

Glass Ceilings or Glass Doors? Wage Disparity Within and Between Firms

Krishna Pendakur and Simon Woodcock*

Department of Economics, Simon Fraser University

April 30, 2008

Abstract

We investigate whether immigrant and minority workers' poor access to high-wage jobs—that is, glass ceilings—is attributable to poor access to jobs in high-wage firms, a phenomenon we call *glass doors*. Our analysis uses linked employer-employee data to measure mean- and quantile-wage differentials of immigrants and ethnic minorities, both within and across firms. We find that glass ceilings exist for some immigrant groups, and that they are driven in large measure by glass doors. For some immigrant groups, the sorting of these workers across firms accounts for as much as half of the economy-wide wage disparity they face.

JEL Codes: J15, J71, J31

Keywords: glass ceilings, wage differentials, immigration, visible minorities, quantile regression, linked employer-employee data

*We thank staff at the BC Inter-University Research Data Center and Yves Decady at Statistics Canada for their help and support with the WES. In addition, we wish to thank our friends and colleagues who have helped us in this research: Christian Dustmann, Nicole Fortin and Jane Friesen; and seminar participants at HEC Montreal. This research was supported by Metropolis British Columbia and by the Canadian Labour Market and Skills Researcher Network. Correspondence to: Department of Economics, Simon Fraser University, 8888 University Dr., Burnaby, BC V5A 1S6, Canada.

1 Introduction

Anti-discrimination policy takes many forms. For example, affirmative action policies primarily target the hiring decision, whereas pay equity policies target wages directly. The extent to which these policies improve the labor market outcomes of disenfranchised workers depends on what kind of barriers these workers face. Affirmative action will be most effective if disenfranchised workers face barriers to employment at “good” firms. Pay equity will be most effective when workers face wage disparity within firms. We investigate whether immigrant and ethnic minority workers’ adverse wage outcomes are driven by poor access to high-wage firms or poor access to high-wage jobs within firms.

Wage disparity can also take many forms. Between-group differences in conditional mean wages have been studied extensively. Recently, researchers have begun to focus on other aspects of the conditional wage distribution. Albrecht et al. (2003), Arulampalam et al. (2007), de la Rica et al. (forthcoming), Pendakur and Pendakur (2006) and others have found evidence that disadvantaged workers in several countries face a *glass ceiling*: a barrier that limits access to high-wage jobs. Because of data limitations, however, these papers cannot distinguish whether these outcomes are driven by poor access to jobs at high-wage firms, or poor wage outcomes within firms.

In this paper, we introduce the idea of a *glass door*. This is a barrier that limits disadvantaged workers’ access to employment at high-wage firms. Just as a glass ceiling truncates the distribution of wages that disadvantaged workers face, a glass door truncates the distribution of firms at which they might find employment. Our main objective is to assess the extent to which exclusion from high-wage jobs – that is, a glass ceiling – is driven by a glass door. To do so, we develop a summary measure of how much the sorting of disadvantaged workers across firms decreases (or increases) their wages. We call this the glass door effect. We show how to measure the glass door effect on average wages, as well as its effect at a given quantile of the wage distribution. In fact, under a particular model of wage determination, the glass door effect is the difference between a within-firm conditional (mean or quantile) wage gap measure, and a corresponding economy-wide measure. As a consequence, it is straightforward to test whether the glass door effect is zero.

Investigating glass doors requires data on multiple employees of multiple firms. To our knowledge, there are no such U.S. survey data. Some administrative data (e.g., the US Census Bureau’s

confidential Longitudinal Employer-Household Dynamics Program database) link employees to their employers, but they contain only limited information about the characteristics of workers, their jobs, and their employers. Consequently, our investigation uses a recent Statistics Canada survey of workers and their employers: the Workplace and Employee Survey (WES). The WES surveys multiple employees of firms, so we can compare wage outcomes of ethnic minority and/or immigrant workers with their Canadian-born white coworkers.

We find that Canadian-born ethnic minority workers do not face significant glass ceilings or glass doors. In contrast, immigrants face substantial conditional mean and conditional quantile wage gaps. Furthermore, these wage gaps are largely accounted for by glass doors.

2 Glass Ceilings, Sticky Floors and Glass Doors

A *glass ceiling* is a barrier to labour market success that operates at the upper end of the wage distribution. For example, a glass ceiling might limit some workers' access to management jobs. This phenomenon has been studied mainly in the context of the male-female wage gap. Albrecht et al. (2003) suggest a method to detect a glass ceiling via quantile regression. They reason that because a glass ceiling limits the ability of disadvantaged workers to "rise to the top," it introduces large wage gaps at upper quantiles of the wage distribution. However, a glass ceiling has little (if any) effect on wage gaps at lower quantiles, so that this pattern is easy to detect by comparing quantile regression estimates at different quantiles. In our work, we distinguish between economy-wide glass ceilings, which limit access to high-wage jobs in the economy as whole, from within-firm glass ceilings, which limit access to high-wage jobs within firms.

Albrecht et al. (2003) found evidence that Swedish women face an economy-wide glass ceiling. Subsequent authors have identified other patterns in wage differentials across quantiles. For example, de la Rica et al. (forthcoming) and Arulampalam et al. (2007) find evidence that women in some European countries face larger wage differentials at the bottom of the conditional wage distribution than at upper quantiles. Pendakur and Pendakur (2006) observe the same for some ethnic minorities in Canada. These authors call this a *sticky floor*, because it is evidence that disadvantaged workers are crowded into very low-wage jobs. de la Rica et al. (forthcoming) argue that a sticky floor could be a consequence of statistical discrimination if employers believe a group of workers are at high risk of quitting. Pendakur and Pendakur (2006) suggest a sticky floor could arise if rents are larger at the bottom of the wage distribution than the top. A

sticky floor could also arise if anti-discrimination policy is most effective at the top of the wage distribution, e.g., if policy primarily targets minority representation in management jobs.

