

Job Prospects and Pay Gaps:

Theory and Evidence on the Gender Gap from U.S. Cities^{*}

Chris Bidner[†] and Ben Sand[‡]

March 30, 2017

Abstract

Are differences in the quality of workers' prospects outside of their current employment relationship influential in generating pay differentials? We consider the role of an economy's industrial structure in generating differences in outside prospects, and apply our analysis to the gender pay gap in the U.S. during the 1980-2010 period. We develop a formal search and matching model that connects outside prospects, industrial structure and wage gaps and use it to guide our subsequent empirical analysis of local labor markets. Our results suggest that an economy's *within-industry* gender pay gap—which also controls for human capital characteristics—is substantially influenced by gender differences in the quality of outside prospects generated by the economy's industrial structure. Our analysis reveals that the relatively sharp narrowing of the gender pay gap during the 1980s is accounted for by the relatively sharp decline in the outside prospects of men during this period.

^{*}Previously circulated as “Industrial structure and the gender gap: Evidence from U.S. cities”. We are grateful for helpful comments received from seminar participants at the University of Adelaide, UC Berkeley, the University of New South Wales, the University of British Columbia, the University of Toronto, McMaster University, UC Merced, Simon Fraser University and from discussions with Paul Beaudry, David Card, Kelly Foley, Nicole Fortin, David Green, Kevin Lang, Thomas Lemieux, Enrico Moretti, and Krishna Pendakur. All errors are our own.

[†]Department of Economics, Simon Fraser University. cbidner@sfu.ca

[‡]Department of Economics, York University. bmsand@econ.yorku.ca

1 Introduction

Outside of perfectly competitive labor markets, a wage gap between otherwise identical workers can arise if the workers have different prospects *outside* of their current employment relationship. In this paper, we argue that this observation is important for understanding gender wage differentials. We focus on the U.S. experience during 1980-2010, a period in which substantial strides toward gender pay equity were made. This period is also one in which various changes in the U.S. economy slashed the employment prospects of male workers. Our analysis sheds light on a deeper connection between these two features, and suggests that much of the relative improvement of women’s pay resulted from the declining outside prospects of men.

A key challenge in evaluating the importance of outside prospects is that they are unobserved. We make progress on this by focusing on differences in outside prospects that are implied by an economy’s industrial structure when viewed through the lens of a model economy with search frictions. The key idea is that an economy’s array of industrial advantages will generally imply quite different outside prospects for men and women. For instance, the characteristics of an economy’s trucking sector will influence the outside prospects of workers in the retail sector, but much more so for men owing to their greater propensity to work in trucking.¹ Thus, when deregulation lowers wages in the trucking sector it also reduces the quality of outside prospects of males relative to females in the retail sector, potentially affecting the gender wage gap in retail. More generally, search frictions will generate rents which vary across industries. The part of the rent that accrues to a worker in a given industry will be shaped by the expected rent that they could obtain in other industries. If there are gender differences in the sorts of industries that workers expect to be employed in, then a given industrial structure generates gender differences in expected rent. This difference, representing a difference in the *quality* of outside prospects, then contributes to within-industry gender pay differentials.

In order for differences in outside prospects to be quantitatively important, it must be that there are non-trivial gender differences in expected rent. To get a first sense of whether this is the case, let d_j denote rent in industry j so that the female-male difference in expected rents—the ‘gender rent gap’—is $\sum_j [\pi_{fj} - \pi_{mj}] \cdot d_j$, where π_{gj} is the probability of a gender g worker being employed in industry j . The contribution of industry j to the gender rent gap therefore depends on how ‘good’ the industry is (d_j) and the ‘femaleness’ ($\pi_{fj} - \pi_{mj}$) of the industry. Using 1980 Census data, Figure 1 plots these two components for aggregated industry groups, using observed employment shares for the probabilities and estimated industry premia for industry rents.² The figure demonstrates a clear pattern: in 1980, ‘male’ jobs tended to be high rent jobs (top left quadrant) and ‘female’ jobs tended to be low rent jobs (bottom right quadrant). This relationship suggests that gender differences in the quality of outside

¹Our approach to formalizing this is to suppose that unemployed workers encounter vacancies randomly, but allow there to be some exogenous probability that the relationship produces insufficient surplus to render it ‘viable’. This probability depends on the worker’s gender and the vacancy’s industry, and as such, the genders effectively encounter a different mix of vacancies when unemployed. To focus on the *quality* of outside prospects, we hold fixed the *quantity* of outside prospects by having genders effectively encounter vacancies at the same rate when unemployed.

²These premia are the estimated coefficients on industry dummies from a log wage regression that also controls for gender, human capital characteristics (education, potential experience), race and occupation. See appendix section B for details.

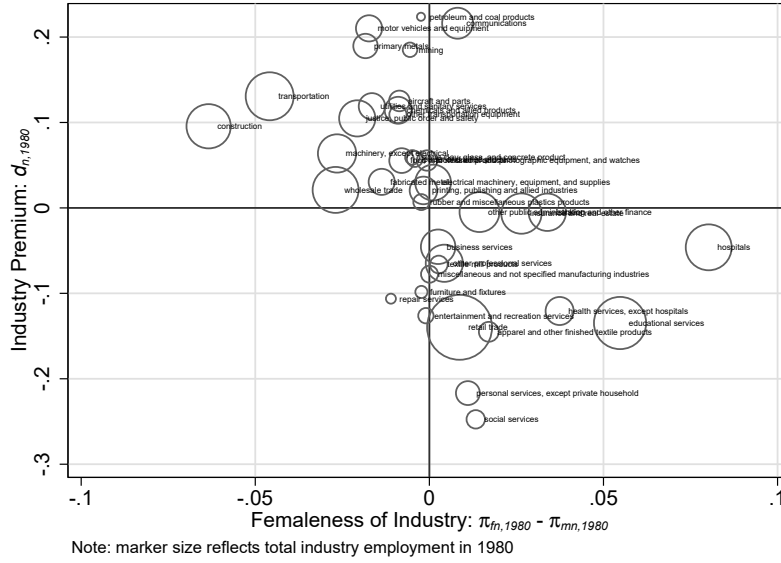


Figure 1: Male Jobs are Good Jobs, 1980

prospects do indeed have the potential to play an important role in shaping within-industry gender wage gaps.

Time series evidence offers a first pass at evaluating the relevance of gender differences in outside prospects. Using data from the Current Population Survey Outgoing Rotation Group (CPS), Figure 2 plots the time pattern of the within-industry gender pay gap in the U.S. during the 1980-2010 period. This is measured by the coefficient on a female dummy in a log wage regression that also controls for human capital variables (education and potential experience), race, occupation and industry indicators.³ The basic pattern of this gap—a sharp narrowing during the 1980s followed by relatively slow narrowing since the mid-1990s—mirrors that of the well-known pattern of the ‘raw’ gender pay gap described extensively in the literature (e.g. Blau and Kahn (2006), Blau and Kahn (1997)). This period was also one in which industrial changes, driven by forces such as globalization and deregulation, are perceived to have hit working males particularly hard. Of specific importance for our purposes, this period saw expected rents for males decline substantially relative to females. Indeed, the gender rent gap, also plotted in Figure 2, significantly narrowed during this period. More importantly, the time pattern of the gender rent gaps shares a striking resemblance to that of the within-industry gender wage gap.⁴ That is, the gender rent gap also narrows particularly dramatically in the 1980s, before stabilizing in the mid-1990s.

Our focus on the declining quality of outside employment prospects for males during this period is a central point of contrast with the various alternative explanations for the narrowing

³See appendix section B for details. The time series of the ‘raw’ gap—the coefficient on a female dummy from a log wage regression that does not control for anything else—is depicted in Figure 10 in the appendix.

⁴Another way to visualize the decline in male rents relative to females is to examine how Figure 1 changes over time. Figure 14 in the appendix does just this, using Census data from 1980, 1990, 2000, and 2010. It shows how the tendency for male jobs to be good jobs weakens over the period, especially between 1980 and 1990.

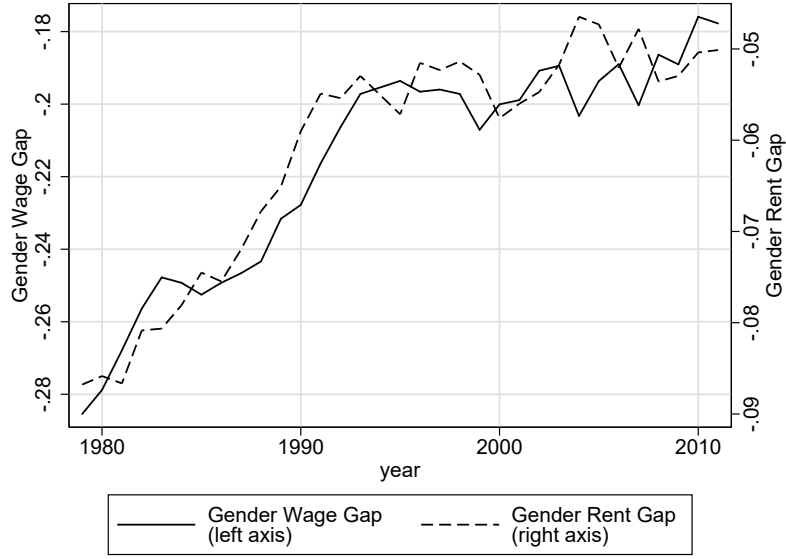


Figure 2: The Gender Pay Gap and Rent Gap Over Time

of the gender pay gap proposed in the literature. Such explanations—e.g. falling discrimination against women, an improvement in average unobserved skills among women owing to selection into the workforce, and an increase in the return to unobserved skills possessed by women because of technological progress—instead emphasize absolute gains in women’s wages.⁵ In evaluating our mechanism against these explanations, it is illustrative to examine the time pattern of real wage *levels*. Figure 3 shows the evolution of male and female real log wages, as well as the difference between the two, for two skill groups using the CPS data. In both plots it is clear that the periods in which the gender difference in wages narrowed the most (roughly, up until the mid-1990s in both cases) were also periods in which real male wages were falling. In short, the deterioration in male wages appears to be important for understanding the narrowing of the gender pay gap. Our explanation is parsimonious in the sense that the industrial changes of the 1980s can simultaneously explain a narrowing of the gender gap *and* a decline in male wage levels without having to invoke some sort of background secular decline in wages.

Whilst the time series evidence is suggestive, the various changes occurring during this period make it difficult to isolate the effect of outside prospects on the gender pay gap. Our analysis will therefore utilize the cross-city implications of industrial structure for within-industry gender pay gaps. As a first step in establishing the viability of this approach, Figure 4 uses Census data on the 100 largest U.S. cities to demonstrate that the average within-industry gender pay gap varies substantially by city.⁶ By plotting the gaps in 1980 and 2000,

⁵For instance, selection is emphasized by Mulligan and Rubinstein (2008), the role of technology is emphasized by Welch (2000), Weinberg (2000), and Beaudry and Lewis (2014). Fortin and Lemieux (2000) draw connections to changes in men’s wage structure.

⁶A city’s gender pay gap is obtained by taking the coefficient on female×city interactions from a log wage regression that also controls for human capital variables, race, occupation, and industry. See appendix section A.1 for further details.

