn, Peter and Hani Mansour (Dec. 2014). "Is Internet Job Search Still Ineffective?" In: *The Economic Journal* 124.581, pp. 1213–1233. ISSN: 1468-0297. DOI: 10.1111/ecoj.12119.

Lang, Kevin and Jee-Yeon K. Lehmann (Dec. 2012). "Racial Discrimination in the Labor Market: Theory and Empirics". In: *Journal of Economic Literature* 50.4, pp. 959–1006. ISSN: 0022-0515. DOI: 10.1257/jel.50.4.959. (Visited on 09/21/2017).

Lavetti, Kurt and Ian Schmutte (Aug. 2016). "Estimating Compensating Wage Differentials with Endogenous Job Mobility". In: *Labor Dynamics Institute*.

Lazear, Edward P. (Dec. 1999). "Culture and Language". In: *Journal of Political Economy* 107.S6, S95–S126. ISSN: 0022-3808. URL: http://www.jstor.org/stable/10.1086/250105 (visited on 05/03/2012).

Legros, Patrick, Andrew F. Newman, and Eugenio Proto (Oct. 2013). "Smithian Growth through Creative Organization". In: *Review of Economics and Statistics* 96.5, pp. 796–811. ISSN: 0034-6535. DOI: 10.1162/REST_a_00421. URL: http://dx.doi.org/10.1162/REST_a_00421 (visited on 09/08/2016).

MacKinnon, James G. and Lonnie Magee (May 1990). "Transforming the Dependent Variable in Regression Models". In: *International Economic Review* 31.2, pp. 315–339. ISSN: 0020-6598. DOI: 10.2307/2526842. URL: http://www.jstor.org/stable/2526842 (visited on 10/10/2014).

Macaluso, Claudia (2017). *Skill Remoteness and Post-Layoff Labor Market Outcomes*. Working Paper. (Visited on 09/22/2017).

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

Marinescu, Ioana and Roland Rathelot (Sept. 2016). *Mismatch Unemployment and the Geography of Job Search*. Working Paper 22672. National Bureau of Economic Research. DOI: 10.3386/w22672.

Matsudaira, Jordan D. (2014). "Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations". In: *The Review of Economics and Statistics* 96.1, pp. 92–102. (Visited on 09/21/2017).

McFadden, Daniel (1973). "Conditional Logit Analysis of Qualitative Choice". In: *Frontiers in Econometrics*. New York: Academic Press, pp. 105–142.

Menezes-Filho, Naércio Aquino, Marc-Andreas Muendler, and Garey Ramey (2008). "The Structure of Worker Compensation in Brazil, with a Comparison to France and the United States". In: *The Review of Economics and Statistics* 90.2, pp. 324–346.

Montgomery, James D. (Dec. 1991). "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis". In: *The American Economic Review* 81.5, pp. 1408–1418. ISSN: 0002-8282. DOI: 10.2307/2006929. URL: http://www.jstor.org/stable/2006929 (visited on 09/27/2013).

Muendler, Marc-Andreas et al. (2004). "Job Concordances for Brazil: Mapping the Classificaçao Brasileira de Ocupaçoes (CBO) to the International Standard Classification of Occupations (ISCO-88)". In: *University of California, San Diego, unpublished manuscript*. (Visited on 09/25/2016).

Munshi, K. (2003). "Networks in the modern economy: Mexican migrants in the US labor market". In: *The Quarterly Journal of Economics* 118.2, p. 549.

*O*NET Resource Center - O*NET 21.0 Production Database*. URL: https://www.onetcenter.org/database.html (visited on 09/27/2016).

Oi, Walter Y. and Todd L. Idson (1999). "Firm Size and Wages". In: *Handbook of Labor Economics*. Vol. 3, Part B. Elsevier, pp. 2165–2214. (Visited on 09/15/2017).

PIS - Programa Integração Social | Caixa. http://www.caixa.gov.br/beneficios-trabalhador/pis/Paginas/default.aspx. (Visited on 09/24/2017).