Previous work in this area has focused on economy-wide glass ceilings (and sticky floors) because available data could not distinguish within-firm from economy-wide wage outcomes. Of course, the economy-wide patterns observed by other authors might arise because disadvantaged workers face a glass ceiling (or sticky floor) within firms. However, Abowd et al. (1999), Woodcock (2007), and others have shown that inter-firm wage differences account for about one third of all wage variation. Thus, an alternative explanation is that these workers do not face within-firm glass ceilings, but some barrier limits their ability to obtain employment at high-wage firms. We call this a *glass door*.

We look for evidence of a glass door by comparing within-firm wage gaps to economy-wide wage gaps (i.e., wage gaps that average over firms). In the next section, we show that this difference – which we call the glass door effect – has a specific and useful interpretation under a particular model of wage determination. Intuitively, if disadvantaged workers have better wage outcomes within firms than they do economy-wide, it indicates that their low wage outcomes are partly the result of how they sort across firms.

Our investigation of the glass door effect considers three features of the conditional wage distribution: conditional means, conditional quantiles and conditional representation. We examine the effect of glass doors on conditional mean wages to assess whether or not immigrant and ethnic minority workers are, on average, employed in firms that pay lower wages than their Canadian-born white counterparts. We estimate the glass door effect at various quantiles to assess whether glass doors contribute to economy-wide glass ceilings and sticky floors. We also measure the representation of immigrant and ethnic minority workers in tails of the economy-wide and within-firm conditional wage distributions. This allows us to directly assess the consequences of glass ceilings and glass doors. Whereas it may be hard to interpret whether or not a wage gap at the top decile is of policy importance, knowing the magnitude of minority under-representation in the top decile may connect more directly with policy discussion.

Aydemir and Skuterud (forthcoming) also use the WES to investigate the role of firms in determining immigrant wage outcomes. They find evidence that male immigrants sort primarily into low-paying establishments, whereas female immigrants sort primarily into low-paying jobs within establishments. There are important differences between their study and this one. First, they focus exclusively on conditional mean wages of immigrants. Consequently, their results are

not informative of the interplay between glass ceilings, sticky floors, and glass doors; they are also not informative of the outcomes of non-immigrant ethnic minorities. More subtly, their measure of the contribution of inter-firm sorting to immigrant wage gaps confounds the consequences of sorting on immigrant status with sorting on other observable characteristics. We return to this point in the next section, after discussing how we measure the glass door effect.

3 Methodology

We begin by comparing average (log) wages of minority (i.e., ethnic minority or immigrant) and majority workers using the linear regression model:

$$E[y_i|x_i, g_i] = x_i'\beta + g_i'\delta. \tag{1}$$

Here, y_i is the log wage of worker i ; x_i is a vector of characteristics that affect wages (e.g., educational attainment, labor market experience, etc.); β measures the returns to those characteristics; g_i is a vector of indicator variables for membership in a minority group; and δ measures the difference in average log wages of minority and majority workers who share the same observed characteristics x_i .

To investigate whether minority workers face an economy-wide glass ceiling or sticky floor, we estimate wage gaps at several quantiles of the conditional wage distribution. We measure the wage gap at the τ^{th} conditional quantile using the quantile regression that satisfies:

$$\Pr[y_i \leq x_i'\beta^\tau + g_i'\delta^\tau | x_i, g_i] = \tau, \tag{2}$$

where β^τ measures the returns to characteristics at the τ^{th} quantile, and δ^τ measures the difference between the τ^{th} quantile of log wages of minority and majority workers, conditional on x_i . Estimates of δ^τ at several quantiles illustrate how wage differentials vary over the conditional wage distribution.

Wage differentials at various quantiles are only partly illustrative of minority workers' access to high-wage jobs. An alternative is to measure the proportion of minority workers in regions of the wage distribution, especially the tails. Following Pendakur, Pendakur, and Woodcock (2008), we call the proportion of a minority group's workers whose wages fall into a region of the population conditional wage distribution their *representation* in that region. Representation

of minority workers in the lower tail of the wage distribution meaningfully quantifies possible crowding into low-wage jobs. Likewise, representation in the upper tail meaningfully quantifies possible exclusion from high-wage jobs.

The representation index of Pendakur, Pendakur, and Woodcock (2008) is intuitive and straightforward to construct. We estimate quantiles of the population wage distribution, conditional on characteristics x_i , from the quantile regression that satisfies:

$$\Pr [y_i \leq x_i' \beta^\tau] = \tau. \quad (3)$$

Note that eq. (3) does not condition on group membership g_i . We use coefficient estimates from eq. (3), $\hat{\beta}^\tau$, to estimate quantiles of the population conditional wage distribution, $\hat{y}_i^\tau = x_i' \hat{\beta}^\tau$, for each worker given their characteristics. The conditional representation of group g below the τ^{th} quantile is

$$r_g^\tau = \frac{1}{N_g} \sum_{i \in g} I(y_i \leq \hat{y}_i^\tau) \quad (4)$$

where N_g is the number of members of group g and I denotes the indicator function. This measures the proportion of workers in group g who earn less than the τ^{th} quantile of the population conditional wage distribution, given their characteristics x_i . We can define representation above the τ^{th} conditional quantile analogously. If $r_g^\tau > \tau$, the proportion of the group's members below the τ^{th} quantile exceeds the population proportion, and we say the group is *over-represented* in that region. Likewise, if $r_g^\tau < \tau$, we say the group is *under-represented* in that region.

Figure 1 illustrates the relationship between quantile wage differentials and the representation index for a particular x_i . The figure shows the cumulative distribution function of wages for a hypothetical population (F_{pop}), a reference group (F_0), and a minority group (F_g). At a given quantile τ , the quantile wage differential between group g and the reference group, δ_g^τ , is the horizontal distance between F_g and F_0 . At the τ^{th} population quantile, y_τ , the vertical distance between F_g and F_0 is the difference between representation of the two groups, $r_g^\tau - r_0^\tau$.