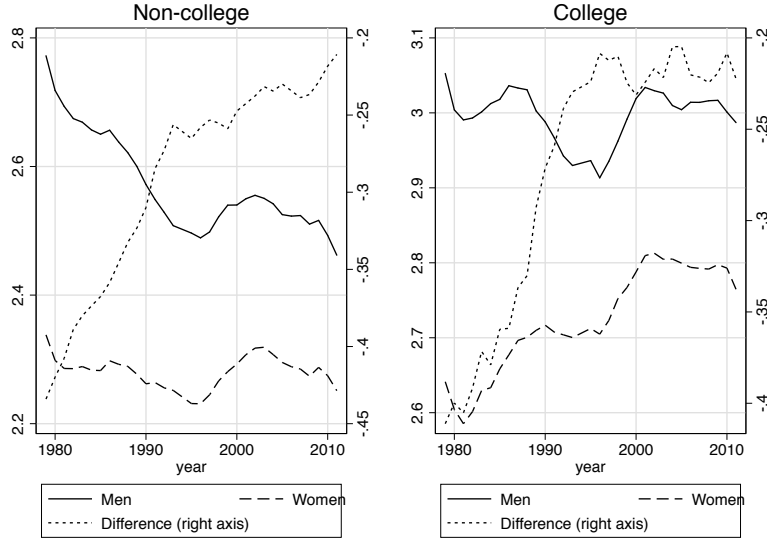


Figure 3: Gender Wage Differentials Narrow when Male Wages Fall

the figure also demonstrates that these gaps are persistent over time: cities with relatively high gaps in 1980 tend to have relatively high gaps in 2000.

In order to explore the implications of the mechanism in a more disciplined manner, we develop a formal model in section 2. The model not only formalizes the mechanism we have in mind, but also provides guidance on how to measure the quality of outside prospects facing a population of workers. The model makes precise the idea that industrial structure has two effects on wage gaps. There is a *direct* effect, arising purely from compositional implications of differences in employment distributions. For instance, men will earn more than women on average (after accounting for differences in human capital) if men are more likely to work in high-paying industries. There is also an *indirect* effect arising from the impact industrial structure has on gender differences in the quality of outside prospects as captured by differences in expected rent. For instance, men will earn more than women on average (after accounting for differences in human capital) *within the same industry* if men have higher quality outside prospects owing to their greater tendency to work in high-paying industries. The theory suggests a way to construct a rent measure that can be used to estimate the magnitude of the indirect effect relative to the direct effect. The magnitude also tells us the extent to which standard decompositions of the wage gap will underestimate the role of industrial structure.

Our empirical approach to estimating the key parameter of interest is based on a generalized triple-differences strategy. We essentially ask whether the cities that experienced a greater narrowing of the gender rent gap were also the cities that experienced a greater narrowing of their within-industry gender pay gaps. We highlight some threats to identification and propose an instrument that addresses them. Finally, we probe the robustness of our results in various ways.

Our estimates indicate that the indirect effect of industrial composition (i.e. the effect via differences in the quality of outside prospects) is at least as large as the direct effect (i.e.

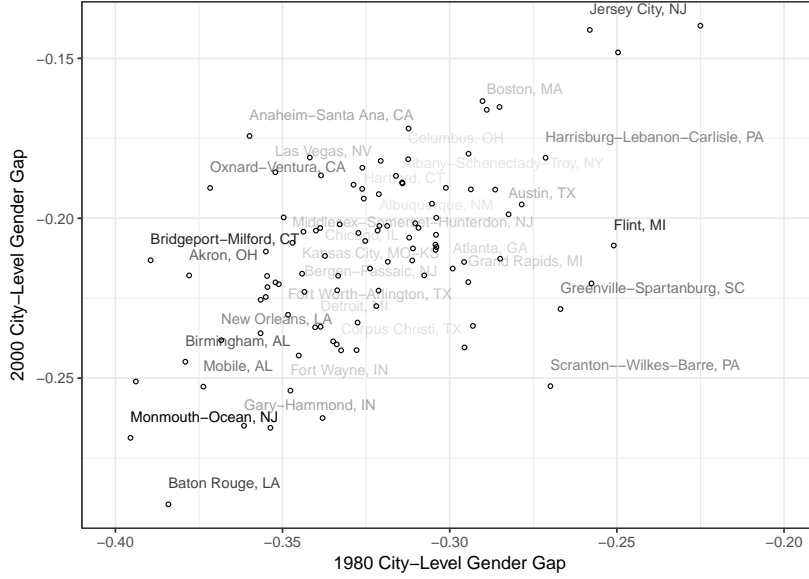


Figure 4: Variation in Gender Wage Gaps Across U.S. Cities, 1980 & 2000

the mechanical effect via composition): around 1.4 times as large. We use our estimates to revisit the CPS time series data from Figure 2 above. Differences in the quality of outside options stemming from industrial structure are able to account for 30-40% of the level of the within-industry gender pay gap. We show that the relatively steep narrowing of the pay gap during the 1980s is accounted for by the relatively steep narrowing of the rent gap during this period. We then explore *why* the narrowing of the rent gap slowed in the 1990s using standard decomposition methods. We find that rents continued to shift from traditionally male to traditionally female industries throughout the 1980-2010 period, but traditionally good jobs ceased experiencing composition shifts away from males toward females in the mid-1990s. Since the mid-1990s the narrowing of the rent gap slowed even further as an offsetting trend appeared—rents shifted away from industries that had seen compositional shifts toward females and toward those that had seen compositional shifts toward males.

1.1 Literature

The enormous literature on gender pay differentials is outlined in numerous excellent surveys, including Goldin (2014), Blau and Kahn (2016), Blau and Kahn (2000), Bertrand (2011), Altonji and Blank (1999) and Olivetti and Petrongolo (2016). Our work is most related to those studies that seek to explain the pay differential persisting *after* controlling for relevant observed characteristics. In examining this, the literature stresses factors such as the level of unobserved skill (Mulligan and Rubinstein (2008)), the return to unobserved skill (Blau and Kahn (1997)), the composition of unobserved skill (Weinberg (2000), Welch (2000), Beaudry and Lewis (2014), Cortes et al. (2016)),⁷ the role of personality traits such as risk aversion

⁷Our work is also related to Beaudry and Lewis (2014) insofar as they examine gender gaps across cities in the U.S. during the same time period. Whereas they are interested in the role played by a city's level of technology adoption (since this influences the return to 'brains' relative to 'brawn'), we are interested in the

and competitiveness, the demand for job amenities (Flabbi and Moro (2012), Goldin (2014)), and the extent of taste-based discrimination.

Whilst such work typically controls for human capital variables, we are focused on the gender pay gap that also persists within industry. However, the main point of contrast is that we are interested in highlighting and quantifying the impact of ‘external’ conditions—the quality of outside employment prospects—on wage gaps. This is of course not to say that ‘internal’ factors are unimportant—indeed, our approach accounts for such factors in a flexible manner. We control for any factor that influences gender pay differentials at the industry \times time level or at the industry \times city level.

Our particular approach to understanding how wage gaps are influenced by differences in outside prospects stemming from industrial structure rests on two key features. The first is that wages are sensitive to outside prospects because of the existence of rents that vary by industry.⁸ Models with this feature have been used to explore a variety of issues (see Green (2015)); e.g. Acemoglu (2001) studies the equilibrium mix of ‘good’ and ‘bad’ jobs. Our work is most closely related to Beaudry et al. (2012). They use such a model to explore the role of an economy’s industrial structure in shaping the wage *level*, highlighting a multiplier effect that emerges in general equilibrium. The second key feature is that there are gender differences in the composition of viable jobs that workers anticipate. This feature is attractive in that it rationalizes the well-known empirical observation that industrial employment distributions vary by gender (e.g. see Olivetti and Petrongolo (2016)).

The main contribution then lies in combining these two features to demonstrate that industrial structure, by influencing differences in outside prospects, has under-appreciated implications for wage *gaps*.⁹ In quantifying the importance of this mechanism, we follow Beaudry et al. (2012) in using variation across U.S. cities. Since we also use the gender dimension we are able to address a central concern of that paper in that we control for unobserved industry \times city \times year factors. Finally, we propose a two-step procedure in order to overcome an inherent problem in constructing the main independent variable.

The general idea that search frictions introduce sources of wage differentials among equally productive workers arises in a number of existing papers. Our particular focus on the role of industrial structure in shaping the quality of outside prospects differentiates us from work in which wage differentials stem from differences in probabilities of transitioning to unemployment (Bowlus (1997), Bowlus and Grogan (2009)), from the existence of prejudiced employers in other jobs (Flabbi (2010)), and from differences in worker productivity in other jobs (Albrecht and Vroman (2002)). Despite the different focus, our framework is general enough to incorporate versions of all of these possibilities as special cases.

city’s industrial structure. Naturally, our empirical work will be careful to mitigate the possibility that a city’s adoption of ‘female-biased’ technology acts as an important omitted variable.

⁸Taking a more ‘micro’ perspective, there is a related literature concerned with understanding the connection between wages and rents at the firm level. See Card et al. (2016) for an overview. Our perspective is more ‘macro’ in the sense that we are interested in the general equilibrium consequences of rent-sharing.

⁹It is important to emphasize that we are concerned with *within*-industry wage gaps. This distinguishes our work from other research that connects industrial structure with wage gaps. For instance Borjas and Ramey (1995) argue that the trade-induced decline of ‘concentrated’ (high rent) industries in the U.S. can help explain an expansion of the skill wage gap since low skilled workers were over-represented in such concentrated industries. Whilst their model and main estimates stress this composition effect, they note the possibility of ‘spillovers’ in wage setting and provide some preliminary evidence suggesting that such spillovers matter. We are squarely concerned with modelling and quantifying the impact of such ‘spillovers’.

As in this paper, a central element of Card et al. (2016) is the idea that search frictions render a worker’s wage sensitive to (i) the surplus available at their specific job, and to (ii) the wage that the worker could earn elsewhere—their outside option. They take a reduced-form approach by assuming that a worker’s outside options are determined by their observed productive characteristics. In contrast, our approach leverages the logical ‘general equilibrium’ implication (i) and (ii): a worker’s wage depends on the surplus available in the current match *and* on the surplus available in the other matches available to the worker. It is by unpacking such general equilibrium implications that we are able to infer a worker’s outside options from their gender and their economy’s industrial structure. Being able to quantify outside options in this way is important because it opens the possibility that women extract less of the rents available in a given job because their outside options are worse and not because they are inherently worse bargainers.¹⁰

2 Theoretical Framework

We focus on the steady state of an economy that exhibits search frictions in the spirit of Pissarides (2000). The economy consists of firms and workers, and unfolds in continuous time. All agents are infinitely lived and discount the future at a rate, r . Workers are one of two genders, $g \in \{m, f\}$. Let ρ_{gc} be the measure of labor force participants of gender g in city c . Each firm belongs to one of N industries, $n \in \{1, \dots, N\}$, and seeks to hire a single worker in order to commence production. Production takes place in one of C local labor markets, $c \in \{1, \dots, C\}$, which are cities in our empirical work. For simplicity we treat the local labor markets as separate economies, but the mechanism is robust to allowing for worker mobility (see Beaudry et al. (2012)).¹¹

Once a worker and firm meet, the pair learn the productivity of the match. Match productivity depends on the characteristics of the worker, the job, and location. For instance, workers may differ with respect to physical strength, people skills, constraints on work hours etc. Jobs may differ with respect to required physical strength, people skills, work hours, etc. (Weinberg (2000), Welch (2000)). Locations may differ with respect physical geography, regulations, culture, etc. The distribution of worker characteristics varies by gender (e.g. males are more likely to possess physical strength, females more likely to face constraints on work hours), the distribution of job characteristics varies by industry (e.g. construction jobs are more likely to require physical strength, and education jobs are more likely to offer suitable work hours), and the distribution of location characteristics varies by city (e.g. coastal cities have ports, Californian cities have abundant sunshine, and cities near Lake Superior have ready access to iron ore).

¹⁰Indeed, Card et al. (2016) find that women extract less of the surplus available within their firm—what they call the “bargaining effect”. They also find that women tend to work at lower-paying firms—what they call the “sorting effect”. Our approach explicitly connects these two effects—a “sorting” effect reasonably suggests that women have worse outside options at any given firm, thereby *generating* a “bargaining” effect (even if there are no gender differences in bargaining skills per se).