Patel, Krishna and Francis Vella (Oct. 2012). "Immigrant Networks and Their Implications for Occupational Choice and Wages". In: *Review of Economics and Statistics*. ISSN: 0034-6535. DOI: 10.1162/REST_a_00327. URL: http://dx.doi.org/10.1162/REST_a_00327 (visited on 09/27/2013).

Peri, Giovanni and Chad Sparber (July 2009). "Task Specialization, Immigration, and Wages". In: *American Economic Journal. Applied Economics* 1.3, pp. 135–169. ISSN: 19457782. DOI: 10.1257/app.1.3.135.

Poletaev, Maxim and Chris Robinson (July 2008). "Human Capital Specificity: Evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984-2000". In: *Journal of Labor Economics* 26.3, pp. 387–420. ISSN: 0734-306X. DOI: 10.1086/588180. URL: http://www.journals.uchicago.edu/doi/10.1086/588180 (visited on 05/03/2016).

Radner, Roy and Timothy Van Zandt (1995). "Information Processing in Firms and Returns to Scale". In: *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability*. Ed. by John O. Ledyard. Springer US, pp. 243–280. ISBN: 978-1-4613-5953-1 978-1-4615-2261-4. URL: http://link.springer.com.ezproxy.bu.edu/chapter/10.1007/978-1-4615-2261-4_10 (visited on 09/08/2016).

Ransom, Michael R. and Ronald L. Oaxaca (2010). "New Market Power Models and Sex Differences in Pay". In: *Journal of Labor Economics* 28.2, pp. 267–289. (Visited on 09/21/2017).

Ransom, Michael R and David P. Sims (2010). "Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri". In: *Journal of Labor Economics* 28.2, pp. 331–355. ISSN: 0734-306X. DOI: 10.1086/649904.

Rivera, Jason M. (Dec. 2013). *Pay Structure and Turnover: Evidence from Brazilian Employer-Employee Matched Data*. SSRN Scholarly Paper ID 2356720. Rochester, NY: Social Science Research Network. (Visited on 09/19/2017).

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

Roy, A. D. (June 1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pp. 135–146. ISSN: 0030-7653. DOI: 10.2307/2662082. URL: http://www.jstor.org/stable/2662082 (visited on 09/27/2013).

Saraiva, Adriana Goncalves. *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. (Visited on 09/24/2017).

Satchi, Mathan and Jonathan Temple (Jan. 2009). "Labor Markets and Productivity in Developing Countries". In: *Review of Economic Dynamics* 12.1, pp. 183–204. ISSN: 1094-2025. DOI: 10.1016/j.red.2008.09.001.

Schmieder, Johannes (Oct. 2013). "What Causes Wage Dispersion: Evidence from New Firms". https://sites.google.com/site/johannesschmieder/Schmieder_WagesNewFirms_WorkingPaper.pdf?attredirects=0&d=1. Working Paper.

Skills Search. URL: https://www.onetonline.org/skills/ (visited on 09/27/2016).

Smith, Adam and Joseph Shield Nicholson (1887). *An Inquiry Into the Nature and Causes of the Wealth of Nations...* T. Nelson and Sons. (Visited on 09/26/2016).

Song, Jae et al. (2015). *Firming Up Inequality.* Working Paper 21199. National Bureau of Economic Research. (Visited on 09/14/2015).

Torres, Sónia et al. (2013). *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). (Visited on 08/10/2016).

Toussaint-Comeau, Maude (2008). "Do ethnic enclaves and networks promote immigrant self-employment?" In: *Economic Perspectives* 32.4, pp. 30–50. ISSN: 1048115X.

Webber, Douglas A. (Apr. 2013). *Firm-Level Monopsony and the Gender Pay Gap.* SSRN Scholarly Paper ID 2254196. Rochester, NY: Social Science Research Network. (Visited on 09/21/2017).

CURRICULUM VITAE

[redacted]

[redacted]

[redacted]