3.1 Glass Door Effects

We estimate glass door effects by comparing economy-wide wage gaps, which average over firms, to within-firm wage gaps. The difference between these estimates is informative of the extent to which wage outcomes are driven by how minority workers sort across employers, versus how

they sort into jobs within employers.

Our within-firm measures condition on employer identity. That is, we augment the regression models described above with firm effects. These encompass both observed and unobserved employer characteristics. In the mean regression case, equation (1) becomes:

$$E[y_i|x_i, g_i, f_i] = x_i'\beta + g_i'\delta + f_i'\psi \quad (5)$$

where f_i is a vector of indicators for each firm and ψ is a vector of firm effects. The firm effects measure inter-firm differences in average wages, conditional on characteristics x_i and group membership g_i .

If the true firm effects (i.e., their population values, as opposed to sample estimates) were observable, we could summarize the glass door effect by regressing $f_i'\psi$ on x_i and g_i . The coefficient on g_i in this regression would measure the average firm effect of each group's members, conditional on their characteristics x_i . This is exactly the glass door effect we seek: a measure of how inter-group differences in wages are affected by inter-group differences in sorting across firms, conditional on characteristics. Of course we can't estimate this regression directly, because the true firm effects are not observable. However, the following proposition shows that we obtain an unbiased estimate of the glass door effect by comparing estimated wage differentials in specifications with and without firm effects. This is not surprising: it simply specializes the well-known result that omitted variable bias can be represented as least squares coefficients in an artificial regression (see, e.g., Greene (2003, pp. 148-149)). Furthermore, this result holds for more general specifications of inter-group conditional wage differences; see the discussion following the appendicized proof.

Proposition 1 (Glass Door Effect) *Assume equation (5) is correctly specified. Let $\hat{\delta}$ and $\hat{\delta}^*$ be the estimated coefficients on group membership in regressions that include and exclude firm effects, respectively. Then $(\hat{\delta}^* - \hat{\delta})$ is an unbiased estimator of δ_f in the regression $E[f_i'\psi|x_i, g_i] = x_i'\beta_f + g_i'\delta_f$.*

Proof. See Appendix ■

Our measure of the glass door effect differs from how Aydemir and Skuterud (forthcoming) measure the effect of inter-firm sorting on mean wage gaps. They compare the average firm effect of immigrants to that of native-born workers. This is the sample analog of the coefficient on g_i in the hypothetical regression of $f_i'\psi$ on g_i only. Because characteristics, x_i , are omitted

from this hypothetical regression, their measure confounds the returns to sorting on observable characteristics (e.g., education or occupation) with sorting on minority group status. In contrast, our measure conditions out the returns to sorting on these observable characteristics.

It is important to note that a zero glass door effect does not imply that firm effects do not belong in the model. Rather, it implies that firm effects are unrelated to group membership, conditional on worker characteristics. The following proposition shows that we can test for the presence of a glass door effect using a Hausman test. The Hausman test applies because under the null of no glass door effect, specifications with and without firm effects both yield consistent estimates of δ , but the estimate with firm effects is inefficient. Under the alternative, however, only the specification with firm effects yields a consistent estimate.

Proposition 2 (Glass Door Test) *Under the null hypothesis that equation (5) is correctly specified with spherical errors, but firm effects are conditionally unrelated to group membership, that is, $H_0 : \delta_f = 0$, the variance of the estimated glass door effect is $\text{Var} [\widehat{\delta}^* - \widehat{\delta}] = \text{Var} [\widehat{\delta}] - \text{Var} [\widehat{\delta}^*]$. Consequently, $Q = (\widehat{\delta}_g^* - \widehat{\delta}_g)' (\text{Var} [\widehat{\delta}_g] - \text{Var} [\widehat{\delta}_g^*])^{-1} (\widehat{\delta}_g^* - \widehat{\delta}_g) \stackrel{a}{\sim} \chi_1^2$ under H_0 , where $\widehat{\delta}_g$ and $\widehat{\delta}_g^*$ are the elements of $\widehat{\delta}$ and $\widehat{\delta}^*$, respectively, corresponding to group g .*

Proof. See Appendix ■

Remark 3 *The variance terms that appear in Proposition 2 are actual (population) variances of the estimated parameters. Under the null, firm effects are conditionally unrelated to group membership, but may still belong in the model. Consequently, the regression that excludes firm effects is mis-specified and standard regression output of its estimated coefficient variances is incorrect. The specification that includes firm effects yields a consistent estimate of the error variance, so we re-scale the sample estimate of $\text{Var} [\widehat{\delta}^*]$ by s^2/s^{2*} , where s^2 and s^{2*} are estimated regression error variances in models with and without firm effects, respectively. This adjustment is conservative, since it increases the estimated variance of the glass door effect.*

Proof. See Appendix ■

In the linear model, eq. (5), the firm effect is common to all employees of the firm. It is therefore usually conceptualized as a pure location shift of the conditional wage distribution. That is, the shape of the wage distribution, conditional on x_i and g_i , is the same at every firm, but its location (mean) differs across firms. We implement firm effects in the quantile regression, eq. (6), the same way:

$$\Pr [y_i \leq x_i' \beta^\tau + g_i' \delta^\tau + f_i' \psi | x_i, g_i, f_i] = \tau. \quad (6)$$

That is, we restrict each firm’s effect to be the same at every quantile of the conditional wage distribution. (Since, in our application, most firms have fewer than 10 surveyed employees, quantile-specific firm effects would be imprecisely estimated.) We implement the restriction by simultaneously estimating regressions for multiple quantiles and imposing cross-equation restrictions on the firm effects, as suggested by Koenker (2004).

Introducing firm effects in the quantile regressions poses some computational challenges. The large number of effects to be estimated, and the implied sparseness of the design matrix of firm effects, pose numerical problems for the linear programming algorithms that are usually used to estimate quantile regressions. So we estimate eq. (6) using Koenker and Ng’s (2005) Frisch-Newton algorithm for sparse quantile regression, for which they provide estimation subroutines in R. We also estimate a version of eq. (6) that excludes group membership indicators to measure within-firm representation.