¹¹Local industrial composition will have an impact on local wages as long as mobility is not perfect. As mobility increases, the mechanism simply operates more at a national level and less at a city level. Even in the extreme case of perfect mobility, the mechanism operates purely at the national level. Our empirical strategy allows for such possibilities—with mobility national level rents will matter (in addition to city-level rents). This will arise as a gender-specific effect (within any given year) which our fixed effects strategy allows for.

To model this in the simplest way, we suppose that for each gender-industry-city the match productivity takes one of two values. With probability λ_{gn} the match is *viable*, in which case the pair can produce output that generates a surplus of P_{gnc} . Otherwise, the match is *non-viable*, in which case a sufficiently small surplus is produced that any proposed division has at least one side preferring continued search.¹² While simple, this formulation allows for many possibilities. Perhaps the most interesting is one in which there are no gender productivity differences among those employed in any industry ($P_{fnc} = P_{mnc}$) yet there are gender differences in employment distributions across industries ($\lambda_{fn} \neq \lambda_{mn}$). We will show how gender pay gaps can arise *within each industry* in such cases.¹³

Let the equilibrium measure of vacancies by industry \times city be denoted $v \equiv \{v_{nc}\}$. Let the equilibrium unemployment rates by gender-city be denoted $u \equiv \{u_{gc}\}$, so that the measure of unemployed workers in city c is $u_c \equiv u_{fc} \cdot \rho_{fc} + u_{mc} \cdot \rho_{mc}$. The total measure of encounters per unit of time between workers and jobs in this market is $M(u_c, v_{nc})$, where M is a constant returns matching function. These encounters are allocated randomly among unemployed workers and vacancies, however the fact that the viability parameter λ_{ng} varies by gender will mean that within each city there is *effectively* a gender-specific rate at which vacancies in industry n arrive. Let q_{gnc} denote this arrival rate. Similarly, let \hat{q}_{gnc} denote the gender-specific effective arrival rate of workers faced by firms in industry n in city c .¹⁴

If we let w_{gnc} be the wage paid to a worker of type g in industry n in city c , then their value function associated with being employed in industry n , E_{gnc} , satisfies the Bellman equation:

$$rE_{gnc} = w_{gnc} + s \cdot (U_{gc} - E_{gnc}), \quad (1)$$

where s is the exogenous separation rate, and where U_{gc} is the value function associated with being an unemployed worker of type g in city c . Normalizing the flow utility while unemployed to zero,¹⁵ the value of U_{gc} satisfies:

$$rU_{gc} = \sum_n q_{gnc} \cdot \max \{E_{gnc} - U_{gc}, 0\}. \quad (2)$$

The value function associated with filling a vacancy in industry n with a worker of type g in city c , J_{gnc} , satisfies:

$$rJ_{gnc} = P_{gnc} - w_{gnc} + s \cdot [V_{nc} - J_{gnc}], \quad (3)$$

where V_{nc} is the value function associated with holding a vacancy in industry n in city c . The value function V_{nc} satisfies:

$$r \cdot V_{nc} = \hat{q}_{fnc} \cdot [J_{fnc} - V_{nc}] + \hat{q}_{mnc} \cdot [J_{mnc} - V_{nc}]. \quad (4)$$

¹²Clearly, a more general model would permit a richer distribution of match productivities. The same qualitative predictions would arise out of such a specification, however at the added cost of far greater complexity.

¹³Note that ‘viability’ may reflect pure gender discrimination in terms of employment barriers (as opposed to wage levels). Other interpretations of ‘viability’ are also possible. For instance, it could arise because of gender identities associated with particular jobs (Akerlof and Kranton (2000)) or could arise from the search technology: e.g. information about job vacancies arrives via social networks (Bayer et al. (2008), Munshi (2003)), and males and females are embedded in different networks (Brashears (2008)).

¹⁴Formally, $q_{gnc} = \lambda_{gn} \cdot \frac{M(u_c, v_{nc})}{u_c} \cdot \frac{u_{gc} \cdot \rho_{gc}}{u_c}$ and $\hat{q}_{gnc} = \frac{M(u_c, v_{nc})}{v_{nc}} \cdot \frac{\rho_{gc} \cdot u_{gc}}{u_c} \cdot \lambda_{gn} = q_{gnc} \cdot \frac{u_c}{v_{nc}}$.

¹⁵This normalization serves to clarify the presentation, but allowing for a gender \times city-specific unemployment flow utility has no impact on either the mechanism or empirical approach.

Wages are determined by generalized Nash bargaining. By letting $\phi \in [0, 1]$ be the workers' relative bargaining power, the wage, w_{gnc} , is such that the value functions satisfy:

$$E_{gnc} - U_{gc} = \frac{\phi}{1 - \phi} \cdot [J_{gnc} - V_{nc}]. \quad (5)$$

So far we have derived outcomes (wages and value functions) for a given u and v . To close the model, we impose the zero profit and steady state conditions. It costs k_{nc} units of capital to open a vacancy in industry n when located in city c . Free entry of firms will ensure that

$$V_{nc} = k_{nc} \quad (6)$$

in equilibrium. That is, for any u the value of v will adjust so that this is satisfied. In a steady state, the measure of workers leaving the unemployed pool to work in industry n equals the measure entering the unemployed pool from industry n . If we let π_{gnc} be the steady state proportion of gender g workers within city c employed in industry n , then the steady state condition is:

$$q_{gnc} \cdot u_{gc} = s \cdot (1 - u_{gc}) \cdot \pi_{gnc}. \quad (7)$$

Summing over n gives the steady-state relationship between a gender's aggregate job arrival rate in a city and their unemployment rate:

$$\sum_n q_{gnc} = s \cdot \frac{1 - u_{gc}}{u_{gc}}. \quad (8)$$

That is, for any v the value of u will adjust so that this is satisfied. From the steady state condition, equations (7) and (8), we have that the steady-state employment shares are given by:

$$\pi_{gnc} = \frac{q_{gnc}}{\sum_j q_{gjc}}. \quad (9)$$

Motivated by the associated empirical observation, we focus on equilibria in which there is positive employment of both genders in each industry.¹⁶ This serves only to simplify the presentation and derivations—the mechanism we stress in no way depends on this. In such a case workers and firms establish employment relationships as soon as they encounter a viable match.

The key outcome from the model is that equilibrium wages of a gender \times city in an industry, w_{gnc} , are tied to the average wage of the gender \times city. Specifically, when the *quantity* of alternative employment opportunities is held fixed across gender \times cities, equilibrium wages satisfy:

$$w_{gnc} = \psi_{gnc} + \delta \cdot \bar{w}_{gc}, \quad (10)$$

where $\bar{w}_{gc} \equiv \sum_j \pi_{gjc} w_{gjc}$ is the average wage of the gender \times city, and $\delta \in (0, 1)$ depends on parameters and the quantity of alternative employment opportunities.¹⁷ This wage relationship resembles an ‘endogenous’ social effect as described by the social effects literature

¹⁶That is, equilibria in which $E_{gnc} + J_{gnc} \geq U_{gc} + V_{nc}$ for all (gnc) .

¹⁷Specifically, the quantity of alternative employment opportunities can be captured by the unemployment rate and, holding the unemployment rate at \bar{u} , we have $\delta = (1 - \phi) \cdot [1 - \bar{u}] / [1 + \bar{u} \cdot (r/s)]$. See appendix section C.2 for further details.

(Manski (1993)). As emphasized in that literature, this relationship produces a ‘social multiplier’ whereby a change in the ψ_{gnc} terms first generates a direct mechanical effect on average wages. But then an indirect effect arises since the resulting rise in average wages raises wages further since $\delta > 0$, which raises wages further, and so on. The total impact on wages can be gleaned by examining the reduced form version of this relationship. This tells us that the wages of a gender \times city will depend on a job-specific component *and* on the average job-specific components they experience in other jobs:

$$w_{gnc} = \psi_{gnc} + \gamma \cdot R_{gc}^*, \quad (11)$$

where $\gamma = \delta/(1 - \delta)$ and

$$R_{gc}^* \equiv \sum_j \pi_{gjc} \cdot \psi_{gjc} \quad (12)$$

is an index of the *quality* of alternative employment opportunities available to the gender \times city in equilibrium.

Since a central goal of our empirical work is to produce a credible estimate of γ , we pause to explain its interpretation before moving on to the empirical implementation.

2.1 Interpreting γ

A key insight from model is that small changes in conditions can, via a *multiplier* effect, induce a large change in the wage gap.¹⁸ For instance, imagine all female workers experienced an increase in their job-specific component, ψ_{fnc} , by x . This could perhaps be due to rising female productivity or declining prejudice. There is a *direct* effect on the wage gap—it narrows by x . But the fact that all female workers benefit from the direct effect means that women are now able to obtain higher wages outside of their current job. This relative improvement in outside prospects raises female wages further, which induces a further relative improvement in outside prospects, and so on. This spillover effect represents the *indirect* effect on the within-industry gender wage gap. In short, the magnitude of γ tells us the magnitude of the indirect effect relative to the direct effect. In terms of the example above, the indirect effect will contribute a further $\gamma \cdot x$ to the narrowing of the pay gap so that the total effect of the change is $(1 + \gamma) \cdot x$.

The same logic extends to the impact of forces ‘external’ to the worker’s employment relationship (i.e. those that leave each ψ_{gnc} unaffected). Specifically, if an economy’s industrial structure changes then there will be a direct effect on the wage gap reflecting the mechanical re-allocation of men and women across better or worse paying jobs. Again this will be accompanied by an indirect effect as the direct effect initiates the cycle of narrowing gaps in wages and outside prospects. Importantly, the direct effect will be a between-industry phenomenon whereas the indirect effect will be a within-industry phenomenon (as all outside prospects are equally valuable, regardless of current industry). Therefore a standard Blinder-Oaxaca decomposition will *underestimate* the total impact of industrial structure on wage gaps since it will only capture the direct effect. The magnitude of γ thus also tells us the extent of this underestimation—the total effect is $(1 + \gamma)$ times as large as that implied by the decomposition.

¹⁸Whilst not focusing on wage gaps, this key insight and the implications we discuss here are central in Beaudry et al. (2012) and are lucidly discussed in Green (2015).

This discussion suggests a counterfactual gender wage gap that can be used to further evaluate the quantitative relevance of the mechanism. This is the within-industry pay gap that would persist if industrial structure did not generate any gender differences in outside prospects, e.g. if the gender differences in the model's viability parameters, λ_{gn} , were eliminated.

We now show these points more formally. To simplify the exposition, consider a particular economy so that, from (11), the wage for gender g workers in industry n can be written, dropping the city subscript, as:

$$w_{gn} = \psi_{gn} + \gamma \cdot R_g^*, \quad (13)$$

where $R_g^* \equiv \sum_j \pi_{gj} \cdot \psi_{gj}$. The average wage for gender g is therefore

$$\bar{w}_g \equiv \sum_j \pi_{gj} \cdot w_{gj} = \sum_j \pi_{gj} \cdot \psi_{gj} + \gamma \cdot R_g^*. \quad (14)$$

We can compute a gender wage gap by taking the gender difference, $\bar{w}^{\text{gap}} \equiv \bar{w}_f - \bar{w}_m$. A standard Blinder-Oaxaca decomposition (Fortin et al. (2011)) can be used to separate the wage gap into a part that is explained by observed differences in industrial employment, E , and a part that remains unexplained, U :

$$\bar{w}^{\text{gap}} \equiv \bar{w}_f - \bar{w}_m \quad (15)$$

$$= \underbrace{\gamma \cdot R^{*\text{gap}} + \sum_j \bar{\pi}_j \cdot [\psi_{fj} - \psi_{mj}]}_U + \underbrace{\sum_j [\pi_{fj} - \pi_{mj}] \cdot \bar{\psi}_j}_E, \quad (16)$$

where $R^{*\text{gap}} \equiv R_f^* - R_m^*$, and where $\bar{\pi}_j$ and $\bar{\psi}_j$ are gender-neutral industry weights and premia.¹⁹ That is, the gender pay differences arise both because the distribution of workers across industries varies by gender (E) and because pay within industries varies by gender (U). The value of U is the gender wage gap depicted in Figure 2. The value of E is the *direct* effect of industrial structure on the wage gap. The *indirect* effect arises because industrial structure also affects $R^{*\text{gap}}$. To see this, note that we can also decompose gender rent differences into a part explained by industrial employment and to an unexplained part:

$$R^{*\text{gap}} = \underbrace{\sum_j \bar{\pi}_j \cdot [\psi_{fj} - \psi_{mj}]}_{\bar{U}} + \underbrace{\sum_j [\pi_{fj} - \pi_{mj}] \cdot \bar{\psi}_j}_E. \quad (17)$$

That is, gender rent differences arise both because the distribution of workers across industries varies by gender (E) and because the job-specific component within industries varies by gender (\bar{U}). That is, \bar{U} captures the direct effect of all factors 'internal' to the employment relationship. Notice that E is the rent gap measure depicted in Figure 2. Notice too that this E is the same as that appearing in the wage gap decomposition. Thus, from (16) and (17), we have

$$U = \bar{U} + \gamma \cdot [E + \bar{U}]. \quad (18)$$

¹⁹Specifically, let Γ be an arbitrary $N \times N$ diagonal weighting matrix. Then, in vector notation, if $\bar{\pi} \equiv (I - \Gamma)\pi_f + \Gamma\pi_m$ then $\bar{\psi} \equiv \Gamma\psi_f + (I - \Gamma) \cdot \psi_m$.

Thus the adjusted wage gap arises because of the direct effect of gender differences in factors ‘internal’ to the employment relationship (\bar{U}) plus the indirect effect of differences in such factors ($\gamma \cdot \bar{U}$) plus the indirect effect of industrial structure ($\gamma \cdot E$).²⁰ This suggests a counterfactual wage gap of interest: the adjusted wage gap that would arise had there been no differences in outside prospects generated by industrial structure:

$$U^* \equiv U - \gamma \cdot E = (1 + \gamma) \cdot \bar{U}. \quad (19)$$

Once a value of γ is estimated, this counterfactual gap can be computed from U and E . As pointed out above, these series have already been estimated using the CPS data and displayed in Figure 2. We report on estimates of U^* in Section 5.

3 Empirical Implementation

3.1 Data

The data for our main empirical analysis comes from the U.S. Census Public Use Micro-Samples (PUMS) for the years 1980, 1990 and 2000. We also use data from the American Community Surveys for the years 2009-2010-2011, that we combine into a single year and refer to as 2010. Our main empirical work excludes the 2010 sample, although little changes when 2010 is included (see section 4.2.3 below).²¹ All analysis is restricted to individuals between the ages of 22 and 54 with at least 1 year of potential labor market experience. We measure hourly wages by dividing annual wage and salary income by the product of weeks worked and usual hours worked per week. In our baseline empirical work, we focus on full-time, full-year workers. It is hoped that this restriction mitigates composition issues that can arise when comparing workers over long periods, particularly for the sample of female workers. However, we examine robustness to this restriction in section 4.2. Our analysis also requires consistent coding of metropolitan areas, years of education and completed schooling, and industrial classification across the Census years. Further details can be found in Appendix A.

3.2 Preliminaries

In order to take the essence of the model to the data we need to incorporate heterogeneous workers. We do so by interpreting w_{gnc} as the wage per effective labor unit and assuming that workers with characteristics X_i have $\exp(X_i' \beta_g + \epsilon_i)$ effective labor units where ϵ_i captures

²⁰The fact that industrial employment distributions simultaneously affect both the explained and unexplained component of gender wage differences (via the direct and indirect effects, respectively) is consistent with the finding that both industrial employment differences and the unexplained component are important in accounting for gender wage gaps (e.g. Blau and Kahn (2016) for the U.S. and Schirle (2015) for Canada).

²¹We exclude the 2010 sample from the main results for four reasons. First, since 2008, the ACS does not contain a continuous measure of weeks worked. Instead, the ACS reports weeks worked in intervals, potentially making our measure of hourly wages less reliable and comparable between genders. Second, we wish to avoid complications that could arise from using labor market outcomes from the Great Recession and its immediate aftermath, as we interpret the model as describing a steady state equilibrium. Third, the fact that the sample spans three years reduces its comparability with the other samples. Finally, despite merging three years the resulting sample size is considerably smaller than the other three years (and especially so after dropping small cells) which may be particularly problematic since it alters the set of fixed effects that are identified.

unobserved individual heterogeneity. A worker's observed log wages are then given by:

$$\ln W_{i,gnc} = X_i' \beta_g + \ln w_{gnc} + \epsilon_i. \quad (20)$$

The values of $\ln w_{gnc}$ are our object of interest. To obtain a measure of these, we estimate (20) capturing $\ln w_{gnc}$ with gender \times industry \times city fixed effects, eliminating industry \times city cells with fewer than 20 observations for both men and women. Our specification for X_i includes controls for a quartic in potential experience; hispanic, black, and immigration dummies; an indicator for whether an individual is observed in a city located in one's birth state—all interacted with education (four categories)—and four occupation dummies. All covariates have coefficients that vary by gender. This procedure is repeated separately for each census year.

To derive the appropriate estimating equation, we turn to the model and take a log approximation of (11) to get:

$$\ln w_{gnc} = \tilde{\psi}_{gnc} + \gamma \cdot \tilde{R}_{gc}^* \quad (21)$$

where $\tilde{R}_{gc}^* \equiv \sum_j \pi_{gjc} \cdot \tilde{\psi}_{gjc}$ and the $\tilde{\psi}_{gnc}$ terms are completely analogous to the ψ_{gnc} terms except that they are expressed as relative to an omitted group because of the log approximation. We cannot hope to use \tilde{R}_{gc}^* as an independent variable since we cannot construct it without knowing $\{\tilde{\psi}_{gjc}\}_j$, and these terms cannot be estimated from (21). To make progress on this, we allow $\tilde{\psi}_{gnc}$ to consist of a gender \times industry effect, and industry \times city effect, and a gender \times city effect:

$$\tilde{\psi}_{gnc} = \iota_{gn} + \zeta_{nc} + \xi_{gc}. \quad (22)$$

Using this, (21) can be written:

$$\ln w_{gnc} = \iota_{gn} + \zeta_{nc} + \gamma \cdot R_{gc} + \nu_{gc} + \varepsilon_{gnc}, \quad (23)$$

where $\nu_{gc} \equiv (1 + \gamma) \cdot \xi_{gc}$, ε_{gnc} is the gender \times industry \times city error component of ϵ_i , and where

$$R_{gc} \equiv \sum_j \pi_{gjc} \cdot [\iota_{gj} + \zeta_{jc}] \quad (24)$$

is a measure of rent. This measure only uses the components of job-specific rent, $\tilde{\psi}_{gnc}$, that are attributable to gender \times industry-specific factors and to industry \times city-specific factors. This is important since, as will become clear, we are able to estimate these components. Finally, to account for differences in the quantity of outside employment opportunities we control for the log unemployment rate of the gender \times city, u_{gc} .²² Our main estimating equation then is based on:

$$\ln w_{gnc} = \iota_{gn} + \zeta_{nc} + \gamma \cdot R_{gc} + \alpha \cdot u_{gc} + \nu_{gc} + \varepsilon_{gnc}. \quad (25)$$

Before this relationship can be estimated, we first need to construct R_{gc} . To this end, the next step our empirical implementation is to collect all the gender \times city factors into a fixed effect, $\chi_{gc} \equiv \gamma \cdot R_{gc} + \alpha \cdot u_{gc} + \nu_{gc}$, and estimate

$$\ln w_{gnc} = \iota_{gn} + \zeta_{nc} + \chi_{gc} + \varepsilon_{gnc}. \quad (26)$$

²²See appendix section C.2 for details of why the unemployment rate is an appropriate measure of a gender \times city's quantity of alternative employment opportunities, for details of the linear approximation that justifies (25), and for details of the log approximation.

We use the resulting estimates of $\{\iota_{gj}, \zeta_{jc}\}_j$, along with observed employment shares, to construct R_{gc} according to (24).²³ We repeat this procedure for each census year.

3.3 Identification

Once we have our rent measure, R_{gc} , we return to (25). In principal, this would allow us to estimate γ using a single cross-section based on:

$$\ln w_{gnc} = D_{gn} + D_{nc} + \gamma \cdot R_{gc} + \alpha \cdot u_{gc} + \nu_{gnc}, \quad (27)$$

where $\nu_{gnc} \equiv \nu_{gc} + \varepsilon_{gnc}$ and the D terms are fixed effects. Given the fixed effects, this specification would use cross-city variation to estimate the relationship between gender rent gaps and within-industry gender wage gaps. As such, the concern with estimating this with OLS is the possibility that unobserved factors affecting within-industry wage gaps in a city are correlated with a city's measured gender rent gap. This could arise for a variety of reasons: (i) cities with a particularly prejudiced attitude toward women pay women relatively poorly within all industries *and* impose upon women barriers to employment in high-paying industries; (ii) cities where women have a particularly high unobserved ability/ambition/competitiveness are those where women are paid relatively highly within all industries *and* have a greater propensity to secure employment in high-paying industries; (iii) cities where women have a particularly high demand for job amenities such as temporal flexibility are those where women are paid relatively little in each industry (because of compensating differential arguments) *and* are particularly concentrated in industries with greater scope to offer such amenities; and (iv) cities with a high level of technology adoption pay women relatively highly within each industry (thanks to womens' larger endowment of 'brains') *and* systematically induce particular industrial structures.

The general issue is that the error term contains gender \times city characteristics that may be correlated with measured rent. Rather than trying to address these in a cross-section setting, our approach begins with using the time dimension in order to control for all (time-invariant) unobserved gender \times city characteristics. Adding a time subscript to each term in (27) and decomposing $\nu_{gnc} = D_{gnc} + \epsilon_{gnc}$, where ϵ_{gnc} is the mean zero time shock to a gender \times industry \times city cell, we get:

$$\ln w_{gnc} = D_{gnt} + D_{nct} + D_{gnc} + \gamma \cdot R_{gct} + \alpha \cdot u_{gct} + \epsilon_{gnc}. \quad (28)$$

The remaining threat is the possibility that the gender \times city \times time component of the error term is correlated with measured rent.²⁴ To clarify this further, we note that given the fixed effects we are adopting a generalized triple-difference strategy. That is, we are estimating γ by making across-city comparisons of changes in within-industry gender pay gaps with changes in

²³Since R_{gc} and $\ln w_{gnc}$ both depend on ι_{gn} and ζ_{nc} , one may be concerned about a mechanical relationship between the two. However, as will become clear, we completely account for this in our regressions by including fixed effects. In other words, there is clearly no mechanical relationship between R_{gc} and $\ln w_{gnc}$ *conditional* on ι_{gn} and ζ_{nc} .