The glass door effect and glass door test of Propositions 1 and 2 have natural analogs in the quantile regression setting. In particular, Angrist et al. (2006) show that quantile regression can be expressed as a weighted least squares problem with particular quantile-specific weights. They also show that omitted variable bias in a quantile regression can be interpreted as a weighted least squares problem. In particular, if eq. (6) is the true conditional quantile function, then the difference between estimates of δ^τ and $\delta^{\tau*}$, where $\delta^{\tau*}$ is the coefficient vector in the quantile regression excluding firm effects, has an interpretation similar to the linear regression case. Just like the linear case, $(\delta^{\tau*} - \delta^\tau)$ estimates the coefficient on group membership in a hypothetical least squares regression of $f_i\psi$ on x_i and g_i . Unlike the linear case, however, this hypothetical regression uses quantile-specific weights. At upper quantiles of the conditional wage distribution, the weights are large (small) for employees of firms with large (small) ψ , and conversely at lower quantiles of the conditional wage distribution. For the exact description of these weights, see the discussion following the proof of Proposition 1 in the Appendix.

We can test for the glass door effect in the quantile setting as well. As in the linear regression case, it is natural to use a Hausman test. Under the Null hypothesis, models with and without firm effects both yield consistent estimates of δ^τ , but the model with firm effects is inefficient. Under the alternative, only the model with firm effects yields a consistent estimate of δ^τ . As in the linear model case, the specification that omits firm effects is mis-specified under the null, so we apply the same conservative adjustment to the estimated variance of the glass door effect in the quantile glass door test. Alternatively, since Angrist et al. (2006) show that estimated

quantile regression coefficients are asymptotically normal even under mis-specification, one could directly bootstrap the asymptotic variance of $(\delta^{\tau^*} - \delta^\tau)$.

4 Data

Our investigation uses the Workplace and Employee Survey (WES). The WES is one of a few linked employer-employee databases worldwide, and the only such data for Canada. In each year between 1999 and 2004, Statistics Canada surveyed a representative sample of approximately 6,000 workplaces. The initial sample of firms was refreshed in 2001 and 2003 to reflect attrition and firm births. In odd-numbered years, Statistics Canada randomly sampled approximately 20,000 employees of these firms. The number of workers sampled from each firm was proportional to size except in small firms, wherein all employees were sampled. Sampled workers were surveyed in that year and the next, and a new sample of workers was drawn in the next odd-numbered year. Because we know the identity of each worker’s employer, the WES is ideal to assess the role of firms in determining wage outcomes.

Our analysis is based on the pooled 1999, 2001, and 2003 cross-sections. We do not use data from even-numbered years for two reasons. First, employee attrition is high in their second survey year and is likely non-random. Second, many sampled workers change employer between survey years and only limited information is collected about their new employer. Also, note that pooling the cross-sections maximizes the number of observed employees of each workplace.

We restrict the sample to non-Aboriginal (i.e., non-Native Indian) workers between 25 and 64 years of age. The restricted sample comprises 58,298 employees of 7,641 workplaces. We observe between one and 55 employees of each firm; the mean number is 7.6, and the median is 6. Estimates that condition on firm effects implicitly exclude firms with only one observed employee, in which case the mean number of employees per firm is 8.4, and the median is 7. We observe 3,064 firms in all three survey years, 2,063 firms in two years, and the remaining 2,514 firms in only one year.

Our outcome measure is the natural logarithm of hourly wages. We estimate several specifications to assess the contribution of observable characteristics, x_i , and employer identity, f_i , to the wage outcomes of immigrants and ethnic minorities. In the main text, we focus on specifications that control for: sex, highest level of schooling (8 categories), marital status (6 categories), number of children (5 categories), a quartic in years of full-time-equivalent labour market experience,

province of residence (7 categories, with maritime provinces combined into a single category), an indicator for residence in a Census metropolitan area, an indicator for full-time employment, occupation (6 categories), an indicator for membership in a union or collective bargaining agreement, and a quadratic in years of seniority with the current employer. A supplementary appendix provides data descriptives and estimation results from alternate specifications.

We estimate wage differentials and representation for six groups, g_i , of men and women. Our six groups are based on whether or not a person is an ethnic minority, an immigrant, and (for immigrants) time since arrival. Anti-discrimination policy in Canada, including both federal Employment Equity and provincial Pay Equity legislation, targets “visible minorities.” Legislation defines these as non-Aboriginal people whose ancestry includes any ethnic or national group other than Canada, the USA, Australia, New Zealand, or any European country. (Aboriginal people constitute a separate equity group.) Essentially, Canadian law targets non-white ethnic minorities. We define our indicator of ethnic minority status likewise. We define immigrants as individuals born outside Canada; and distinguish between recent immigrants (10 or less years since arrival in Canada) and non-recent immigrants. The reference group in all cases is Canadian-born individuals who are not visible minorities (that is, who are white). Consequently, our six groups of workers are: Canadian-born white workers, Canadian-born visible minority workers, recent white immigrants, recent visible minority immigrants, non-recent white immigrants, and non-recent visible minority immigrants.

We estimate all specifications that exclude firm effects separately by sex. We pool men and women in specifications that include firm effects. In this case, we interact all covariates (except the firm identifiers) with sex. This restricts firm effects to be the same for all employees of a firm, so they reflect a pure location shift of the conditional wage distribution.

We estimate all specifications using employee sample weights provided by Statistics Canada. Although our quantile regression specifications with firm effects are estimated using subroutines optimized for the highly parameterized and sparse nature of the problem, we still face computational constraints that limit the number of firms we can include in the quantile regressions. This forces us to subsample from the data. Our reported estimates for quantile regressions with firm effects are averaged over 50 subsamples. Each subsample consists of all surveyed employees of 1,500 randomly sampled firms. We sample firms with replacement in each subsample. To increase precision of the estimated firm effects, we sample firms with probabilities proportional to the number of surveyed employees and adjust the employee sample weights accordingly.