²⁴That is, decompose $\epsilon_{gnc} = \tau_{gct} + \tilde{\epsilon}_{gnc}$ where $\tilde{\epsilon}_{gnc}$ is mean zero within each gender \times city \times time cell. The endogeneity threat arises because of the possibility that τ_{gct} is correlated with R_{gct} . As will become clear, the concern more specifically is that the time change in the gender difference in τ_{gct} is correlated with the time change in the gender difference in R_{gct} .

gender rent gaps. In other words, by letting gender differences be denoted $X_k^{\text{gap}} \equiv X_{fk} - X_{mk}$ and time differences be denoted by $\Delta X_{kt} \equiv X_{k,t} - X_{k,t-1}$, we have:

$$\Delta \ln w_{nct}^{\text{gap}} = \Delta D_{nt}^{\text{gap}} + \gamma \cdot \Delta R_{ct}^{\text{gap}} + \alpha \cdot \Delta u_{ct}^{\text{gap}} + \Delta \epsilon_{nct}^{\text{gap}}. \quad (29)$$

Since $\Delta R_{ct}^{\text{gap}}$ only varies at the city level (within a given year), any remaining endogeneity concerns arise because of the possibility that changes in a city's average within-industry gender wage gap are driven by omitted factors that are correlated with changes in the city's measured gender rent gap.

There are two reasons why omitted city-specific factors may be correlated with changes in a city's measured gender rent gap. First, omitted factors may be correlated with changes in gender employment share differences. This may arise because of factors within or outside the model. Within the model, suppose that females in some city experience a positive productivity shock. This causes an increase in the wages of females relative to males within each industry in the city, but it will also influence the relative profitability of industries and thus their equilibrium employment shares. This will in general affect the gender difference in employment distributions and thus measured gender rent differentials. Outside the model, suppose that a city experiences a particularly large increase in female labor supply—perhaps because of participation or migration—which could affect the measured wage gap via selection arguments. The additional women may systematically enter particular industries and thereby affect measured female rent because of changes to observed female employment shares.²⁵

Second, omitted factors may be correlated with changes in the city component of industry premia, $\{\zeta_{nct}\}_n$. For instance, the particularly large increase in female labor supply discussed above may be a response to changes in a city's industrial advantages.

To address these issues we propose an instrument that completely discards variation due to changes in employment shares and to changes in the city component of industry premia. Specifically, we fix employment shares at a base-year level and apply a vector of industry premia that only varies across time. Using 1980 as the base year, our instrument is:

$$Z_{gct} \equiv \sum_n \pi_{gnc,1980} \cdot d_{nt}, \quad (30)$$

where d_{nt} is a time-varying gender-neutral national industry premium.²⁶ Since

$$\Delta Z_{ct}^{\text{gap}} = \sum_n \pi_{nc,1980}^{\text{gap}} \cdot \Delta d_{nt}, \quad (31)$$

the instrument isolates variation in $\Delta R_{ct}^{\text{gap}}$ stemming from over-time changes in national gender-neutral industry premia applied to gender differences in base-year employment distributions.²⁷

²⁵This would downward bias OLS estimates if the additional women were of higher average unobserved ability and tended to enter 'female' industries (since these tend to be low-paying industries).

²⁶To obtain these premia, we simply estimate $\kappa_{gnc} = d_{ct} + d_{gt} + d_{nt} + \epsilon_{gnc}$ where $\kappa_{gnc} \equiv \iota_{gnt} + \zeta_{nct}$, the d terms are fixed effects, and ϵ_{gnc} is the error term. The resulting estimates of the d_{nt} terms are the national gender-neutral industry premia for year t that we use. Discarding the city-component of industry premia also helps address the particularly large bias arising from measurement error in panel models. It is primarily for this reason that we also discard the gender-component of the industry premia.

²⁷We choose to use 1980 as the base year for practical reasons. In principle, we could have chosen an earlier Census, such as the 1970 Census for which SMSA definitions are available. However, the smaller sample sizes in the 1970 Census made computing gender \times city \times industry employment shares for full-time, full-year workers difficult, as it resulted in many missing cells.

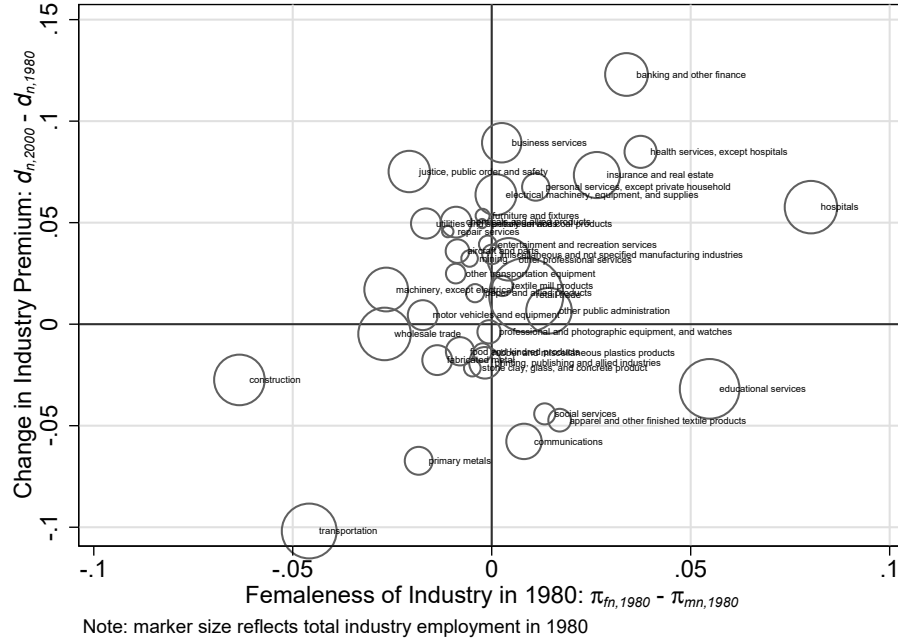


Figure 5: Instrument: Industry-level variation

To get an overview of which industries experienced premia changes and which were particularly ‘female’ in 1980 (as measured by $\pi_{nc,1980}^{\text{gap}}$) we depict these variables at the national level in Figure 5. We see that, broadly speaking, relative rents shifted from ‘male’ industries toward ‘female’ industries during the 1980-2000 period.

In order to satisfy the exclusion restriction we need to ensure that omitted city characteristics that affect the changes in a city’s within-industry gender wage gap are not correlated with $\Delta Z_{ct}^{\text{gap}}$ —i.e. not systematically related to the city’s industrial structure in 1980. To address this we control for base-year values of city-level characteristics. Specifically, we add these controls to the specification given by (28), allowing for the effect to vary with gender \times year.

The main omitted variable that we are concerned about is the extent of technology adoption. Beaudry and Lewis (2014) argue that cities that adopted technology (e.g. the PC) relatively extensively during the 1980s and 1990s experienced a narrowing of their adjusted gender pay gap (as such technologies are argued to benefit brains over brawn; Weinberg (2000)). They show how PC adoption was more extensive in cities that were more educated in 1980. Controlling for city-level aggregate education is important since a city’s industrial composition (and thus, at least potentially, our instrument) varies systematically with its aggregate education. Failing to account for this will bias our estimates downward to the extent that highly-educated cities saw a relatively small (instrumented) narrowing of the gender rent gap. We follow Beaudry and Lewis (2014) in measuring city education as the log ratio of college to high school equivalents in 1980.²⁸

²⁸ Beaudry and Lewis (2014) follow Card (2009)—those with less than high school are treated as contributing 70% of a high school worker, and those with some post-secondary are treated as contributing 60% of a high school worker and 40% of a college worker. Our results are not sensitive to this particular measure—alternatives

Changes in a city’s gender wage gap could be driven by changes in discriminatory attitudes toward women. This would require that women’s progress over time occur unevenly across cities. For instance a decline in discriminatory attitudes may emerge in larger cities before spreading to smaller ones. Thus, smaller cities in 1980 may be expected to experience a more dramatic narrowing of their gender pay gaps over time as they ‘catch up’ with larger cities. This is problematic to the extent that a city’s industrial structure systematically varies with city size. To address this sort of confounding effect, we control for city size (aggregate employment) in 1980.

Finally, it could be that changes in a city’s gender wage gap are driven by changes in selection. For instance, trends in women’s labor force participation may unfold unevenly across cities. Since there is more scope for changes in female participation in places that started with low female participation rates, it could be that the effect of selection on gender wage gaps is most pronounced in such places. This is problematic to the extent that a city’s industrial structure varies systematically with its female labor force participation rate. To address this possible confounding effect we control for the city’s female participation rate in 1980.²⁹

4 Results

4.1 Main Results

Table 2 presents the results from estimating (28) using data from 1980 to 2000. The dependent variable in all specifications is the regression adjusted wages obtained from (20). The first two columns of Table 2 contain OLS estimates. These specifications, and all those that follow, include D_{gnt} , D_{nct} , and D_{gnc} fixed effects. Thus, identification of γ comes from a generalized triple-difference strategy which compares across city changes in within-industry gender pay gaps with changes in gender differences in R_{gct} . The OLS estimates of γ in columns 1 and 2 are small and statistically insignificant. However, as discussed in section 3, the wage equation derived from the model will produce valid estimates of the causal effect of rents on wages only if time-varying, city-specific factors are uncorrelated to changes in city rent-gaps.

Columns 3-6 of Table 2 report 2SLS estimates of γ . These estimates exploit variation in R_{gct} predicted by the instrumental variable Z_{gct} presented in section 3. Recall that this instrument isolates within-city movements in R_{gct} by using *national*-level movements in gender-neutral industrial premia, weighted by base-year employment shares. The instrument is a strong predictor of the evolution of rent gaps within-cities, as shown in the bottom panel of Table 2, which contains the first-stage results. In particular, in columns 3-6, the first-stage coefficient on Z_{gct} is highly statistically significant, with an F -statistics all above 25.

The point estimates of γ in columns 3-6 vary between 0.73 - 1.54. In column 6, which includes all of our control variables, the point estimate is 1.39 and is highly statistically significant.³⁰ The magnitude of $\hat{\gamma}$ has an intuitive interpretation. First, a $\hat{\gamma} > 0$ reflects

such as BA share give very similar results.

²⁹We return to selection issues in our robustness checks in section 4.2.

³⁰Note that in this table, and all those that follow, we compute standard errors that are heteroskedasticity-robust and clustered at the city-level. Given that the level of variation in R_{gct} is at the gender \times city \times year level, we have also explored clustering by gender \times city but this made little difference in practice. Our chosen level of clustering could be considered conservative, and allows for correlated errors between genders, within a

indirect or general equilibrium forces implied by the model—within-industry gender pay gaps fall more in cities where women’s employment in higher-paying industries improves relative to men’s. Specifically, the estimate from column 6 in Table 2 indicates that the indirect effect of industrial composition is more than one-and-a-third times as large as the direct effect. Second, $\hat{\gamma}$ can be linked to the standard Blinder-Oaxaca decomposition, as discussed in section 2.1. For instance, our estimate from column 6 in Table 2 suggests that differences in industrial employment is more than two-and-a-third times ($1+\gamma=2.39$) as important for explaining the gender gap as the standard decomposition would imply. Finally, our results can be linked to the ‘social multiplier’ described by the social effects literature, as discussed in section 2 (Manski, 1993). In particular, since $\gamma = \delta/(1 - \delta)$, our estimates suggest that $\delta \approx 0.58$.³¹ This effect highlights the strategic complementarity of within-gender, across industry wages.

4.2 Other Explanations and Robustness

4.2.1 Excluding College Workers

Our motivating observation in section 1 that male real wages suffered absolute declines since 1980 applies less strongly to those with college degrees (see Figure 3). Alternative arguments for narrowing gender gap, such as those based on technological change, presumably apply more readily to college educated workers. For these reasons we find it instructive to examine the robustness of our results to excluding college educated workers.

Table 5 estimates equation (28) using only data on non-college workers over the 1980-2000 period. The layout of the table is exactly the same as Table 2. The 2SLS results in columns 3-6 vary between 1.82 and 1.98, depending on the set of controls used. The bottom panel of the table shows the first-stage results using Z_{gct} as an instrument, which again indicate a strong first-stage relationship. The results from this table suggest that the mechanism we emphasise is potentially even more important for less educated workers.

4.2.2 Rents, Profitability and Discrimination

A body of research posits that firms are better able to indulge in (costly) taste-based discrimination when they are in a high-rent industry; e.g. see Black and Brainerd (2004), Black and Strahan (2001), and Ashenfelter and Hannan (1986). Our results cannot be confounded by the basic version of this story since the level of prejudice prevailing in an industry—whether due to available rents or any other reason—is captured by the gender \times industry \times year fixed effects. In other words, adjusted wage gaps at the city level are by definition not affected by the extent of industry-specific discrimination. However, our results are possibly explained by a slightly extended version whereby firms *can* discriminate in a high-rent industry and *will* discriminate (with a higher probability) when they are surrounded by other discriminatory firms.³² Thus, cities with many high-rent industries will have more discrimination within each high-rent industry and thus will experience a wider adjusted gender wage gap as a result.

city, as well as over time.

³¹In fact, we could estimate (10) directly. The structure of (10) implies an inherent endogeneity problem because of \bar{w}_g , but equation (14) suggests that Z_{gct} can be used as an instrument for \bar{w}_{gc} to estimate δ directly. We report the results from this exercise in Appendix E.3, Table E1.

³²The latter aspect could arise for a variety of reasons. For instance, because workers that are discriminated against have less scope to leave a discriminatory firm for a non-discriminatory one, or perhaps because of greater difficulty in containing the cultural transmission of discriminatory attitudes.

Table 1: Estimates: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{gct}	-0.062 (0.13)	0.12 (0.15)	0.73** (0.36)	1.54** (0.45)	1.42** (0.39)	1.39** (0.40)
u_{gct}	-0.013 (0.011)	-0.0053 (0.011)	-0.012 (0.013)	-0.014 (0.015)	-0.0086 (0.015)	-0.0075 (0.015)
Observations	16476	16476	16476	16476	16476	16476
R^2	0.990	0.990	0.990	0.990	0.990	0.990
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls						
Education		Yes		Yes	Yes	Yes
Size		Yes			Yes	Yes
Participation		Yes				Yes
First-Stage:						
Z_{gct}			2.18** (0.37)	1.73** (0.34)	1.78** (0.34)	1.75** (0.33)
F -Stat.			34.96	26.35	26.68	28.50
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

Table 2: Estimates: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{gct}	0.077 (0.24)	0.40* (0.24)	1.30** (0.55)	2.19** (0.53)	1.99** (0.46)	1.90** (0.47)
u_{gct}	-0.013 (0.011)	-0.0041 (0.011)	-0.0085 (0.012)	-0.0072 (0.013)	-0.0026 (0.014)	-0.00056 (0.013)
Observations	16476	16476	16476	16476	16476	16476
R^2	0.990	0.990	0.990	0.990	0.990	0.990
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls						
Education		Yes		Yes	Yes	Yes
Size		Yes			Yes	Yes
Participation		Yes				Yes
First-Stage:						
Z_{gct}			1.41** (0.18)	1.21** (0.17)	1.27** (0.17)	1.28** (0.17)
F -Stat.			59.50	52.45	55.66	55.35
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

Table 3: Basic Results, base Workers: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{cg}	-0.071 (0.13)	0.0045 (0.12)	0.99** (0.41)	1.27** (0.38)	0.00012 (0.39)	0.49 (0.42)
Non-Col. $\times R_{cge}$	0.21 (0.18)	0.13 (0.18)			1.51** (0.55)	1.25** (0.59)
u_{gc}	-0.0028 (0.011)	0.0058 (0.011)	0.0057 (0.012)	0.014 (0.013)	0.0058 (0.012)	0.013 (0.014)
Observations	63874	63874	63874	63874	63874	63874
R^2	0.983	0.983	0.983	0.983	0.983	0.983
Partial R^2			.04	.077	.	.
Fixed Effects:						
City \times Ind. \times Educ. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Educ. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Educ. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls						
Education		Yes		Yes		Yes
Size		Yes		Yes		Yes
Participation		Yes		Yes		Yes
First-Stage:						
			1.35** (0.22)	1.29** (0.23)		
F -Stat.			39.56	31.07	13.71	16.75
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (*) and (**) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The Partial R^2 shows the fraction of the cross-city variation of the gender gap that is explained by R_{gc} after taking into account all other regressors. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

Table 4: Estimates for Non-College Workers: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{gct}	0.21 (0.16)	0.27* (0.16)	1.82** (0.61)	1.95** (0.58)	1.89** (0.56)	1.98** (0.57)
u_{gct}	0.0051 (0.013)	0.0039 (0.013)	0.010 (0.018)	0.0020 (0.017)	0.0044 (0.017)	0.0018 (0.019)
Observations	15264	15264	15264	15264	15264	15264
R^2	0.988	0.988	0.987	0.987	0.987	0.987
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls						
Education		Yes		Yes	Yes	Yes
Size		Yes			Yes	Yes
Participation		Yes				Yes
First-Stage:						
Z_{gct}			1.53** (0.29)	1.48** (0.29)	1.47** (0.31)	1.44** (0.30)
F -Stat.			26.83	25.50	22.51	22.85
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

Table 5: Estimates for Non-College Workers: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{gct}	0.21 (0.16)	0.27* (0.16)	1.99** (0.71)	2.19** (0.67)	2.16** (0.67)	2.34** (0.70)
u_{gct}	0.0051 (0.013)	0.0039 (0.013)	0.011 (0.019)	0.0022 (0.018)	0.0044 (0.019)	0.0014 (0.021)
Observations	15264	15264	15264	15264	15264	15264
R^2	0.988	0.988	0.987	0.987	0.987	0.987
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls						
Education		Yes		Yes	Yes	Yes
Size		Yes			Yes	Yes
Participation		Yes				Yes
First-Stage:						
Z_{gct}			1.31** (0.31)	1.25** (0.31)	1.23** (0.32)	1.20** (0.31)
F -Stat.			17.88	16.29	14.77	14.56
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

To provide evidence against this sort of channel, we explore a key implication of the extended version of the above story. Namely, the effect of a wider gender rent differential (if acting as a proxy for a higher rent level) should be larger in high rent industries. Our mechanism makes no such prediction since taste-based discrimination plays no role—wider gender rent differentials should affect gender wage differentials within all industries.

To explore this empirically, for each year we divide industries into a high paying (“good”) sector and low paying (“bad”) sector. We do so on the basis of the estimated gender-neutral national premia, and choose the cut-off premium so that the size of the sectors are roughly equal.³³ We then perform our above analysis on samples of workers in the good and bad sectors separately to see whether rent differentials have a stronger impact in the good sector.

In Table 6, we estimate (28) separately for those industries which were in the good and bad sectors in 1980, using the years 1980-2000. Columns 1-3 show the 2SLS results for our baseline sample, and columns 4-6 show the results for non-college workers. In both cases it is clear that rents have an important effect on wages in the low-paying sector. In fact, the impact appears to be *larger* in the low-paying sector than in the high-paying sector—particularly for non-college workers. This finding suggests that the link between industrial structure and gender wage gaps that we are stressing is not being driven by good jobs creating an environment amenable to taste-based discrimination.

4.2.3 Alternative Specifications and Sample Restrictions

In Table 7 we probe the robustness of our baseline empirical results to a variety of alternative specifications and time periods. First, as discussed above, our baseline estimates focus on the period from 1980-2000. However, as shown in Figures 2 and 3, much of the relative improvement in women’s wages and rents at the national level are concentrated in the 1980s. Thus, we would expect our mechanism to be particularly relevant during the 1980s. In column 1 of Table 7 we re-estimate our model using only this time period. The estimated coefficient on R_{get} is 1.13, which is slightly lower than our preferred specification in Table 2 but this difference is not statistically significant.

In the column 2 of Table 7, we use the 2010 ACS data to construct a panel from 1980-2010. Using this longer panel results in slightly lower, but statistically significant estimates of γ . This longer panel also allows us to include gender-city time trends in our specification. Recall that the main threat to our identification strategy is the omission of time-varying gender-city variables. The inclusion of linear gender-city time trends helps to account, as an example, for adoption of gender-bias technology that is smooth over time. Including the time trends reduces the available variation in R_{get} that we have to work with, but still results in a statistically significant estimate of γ of 0.83 (column 3).

In our baseline empirical work, we chose to limit our sample to those working full-time, full-year in order to focus on a group of workers whose composition over time might be reasonably stable. In column 4 of Table 7 we assess the robustness to including both full- and part-time workers in our sample. One advantage of doing so is that adding part-time workers increases the sample size we have to work with. A disadvantage is that, by treating the wages of all workers equally irrespective of how many hours they supply, the wage distribution will not be representative of the total number of hours worked in the economy. This might be of particular concern when analysing gender differences in wages if women tend to work less

³³See section E.2 for actual classifications.

Table 6: Estimates for Good Jobs vs. Bad Jobs: Years 1980-2000

	All			Non-College		
	(1)	(2)	(3)	(4)	(5)	(6)
	All Jobs	Good Jobs	Bad Jobs	All Jobs	Good Jobs	Bad Jobs
R_{gct}	1.39** (0.40)	0.99** (0.39)	1.60** (0.47)	1.98** (0.57)	1.10** (0.47)	2.48** (0.69)
u_{gct}	-0.0075 (0.015)	-0.0068 (0.015)	-0.0071 (0.018)	0.0018 (0.019)	0.013 (0.017)	-0.0035 (0.024)
Observations	16476	8372	8104	15264	6992	8272
R^2	0.990	0.984	0.988	0.987	0.982	0.982
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage:						
Z_{gct}	1.75** (0.33)	1.76** (0.32)	1.75** (0.34)	1.44** (0.30)	1.45** (0.31)	1.44** (0.30)
F -Stat.	28.50	30.52	27.15	22.85	22.46	22.59
p -val	0.00	0.00	0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The Partial R^2 shows the fraction of the cross-city variation of the gender gap that is explained by R_{gc} after taking into account all other regressors. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

than men. One potential compromise is to weight each observation by the number of hours worked, which gives more importance to full-time workers. We show hours-weighted results in column 5. In columns 4 and 5, the estimate of γ is 0.91 and 1.36, respectively. These estimates are on the same order of magnitude as our preferred baseline results which focus on the restricted sample of full-time workers.

We then proceed to further address two selection issues that potentially persist despite our IV strategy and controls. The first issue is women’s self-selection into the workforce. Not only has female labor market participation changed dramatically over time (Mulligan and Rubinstein (2008)), it also has substantial variation across cities within a given year. A potential concern for our framework is that participation decisions are related to R_{gct} . For example, one might expect women to participate in the workforce at a higher rate in cities where R_{fct} is high (since this represents favourable employment opportunities) compared to when R_{fct} is low. Suppose, further, that women positively select into the labor force (higher ability women tend to be drawn into the workforce before lower ability women). In this case, the average quality of women in the labor market will be, all else equal, negatively related with R_{fct} since marginal quality will fall as participation expands. Note that in this scenario, female self-selection could work against us finding evidence of spill-over effects since selection bias of this form will be associated with higher gender pay gaps at the city level but potentially lower rent differentials. It is for this reason that our base line controls include female labor force participation.