Throughout, we estimate standard errors following Statistics Canada’s recommended procedure, using 100 sets of provided bootstrap sample weights. In the case of quantile regressions with firm effects, including the quantile regressions that underlie the within-firm representation index, we bootstrap the entire estimation procedure (including subsampling and averaging over subsamples) for each set of bootstrap weights.

5 Results

Table 1 presents our estimates for Canadian-born visible minorities. These workers face substantial within-firm mean wage differentials: about $-.05$ log points for men and $-.06$ log points for women. However, they do not face large glass doors: the mean glass door effect is $-.02$ log points for women in this group, and $.037$ log points for men. The glass door test statistic is negative for men, which is a fairly common finite-sample occurrence for Hausman tests and is usually interpreted as rejection of the null. Evidently these men are employed at higher-wage firms than their white counterparts, on average, and consequently face no economy-wide wage gap. In contrast, women sort into slightly lower-paying firms than their white counterparts, which accounts for about one quarter of the economy-wide wage gap they face.

There is no strong evidence that Canadian-born visible minorities face an economy-wide glass ceiling or sticky floor. Women in this group fare worse in the upper half of the wage distribution than they do in the lower half, but the disparity hits too low in the distribution (somewhere below the median) to characterize this as a glass ceiling in the classic sense. This upper-tail wage disparity does, however, appear to reflect inter-firm sorting. The point estimates suggest that the quantile glass door effect is positive at lower quantiles but negative at upper quantiles, so that glass doors contribute to the economy-wide pattern in quantile wage disparity. However, the quantile glass door effects are either small in magnitude or statistically insignificant at all reported quantiles.

Tables 2 and 3 present our estimates for immigrants to Canada. Both recent (Table 3) and non-recent (Table 2) immigrants face substantial economy-wide and within-firm mean wage gaps in comparison to Canadian-born white workers. Unsurprisingly, mean wage gaps are larger for recent immigrants than non-recent immigrants, and larger for visible minority immigrants than white immigrants. Immigrant workers also face substantial glass doors that account for between one quarter and one half of the economy-wide wage gaps they face. Mean glass door effects

are larger for recent immigrants than non-recent immigrants, and larger for visible minority immigrants than white immigrants. The former suggests that inter-firm sorting is one avenue whereby immigrant wage outcomes improve with time spent in the host country.

Non-recent male immigrants to Canada face an economy-wide glass ceiling and are consequently under-represented in the upper decile of the wage distribution. The under-representation is particularly severe for visible minorities: only 4.1% earn wages in the upper decile of the conditional wage distribution. Glass doors contribute to the economy-wide glass ceiling these men face by reducing wages more in the upper part of the wage distribution than at the bottom. The role of glass doors is particularly large for visible minorities: glass doors reduce their wages by 0.033 log-points at the bottom decile, and by nearly twice this much at the top decile. In contrast, non-recent female immigrants do not face economy-wide glass ceilings, and face glass door effects that are similar across the quantiles.

Recent male immigrants and recent white female immigrants face economy-wide wage gaps that are larger at lower quantiles of the wage distribution than at upper quantiles, which suggests they face a sticky floor. As a consequence, these workers are substantially over-represented in the bottom decile. For visible minority recent immigrant men, this over-representation is severe: 28.9% earn wages in the bottom decile of the conditional wage distribution. Glass doors contribute substantially to the economy-wide sticky floor by reducing wages more at the bottom of the wage distribution than the top. In the case of recently arrived visible minority men, glass doors account for the entire economy-wide sticky floor: these men face within-firm wage gaps that are effectively constant across quantiles.

The fact that recent male immigrants, especially visible minorities, face a sticky floor and non-recent male immigrants face a glass ceiling is telling. It suggests that recently arrived men sort initially into very low-wage jobs. Over time, their outcomes improve as they move up the wage distribution. Eventually, however, they hit a glass ceiling that prevents further progress. At each stage, inter-firm sorting is important – exacerbating the sticky floor they face shortly after arrival, and contributing to the glass ceiling they face later on.

References

- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67(2), 251–334.

- Albrecht, J., A. Bjorklund, and S. Vroman (2003). Is there a glass ceiling in Sweden? *Journal of Labor Economics* 21(1), 145–177.
- Angrist, J., V. Chernozhukov, and I. Fernández-Val (2006). Quantile regression under misspecification, with an application to the U.S. wage structure. *Econometrica* 74(2), 539–563.
- Arulampalam, W., A. L. Booth, and M. L. Bryan (2007). Is there a glass ceiling over Europe? Exploring the gender pay gap across the wages distribution. *Industrial and Labor Relations Review* 60(2), 163–186.
- Aydemir, A. and M. Skuterud (forthcoming). The immigrant wage differential within and across establishments. *Industrial and Labor Relations Review*.
- de la Rica, S., J. J. Dolado, and V. Llorens (forthcoming). Ceilings or floors? gender wage gaps by education in Spain. *Journal of Population Economics*.
- Greene, W. H. (2003). *Econometric Analysis* (5th ed.). New Jersey: Prentice Hall.
- Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis* 91, 74–89.
- Koenker, R. and P. Ng (2005). A Frisch-Newton algorithm for sparse quantile regression. *Acta Mathematicae Applicatae Sinica (English Series)* 21(2), 225–236.
- Pendakur, K. and R. Pendakur (2006). Glass ceilings for ethnic minorities.
- Pendakur, K., R. Pendakur, and S. Woodcock (2008, January). Representation and severity in the distribution of income. Mimeo.
- Woodcock, S. D. (2007). Match effects. SFU Economics Discussion Paper dp07-13.