In Table 7 we attempt to control for women’s self-selection into the labor force using an alternative method. In particular, we borrow the approach from recent work by Mulligan and Rubinstein (2008), who show that the nature of self-selection into the labor market by women explains much of the over-time improvement in gender pay gaps. In their paper, they use a Heckman-two step selection correction estimator to estimate selection corrected pay gaps. While exclusion restrictions for this type of estimator are notoriously hard to come by, we follow Mulligan and Rubinstein (2008) in exploiting the presence of young children at home along with family structure. To implement their approach in our framework, we estimate a first-stage probit equation where the dependent variable is a binary indicator for full-time employment. We control for the same X_i vector as discussed above, but add to this indicators for the presence of young children, marital status, and their interaction. Using the estimated coefficients from this procedure, we construct the Inverse Mills Ratio (IMR), and include this along with marital status in our first-stage estimation of (20). In implementing this, we assume, as in Mulligan and Rubinstein (2008), that selection bias for men is zero and allow the effect of the included IMR to vary by education. Further details on this procedure can be found in Appendix D.1. We interpret the regression adjusted wages from this procedure to be corrected for non-random female labor force participation. Column 6 of Table 7 contains the results of this exercise. We find that this additional control has little effect on our baseline results.³⁴

³⁴Our first-stage results with the inclusion of the IMR ratio are generally consistent with Mulligan and Rubinstein (2008). In particular, we find that the coefficient on the IMR switches signs from negative in the early Census years to positive in the most recent, particularly for those with less education. The fact that the inclusion of the IMR does not impact our point estimates of γ may not be surprising since we rely on quite different variation compared to Mulligan and Rubinstein (2008), and our control for city-level participation does not have much impact. It should also be noted that there is an ongoing debate about the Mulligan and Rubinstein (2008) claim that the narrowing of the gender gap is completely explained by selection. Machado (2012) and Herrmann and Machado (2012) question the identification strategy of Mulligan and Rubinstein

A second selection issue is that workers could self-select into cities (Dahl (2002)): places with high R_{gct} could systematically attract gender g workers with particular unobserved abilities. This may be particularly important in the wake of the decline in manufacturing where affected cities saw sizeable population declines. This will produce an upward bias if it is the high ability workers that re-locate in high rent cities. Although plausible, it could also be argued that migration will be most pronounced among low ability workers since these are the workers most likely to lose their jobs. Furthermore this sort of threat is less problematic in our context to the extent that re-location occurs at the *family* level and that workers sort into families in a manner that is positive assortative on ability. For instance, if a high ability male relocates then there will be a tendency for him to be accompanied by a high ability female. In other words, selective migration will tend to occur on the basis of absolute average (across genders) rent levels whereas our identification comes from gender differences in rent levels. In any case, we adopt the selection correction procedure suggested by Dahl (2002), who estimates a local college premia of young men while accounting for non-random sorting across U.S. states. Dahl develops a two-step selection correction estimator to deal with the high dimension scenario of workers choosing to live and work in any state, given their birth state, that is based on an index sufficiency assumption that is implemented by estimating and including a reduced set of selection probabilities into the estimated second-step wage equation. We implement Dahl’s two-step procedure in our context by including, for example, a low-order polynomial of the estimated probability that an individual was born in state s and is observed in city c . A detailed discussion of our selection correction procedure is presented Appendix D.2. Note that both the Dahl (2002) and the Mulligan and Rubinstein (2008) selection corrections can be implemented simultaneously, and we present these results in column 7 of Table 7. Again, we find that this additional control has little effect on our point estimates.

In the last three columns of Table 7, we assess the robustness of our results to focusing on younger workers, larger cities, and an alternative definition of cities. In particular, in column 8, we re-estimate our model focusing on younger workers (less than 35 years old). On the one hand, this group may be less sensitive to the local forces emphasised in our framework if younger groups are more mobile. On the other hand, this group maybe more sensitive if younger workers’ wages are more directly related to current labor market conditions (Oreopoulos et al., 2012). As shown in column 8, this sample restriction produces an estimate of γ of 1.02, which is in line with our baseline estimates. Column 9 of Table 7 restricts the analysis to the largest 50 cities in the US, which results in a slightly larger estimate of γ of 1.50. In column 10, we define cities using 1999 CMSA definitions rather than 1990 SMSA definitions. CMSA definitions produce larger cities, on average, because these definitions may combine several SMSAs. Working with this alternative definition has virtually no impact on our estimates.

(2008), suggesting that imposing the same selection rule for all women is overly restrictive and that family structure is correlated to pre-labor market characteristics. Machado (2012) shows that the gender gap declines for a sub-sample of ‘always working’ women (women with children who would likely work if they did not have children) and Chandrasekhar et al. (2012) use a bounding technique to show that a decline in the gender gap cannot be rejected in the Mulligan and Rubinstein (2008) data. Bar et al. (2015) argue that Mulligan and Rubinstein (2008) results are biased and overstate the importance of selection in explaining the decline in the gender gap.

Table 7: Alternative Specifications

2SLS										
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
R_{gct}	1.13** (0.34)	0.91** (0.28)	0.83** (0.34)	0.91** (0.33)	1.17** (0.36)	1.31** (0.41)	1.29** (0.39)	1.02** (0.38)	1.50** (0.55)	1.36** (0.43)
u_{gct}	-0.016 (0.017)	0.0036 (0.012)	-0.013 (0.012)	-0.011 (0.014)	-0.0070 (0.015)	-0.0096 (0.017)	-0.0086 (0.017)	-0.0058 (0.013)	-0.011 (0.017)	0.0054 (0.016)
Observations	10424	21198	21198	17528	17528	14524	14524	13996	9630	14838
R^2	0.993	0.988	0.989	0.989	0.990	0.989	0.988	0.983	0.991	0.991
Fixed Effects:										
City× Ind.× Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Female× Ind.× Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Female× Ind.× City	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Female× City trends	No	No	Yes	No	No	No	No	No	No	No
Controls	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Years	1980-1990	1980-2010	1980-2010	1980-2000	1980-2000	1980-2000	1980-2000	1980-2000	1980-2000	1980-2000
Sample				Inc. PT	Hours			Young	50 Largest	CMSA
Selection Correction						IMR	IMR,Dahl			
First-Stage:										
	2.81** (0.50)	1.45** (0.24)	1.62** (0.38)	1.72** (0.36)	1.76** (0.34)	2.00** (0.37)	2.04** (0.38)	1.87** (0.41)	1.82** (0.43)	1.77** (0.34)
F -Stat.	32.01	35.44	18.52	22.35	26.62	28.47	29.33	20.78	18.28	26.43
p -val	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities Census and ACS data. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. All columns estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure.

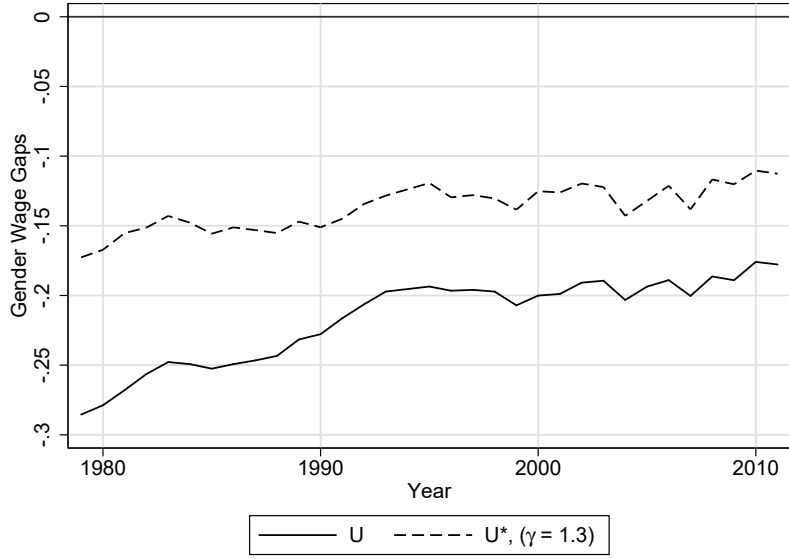


Figure 6: Counterfactual gender wage gaps over time

5 Implications for Time Series Evidence

In order to elaborate upon the quantitative implications of our estimates, we return to the time series evidence presented in the Introduction. In that section we illustrate the time pattern of the adjusted wage gap and of a rent gap (Figure 2). Given the subsequent analysis, the rent gap that we plot is actually more precisely described as the explained component³⁵ of the rent gap and thus represents E_t from section 2.1. Similarly, the adjusted wage gap that we plot measures the unexplained component, U_t , from section 2.1. In this section we use these series, along with our estimates of γ , to produce the counter-factual wage gap suggested in section 2.1. This exercise suggests that the relatively rapid narrowing of the rent gap in the 1980s is able to fully account for the relatively rapid narrowing of the adjusted wage gap during this period. We then attempt to understand *why* the rent gap narrowed relatively rapidly during this period by decomposing changes into parts due to changes in gender-specific employment patterns and to changes in industry premia.

5.1 Counterfactual Gender Wage Gap

In section 2.1 we identified U^* as a counter-factual gender gap of interest. It tells us the gender pay gap remaining after taking out the impact of industrial structure on the quality of outside prospects. This series is plotted using the CPS data in Figure 6 using $\gamma = 1.3$.

There are two main points illustrated by the figure. First, the effect of industrial structure on adjusted wage gaps is substantial. Since U^* is around 65% of the magnitude of the adjusted gap, we conclude that differences in the quality of outside prospects generated by industrial

³⁵This is because when computing the rent gap from the CPS data we used gender-neutral industry premia. Full rent differences would use gender-specific premia, but recall that such a measure is infeasible since only relative industry premia are identified.



Figure 7: Gender Wage Gaps Relative to 2010

structure can account for around 35% the adjusted wage gap in any given year.

Second, the relatively steep narrowing of the rent gap in the 1980s fully accounts for the relatively steep narrowing of the pay gap during the 1980s. To see this, note that U^* does not display the relatively rapid decline displayed by the adjusted gap during 1980s. The adjusted gap narrows by around 5 log points in the 1980s and around 2.5 log points in the 1990s and the 2000s. In contrast, U^* declines by around 2 log points in the 1980s, 1990s and the 2000s. The rapid narrowing of the adjusted gender gap in the 1980s relative to later periods is therefore fully accounted for by the relatively rapid narrowing of the rent gap during the 1980s. Figure 7 provides a clearer visual representation of this by plotting the gaps relative to their value in 2010. Whereas the transition of the adjusted gap to the 2010 involves relatively rapid change during the 1980s, the transition of U^* is roughly even throughout the period.

5.2 Decomposing Rent Differences Over Time

If the rent gap plays such an important role, then one would naturally like to better understand the drivers of this variable's evolution. To this end, we decompose the gender rent gap to get a sense of the relative importance of changing employment distributions and changing industrial rents. As described in the Introduction, given a vector of date t industry premia d_t and a vector of national employment distributions by gender, π_{gt} , the national gender rent gap is

$$E_t \equiv \underbrace{[\pi_{ft} - \pi_{mt}]}_{\pi_t^{\text{gap}}} \cdot d_t. \quad (32)$$

To help interpret changes in E_t recall that a typical element of π_t^{gap} —i.e. $\pi_{fnt} - \pi_{mnt}$ —is a measure of the ‘femaleness’ of industry n in year t .³⁶ Notice that the average femaleness across industries is zero by construction. We can decompose the evolution of E_t into parts due to changes in industry premia (“prices”), to changes in employment distributions (“shares”), and to the joint change in these (“interaction”). Specifically, pick some fixed year, \bar{t} , and decompose:

$$\Delta E_t \equiv \underbrace{E_t - E_{\bar{t}}}_{\text{total}} = \underbrace{[\pi_t^{\text{gap}} - \pi_{\bar{t}}^{\text{gap}}] \cdot d_{\bar{t}}}_{\text{shares}} + \underbrace{\pi_{\bar{t}}^{\text{gap}} \cdot (d_t - d_{\bar{t}})}_{\text{prices}} + \underbrace{[\pi_t^{\text{gap}} - \pi_{\bar{t}}^{\text{gap}}] \cdot [d_t - d_{\bar{t}}]}_{\text{interaction}}. \quad (33)$$

The “shares” component tells us the change in the rent gap had only employment distributions changed, the “prices” component tells us the change in the rent gap had only premia changed, and the “interaction” component tells us the additional consequences of both changing. From another perspective, the “prices” component captures the extent to which relative rents shifted toward initially female jobs. This is the variation used by our instrument, Z_{gct} discussed in section 3.3. The “shares” component captures the extent to which jobs that became more female were initially good jobs. The “interaction” component captures the extent to which jobs that became more female were ‘improving’ jobs (i.e. those experiencing a relatively large change in industry premia).