6 Appendix

6.1 Proofs of Propositions

Proof of Proposition 1. Suppose that eq. (5) is correctly specified, and rewrite it in matrix notation as $E[y|X, G, F] = X\beta + G\delta + F\psi$. The least squares estimator of δ in this regression is unbiased, i.e., $E[\widehat{\delta}] = \delta$. When the estimated equation omits firm effects, however, standard results for omitted variables bias imply $E[\widehat{\delta}^*] = \delta + (G'M_XG)^{-1}G'M_XF\psi$, where $M_Z = I - Z(Z'Z)^{-1}Z'$ for some matrix Z and I is the identity matrix. Consequently, $E[\widehat{\delta}^* - \widehat{\delta}] = (G'M_XG)^{-1}G'M_XF\psi$, which is exactly the least squares estimate of δ_F in the regression $E[F\psi|X, G] = X\beta_F + G\delta_F$. ■

The intuition underlying the proof of Proposition 1 is straightforward. Note that M_Z is idempotent, so we can rewrite the bias term as $(G'M_XG)^{-1}G'M_XF\psi = (\widetilde{G}'\widetilde{G})^{-1}\widetilde{G}'F\psi$, where $\widetilde{G} = M_XG$ is a matrix of residuals in the least squares regression of the group membership indicators on X . Therefore $(\widetilde{G}'\widetilde{G})^{-1}\widetilde{G}'F\psi$ can be interpreted as estimated coefficients in the least squares regression of $F\psi$ on the component of group membership that is unrelated to individual characteristics, \widetilde{G} . In the simple case where X is orthogonal to G and F , so that $G'M_XG = G'G$ and $G'M_XF = G'F$, the bias is $(G'G)^{-1}G'F\psi$. In this case, bias in the group g wage differential, $\widehat{\delta}_g^*$, is the average firm effect of all employers of group g 's members: $E[\widehat{\delta}_g^*] - \delta_g = N_g^{-1}\sum_{i \in g}\psi_{(i)}$, where N_g is the sample number of workers in group g , and $\psi_{(i)}$ is the firm effect of worker i 's

employer. More generally, the bias is interpreted as the average firm effect of group g 's employers, net of the component that is explained by sorting on observable characteristics, X .

Our measure of the glass door effect is easily extended beyond the simple dummy-variable specification of group membership. Consider the fully interacted model where $w'_i = [1 \ g'_i] \otimes x'_i$, and let W be the matrix with rows w'_i . Assume that the correctly specified model is $E[y|X, G, F] = W\beta + F\psi$. In this model, the within-firm wage gap depends on X , and can be evaluated by comparing predicted values for members of a particular group to predicted values for the reference group. The least squares estimator of β in this regression is unbiased. When the estimated equation omits firm effects, the omitted variables bias is $E[\widehat{\beta}^* - \beta] = (W'W)^{-1}W'F\psi$, which is the coefficient in an artificial regression of $F\psi$ on W . Just as the within-firm wage gap depends on X , so too does the glass door effect, which is evaluated by comparing predicted values of $w'_i(\widehat{\beta}^* - \beta)$ for members of each group to predicted values for the reference group.

To see the quantile analog of Proposition 1, let $\theta^{\tau*}$ and θ^τ be coefficient vectors on $[X \ G]$ in quantile regressions that exclude and include firm effects, respectively, at the τ^{th} quantile. Angrist et al. (2006, p. 547) show that the bias due to omitted firm effects takes the form

$$(\theta^{\tau*} - \theta^\tau) = E \left[[X \ G]' \tilde{W}_\tau ([X \ G \ F]) [X \ G] \right]^{-1} E \left[[X \ G]' \tilde{W}_\tau ([X \ G \ F]) F\psi \right]$$

where $\tilde{W}_\tau ([X \ G \ F]) = \text{diag}(\tilde{w}_\tau ([X \ G \ F]))$, $\tilde{w}_\tau ([X \ G \ F]) = \frac{1}{2} \int_0^1 f_{\varepsilon_\tau} (u ([X \ G] \theta^{\tau*} - [X \ G \ F] [\theta^{\tau'} \ \psi']')) du$, $\varepsilon_\tau = Y - [X \ G \ F] [\theta^{\tau'} \ \psi']'$, and f_{ε_τ} is the conditional density of ε_τ given $[X \ G \ F]$. Given the weights, this bias term is the coefficient vector in a weighted least squares regression of $F\psi$ on $[X \ G]$. The weights are the same as those in the weighted least squares representation of the mis-specified quantile regression of Y on $[X \ G]$.

Proof of Proposition 2. When equation (5) is correctly specified and errors are spherical, then $y = X\beta + G\delta + F\psi + \varepsilon$ where ε is the error vector satisfying $E[\varepsilon|X, G, F] = 0$ and $E[\varepsilon\varepsilon'|X, G, F] = \sigma^2 I$. The estimators are $\widehat{\delta}^* = \delta + (G'M_X G)^{-1} G'M_X (F\psi + \varepsilon)$ and $\widehat{\delta} = \delta + (G'M_{[X \ F]} G)^{-1} G'M_{[X \ F]} \varepsilon$. Under the null, $\delta_f = (G'M_X G)^{-1} G'M_X F\psi = 0$ which implies $\widehat{\delta}^* = \delta + (G'M_X G)^{-1} G'M_X \varepsilon$ and

$$\begin{aligned} \text{Cov} \left[\widehat{\delta}^*, \widehat{\delta} \right] &= E \left[(G'M_X G)^{-1} G'M_X \varepsilon \varepsilon' M_{[X \ F]} G (G'M_{[X \ F]} G)^{-1} \right] \\ &= \sigma^2 (G'M_X G)^{-1} G'M_X M_{[X \ F]} G (G'M_{[X \ F]} G)^{-1}. \end{aligned}$$

Note that $G'M_X M_{[X \ F]} = G'M_{[X \ F]}$. The intuition for this result is straightforward. $G'M_X$ gives the residuals from least squares regression of G on X . Regressing these on X and F yields residuals $(G'M_X M_{[X \ F]})$ that are the same as those obtained from the regression of G on X and F directly $(G'M_{[X \ F]})$. It follows that $\text{Cov} \left[\widehat{\delta}^*, \widehat{\delta} \right] = \sigma^2 (G'M_X G)^{-1} = \text{Var} \left[\widehat{\delta}^* \right]$ under the null, so that $\text{Var} \left[\widehat{\delta}^* - \widehat{\delta} \right] = \text{Var} \left[\widehat{\delta}^* \right] + \text{Var} \left[\widehat{\delta} \right] - 2\text{Cov} \left[\widehat{\delta}^*, \widehat{\delta} \right] = \text{Var} \left[\widehat{\delta} \right] - \text{Var} \left[\widehat{\delta}^* \right]$. Under standard regularity conditions, $\widehat{\delta}$ and $\widehat{\delta}^*$ are asymptotically normal and hence so is their difference, so that $Q \sim \chi_1^2$ under the null. ■

Proof of Remark 3. Let N denote the number of observations, k^* denote the rank of $[X \ G]$, and e^* denote the least squares residual vector in the mis-specified model. Then the usual estimator of the error variance in the mis-specified model, $s^{*2} = e^{*'} e^* / (N - k^*)$, has expectation $E[s^{*2}] = \sigma^2 + (\psi' F' M_{[X \ G]} F \psi) / (N - k^*)$. The term in parentheses (the bias) is non-negative. ■

Figure 1: Quantile Wage Differentials and Representation

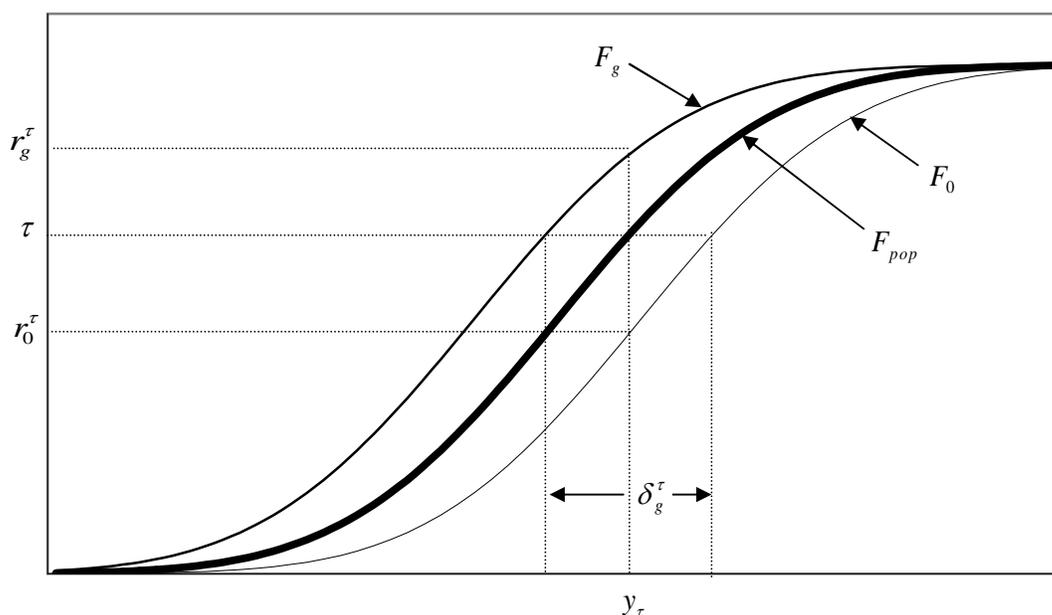


TABLE 1
ESTIMATES FOR CANADIAN-BORN VISIBLE MINORITIES

	MEN			WOMEN		
	Economy Wide (1)	Within Firms (2)	Glass Door Effect (3)	Economy Wide (4)	Within Firms (5)	Glass Door Effect (6)
Mean wage differential	-0.012 (0.019)	-0.049** (0.012)	0.037 [-]	-0.083** (0.021)	-0.063** (0.019)	-0.020* [0.069]
Quantile differential						
10th percentile	-0.026 (0.023)	-0.083** (0.025)	0.057** [0.002]	0.012 (0.048)	-0.065 (0.067)	0.077 [0.179]
Median	-0.029 (0.034)	-0.024 (0.039)	-0.006 [0.847]	-0.115** (0.038)	-0.088** (0.025)	-0.027 [-]
90th percentile	0.079 (0.060)	-0.028 (0.033)	0.107 [-]	-0.084** (0.018)	-0.022 (0.047)	-0.062 [0.174]
Representation						
Below 10th percentile	0.102 (0.006)	0.112 (0.014)		0.082** (0.000)	0.114 (0.021)	
Below median	0.491 (0.006)	0.508 (0.030)		0.604** (0.031)	0.605** (0.037)	
Above 90th percentile	0.145** (0.011)	0.107 (0.014)		0.075** (0.004)	0.116 (0.021)	
s^2	0.390	0.288		0.349	0.288	

Notes: Standard errors are in parentheses, p-values for glass door test are in brackets. Missing p-value indicates a negative Hausman test statistic. Reference category for differentials is Canadian born, white. For differentials and glass door test, ** indicates statistically significant at 5% level and * indicates statistically significant at 10% level. For representation, ** indicates difference from population quantile is statistically significant at 5% level and * indicates this difference is statistically significant at 10% level. Columns (1) and (4) are based on separate regressions for 32,898 men and 25,400 women. Columns (2) and (5) are based on pooled regressions for men and women, with all controls except firm effects interacted with sex, averaged over 50 random samples of 1,500 firms with at least 2 employees. Mean number of observations in the 50 random samples is 17,906.

TABLE 2
ESTIMATES FOR NON-RECENT IMMIGRANTS (> 10 YEARS SINCE IMMIGRATION)

	MEN			WOMEN		
	Economy Wide (1)	Within Firms (2)	Glass Door Effect (3)	Economy Wide (4)	Within Firms (5)	Glass Door Effect (6)
WHITE						
Mean wage differential	-0.067** (0.008)	-0.047** (0.006)	-0.020** [0.000]	-0.039** (0.008)	-0.019** (0.006)	-0.020** [0.000]
Quantile differential						
10th percentile	-0.059** (0.022)	-0.051** (0.014)	-0.007 [-]	-0.058** (0.015)	-0.036 (0.029)	-0.022 [0.398]
Median	-0.066** (0.008)	-0.032* (0.016)	-0.035** [0.023]	-0.041** (0.010)	-0.027** (0.009)	-0.015** [0.002]
90th percentile	-0.075** (0.012)	-0.033** (0.015)	-0.042** [0.000]	-0.038** (0.013)	-0.011 (0.021)	-0.026 [0.149]
Representation						
Below 10th percentile	0.105 (0.006)	0.109 (0.015)		0.113 (0.012)	0.104 (0.006)	
Below median	0.530** (0.009)	0.508 (0.010)		0.513 (0.008)	0.520** (0.009)	
Above 90th percentile	0.081** (0.003)	0.096 (0.005)		0.081* (0.011)	0.094 (0.007)	
VISIBLE MINORITY						
Mean wage differential	-0.192** (0.010)	-0.098** (0.007)	-0.094** [0.000]	-0.101** (0.009)	-0.051** (0.010)	-0.050** [0.000]
Quantile differential						
10th percentile	-0.112** (0.034)	-0.080** (0.021)	-0.033** [-]	-0.107** (0.017)	-0.046** (0.013)	-0.061** [0.000]
Median	-0.195** (0.015)	-0.111** (0.022)	-0.084** [0.000]	-0.092** (0.010)	-0.050** (0.014)	-0.042** [0.000]
90th percentile	-0.180** (0.017)	-0.119** (0.027)	-0.061** [0.011]	-0.109** (0.016)	-0.040 (0.048)	-0.069 [0.141]
Representation						
Below 10th percentile	0.135** (0.012)	0.109 (0.010)		0.128** (0.008)	0.103 (0.007)	
Below median	0.664** (0.010)	0.580** (0.015)		0.588** (0.014)	0.518 (0.014)	
Above 90th percentile	0.041** (0.002)	0.067** (0.008)		0.068** (0.008)	0.095 (0.008)	
s^2	0.390	0.288		0.349	0.288	

Notes: Standard errors are in parentheses, p-values for glass door test are in brackets. Missing p-value indicates a negative Hausman test statistic. Reference category for differentials is Canadian born, white. For differentials and glass door test, ** indicates statistically significant at 5% level and * indicates statistically significant at 10% level. For representation, ** indicates difference from population quantile is statistically significant at 5% level and * indicates this difference is statistically significant at 10% level. Columns (1) and (4) are based on separate regressions for 32,898 men and 25,400 women. Columns (2) and (5) are based on pooled regressions for men and women, with all controls except firm effects interacted with sex, averaged over 50 random samples of 1,500 firms with at least 2 employees. Mean number of observations in the 50 random samples is 17,906.

TABLE 3
ESTIMATES FOR RECENT IMMIGRANTS (≤ 10 YEARS SINCE IMMIGRATION)

	MEN			WOMEN		
	Economy Wide (1)	Within Firms (2)	Glass Door Effect (3)	Economy Wide (4)	Within Firms (5)	Glass Door Effect (6)
WHITE						
Mean wage differential	-0.068** (0.015)	-0.028** (0.014)	-0.040** [0.000]	-0.096** (0.018)	-0.060** (0.017)	-0.036** [0.000]
Quantile differential						
10th percentile	-0.085** (0.035)	-0.089** (0.025)	0.004 [-]	-0.147** (0.026)	-0.055* (0.031)	-0.092** [0.000]
Median	-0.114** (0.017)	-0.045 (0.033)	-0.069** [0.026]	-0.106** (0.024)	-0.052** (0.015)	-0.054 [-]
90th percentile	-0.054 (0.040)	-0.056** (0.028)	0.002 [-]	-0.047 (0.076)	-0.020 (0.122)	-0.028 [0.797]
Representation						
Below 10th percentile	0.102 (0.007)	0.114 (0.012)		0.143** (0.004)	0.111 (0.014)	
Below median	0.557** (0.014)	0.495 (0.023)		0.580** (0.009)	0.537 (0.025)	
Above 90th percentile	0.095 (0.008)	0.081* (0.010)		0.103 (0.009)	0.099 (0.015)	
VISIBLE MINORITY						
Mean wage differential	-0.312** (0.014)	-0.179** (0.011)	-0.133** [0.000]	-0.220** (0.014)	-0.162** (0.013)	-0.058** [0.000]
Quantile differential						
10th percentile	-0.317** (0.025)	-0.200** (0.048)	-0.117** [0.009]	-0.136** (0.019)	-0.144** (0.043)	0.007 [0.856]
Median	-0.350** (0.020)	-0.174** (0.043)	-0.176** [0.000]	-0.263** (0.013)	-0.152** (0.022)	-0.111** [0.000]
90th percentile	-0.253** (0.038)	-0.193** (0.044)	-0.060* [0.074]	-0.223** (0.048)	-0.158** (0.034)	-0.065 [-]
Representation						
Below 10th percentile	0.289** (0.024)	0.205** (0.014)		0.151** (0.011)	0.181** (0.016)	
Below median	0.756** (0.010)	0.647** (0.021)		0.726** (0.009)	0.635** (0.025)	
Above 90th percentile	0.026** (0.005)	0.051** (0.006)		0.033** (0.010)	0.041** (0.009)	
s^2	0.390	0.288		0.349	0.288	

Notes: Standard errors are in parentheses, p-values for glass door test are in brackets. Missing p-value indicates a negative Hausman test statistic. Reference category for differentials is Canadian born, white. For differentials and glass door test, ** indicates statistically significant at 5% level and * indicates statistically significant at 10% level. For representation, ** indicates difference from population quantile is statistically significant at 5% level and * indicates this difference is statistically significant at 10% level. Columns (1) and (4) are based on separate regressions for 32,898 men and 25,400 women. Columns (2) and (5) are based on pooled regressions for men and women, with all controls except firm effects interacted with sex, averaged over 50 random samples of 1,500 firms with at least 2 employees. Mean number of observations in the 50 random samples is 17,906.