Since the “shares” component is the only component to use a time-invariant set of industry premia, only it can be further decomposed to give a sense of the evolution of wage *levels*. Specifically, we can measure whether a larger “shares” component is due to women moving into initially better jobs or due to men moving into initially worse jobs. The decomposition is:

$$[\pi_t^{\text{gap}} - \pi_{\bar{t}}^{\text{gap}}] \cdot d_{\bar{t}} = \underbrace{[\pi_{ft} - \pi_{f\bar{t}}] \cdot d_{\bar{t}}}_{\text{shares (female)}} - \underbrace{[\pi_{mt} - \pi_{m\bar{t}}] \cdot d_{\bar{t}}}_{\text{shares (male)}} \quad (34)$$

The “prices”, “shares”, and “interaction” components are displayed in Figure 8, using the CPS data with $\bar{t} = 1980$. We see that price changes have played a very important role, especially since the mid 1980s. Both prices and shares contributed to the narrowing of the rent gap between the mid-1980s and mid-1990s, but the effect of prices was twice as large. The interaction effect played no role during this episode.

Interestingly, the slowdown in the narrowing of the rent gap (and thus the pay gap) since the mid-1990s seems entirely due to a slowdown in the “shares” component and an actual decline in the “interaction” component. That is, rent changes continued to favour industries that were initially more female (“prices”), but this was offset by two forces. The jobs that were becoming increasingly female (i) ceased being those that were relatively good (“shares”), and (ii) started becoming those that were ‘declining’ (in the sense of experiencing the largest declines in industry premia). Figures 11 - 13 in the appendix give a sense of which industries are responsible for each of these changes.

To unpack the “shares” component, Figure 9 shows its decomposition into male and female parts. We see that the rise of the “shares” component until the mid-1990s was due mostly to men shifting out of jobs that were initially good. It appears that women saw a modest shift

³⁶Figure 14 in the appendix displays an industry’s premium against this measure of femaleness for each of the census years.

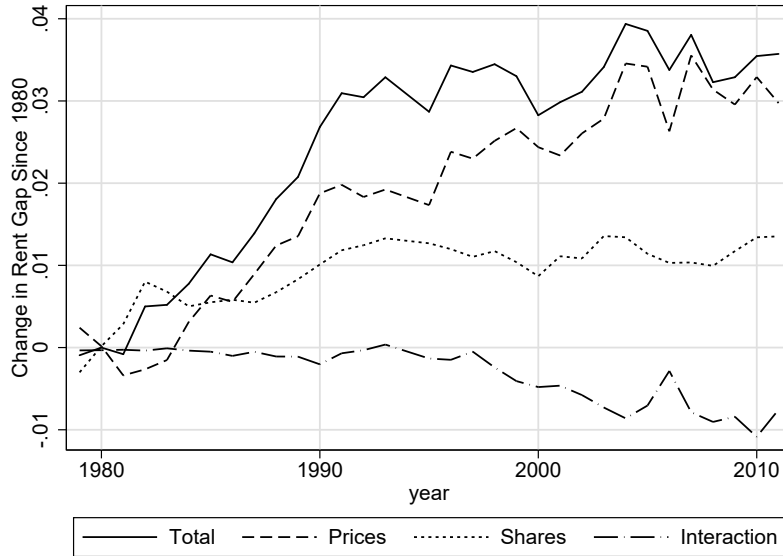


Figure 8: Decomposing Changes in the Rent Gap

toward better jobs, but by the mid-1980s women too had started to shift out of jobs that were initially good. The period between the mid-1980s and the mid-1990s thus saw initially good jobs become more female because men were moving out of such jobs faster than were women. By the mid-1990s men and women were shifting out of initially good jobs at around the same pace, leaving the overall “shares” component relatively stable. This decomposition provides evidence that is consistent with the rapid narrowing of the adjusted pay gap in the 1980s being more about men losing out than women gaining. Indeed, by the mid-1990s when the narrowing slowed, men and women had both experienced absolute declines in the “shares” component of rent but men had lost almost three times as much.

To summarize, the period until the mid-1990s was characterized by initially male jobs losing rents relative to initially female jobs, and to a lesser extent because initially good jobs became more female (owing to men shifting out of such jobs to a greater extent than women). Since the mid-1990s rents continue to shift toward initially female jobs, but this is partially offset by rents also shifting toward jobs that had become less female.

6 Conclusions

Our broad goal in this paper was to explore whether differences in the quality of alternative employment opportunities stemming from an economy’s industrial structure are relevant for understanding gender pay differences. Our results suggest that such factors are indeed very important—this ‘indirect’ effect of a city’s industrial structure is estimated to be around 1.4 times as large as the direct compositional effect of industrial composition.

But how does our analysis add to existing perspectives? In one sense, the implication is negative: the celebrated ‘gains’ of women during the 1980s may be more accurately characterized as the ‘losses’ of men, owing to declining male outside prospects during this period. But

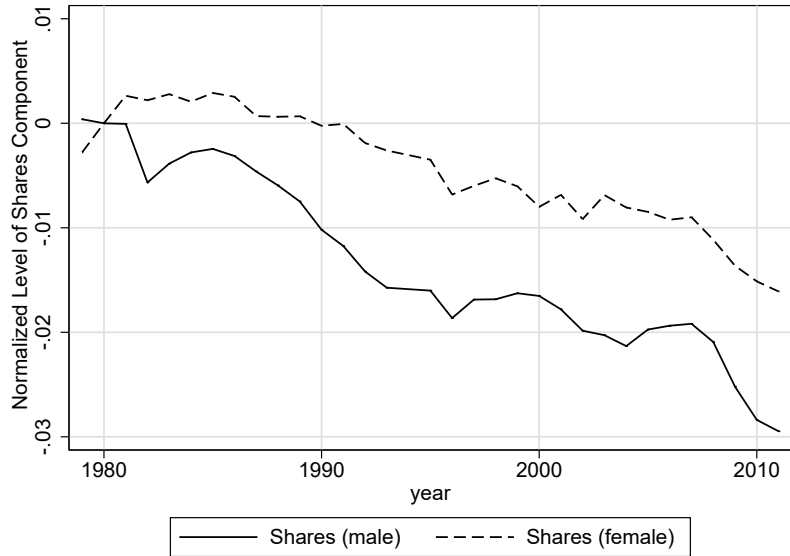


Figure 9: Decomposing Shares Component

in another sense the implication is very positive: the future impact of technological changes on the gender gap are likely to be substantially understated. For instance, Goldin (2014) argues that gender wage equity will require a change in the nature of work, specifically to accommodate greater temporal flexibility. Our analysis suggests that the impact of such a change is under-appreciated—as technologies alter the nature of work in ways that boost the representation and pay of women in some sectors, there will be the sort of direct effect highlighted by Goldin (2014). However we can expect a further indirect effect, of at least the same magnitude, owing to the fact those affected sectors offer higher quality outside employment opportunities to all women.

This latter aspect highlights policy implications that go beyond equal pay legislation—we estimate that the within-industry pay gap would narrow by around 35% if gender differences in the factors driving industrial employment distributions were eliminated. However, to achieve this there is clearly a need for future research aimed at explaining *why* some industries are more ‘male’ than others. Such research is particularly valuable since gender pay gaps influence a range of other socially important outcomes such as rates of marriage, labor force participation, and domestic violence (Bertrand et al. (2015), Aizer (2010)).

Finally, we reiterate that our analysis can be easily applied to any other pay gap of interest, such as racial pay gaps. We are currently exploring whether the analysis here is able to shed light on the skill wage gap and are hopeful that this sort of analysis will prove useful in other countries and time periods.

Appendix

A Data

A.1 Census Data

The Census data was obtained with extractions done using the IPUMS system (see Ruggles et al. (2010)). The files were the 1980 5% State (A Sample), 1990 State, 2000 5% Census PUMS, and the 2009-2010-2011 American Community Surveys. The initial extraction includes all individuals aged 22 - 54 not living in group quarters. All calculations are made using the sample weights provided. We focus on the log of hourly wages, calculated by dividing wage and salary income by annual hours worked (usual hours worked \times annual weeks worked). We impute incomes for top coded values by multiplying the top code value in each year by 1.5. Since top codes vary by State in 1990 and 2000, we impose common top-code values of 140,000 in 1990 and 175,000 in 2000.

A consistent measure of education is not available for these Census years. We use indicators based on the IPUMS recoded variable `EDUCREC` that computes comparable categories from the 1980 Census data on years of school completed and later Census years that report categorical schooling only. To calculate potential experience (age minus years of education minus six), we assign group mean years of education from Table 5 in Park (1994) to the categorical education values reported in the 1990 and 2000 Censuses and the ACS.

Census definitions of metropolitan areas are not comparable over time since, in general, the geographic areas covered by them increase over time and their definitions are updated to reflect this expansion. The definition of cities we use attempts to maximize geographic comparability over time and roughly correspond to 1990 definitions of MSAs provided by the U.S. Office of Management and Budget.³⁷ To create geographically consistent MSAs, we follow a procedure based largely on Deaton and Lubotsky (2003) which uses the geographical equivalency files for each year to assign individuals to MSAs or PMSAs based on FIPs state and PUMA codes (in the case of 1990 and 2000) and county group codes (for 1970 and 1980). Each MSA label we use is essentially defined by the PUMAs it spans in 1990. Once we have this information, the equivalency files dictate what counties to include in each city for the other years. This process results in our having 152 MSAs that are consistent across all our sample years. Code for this exercise was generously provided by Ethan G. Lewis. Our definitions differ slightly from those in Deaton and Lubotsky (2003) in order to improve the 1980-1990-2000 match. We further restrict our analysis to the 100 largest cities based on population estimates from 1980.

We use an industry coding that is consistent across Censuses and is based on the IPUMS recoded variable `IND1990`, which recodes census industry codes to the 1990 definitions. We aggregate this variable into 45 detailed industry groups based on standard BLS definitions

For all analysis using wage data, we further restrict the sample to those (1) currently employed at the time of the census, (2) with positive wage and salary income, (3) those who are full-time, full year worker, defined as a worker who usually works 35 hours per week and worked at least 40 weeks in the year prior to the Census. The ACS does not contain a continuous measure of weeks worked in the year prior to the survey after 2008. Instead, an

³⁷See <http://www.census.gov/population/estimates/pastmetro.html> for details.

interval of weeks worked is provided. We convert this into a continuous measure by taking the mid-point in each interval.

A.2 Merged Outgoing Rotation Group Current Population Survey

MORG CPS data from 1979-2014 are downloaded from the NBER³⁸. Our initial extractions included all individuals between the ages of 22-55. Prior to 1992, education was reported as the number of completed years. In 1992 and after, education is reported in categories as the highest grade/degree completed. We convert categories to years of completed school in the post-1991 data based on Park (1994). We further restrict our analysis to full-time workers by selecting workers who work at least 30 hours in a week. The construction of our wage data closely follows Lemieux (2006). Wage data is based on those who report employment in reference week. In all wage calculations, we set allocated wages to missing. Our hourly wage measure is based on reported hourly wage for those who report hourly payment and not adjusted for topcoding. For workers who are not paid hourly:

1. We use edited weekly earnings. For the years 1984-1986, we use unedited earnings due to the higher topcode value.
2. Adjust topcoded wages by a factor of 1.4.
3. Divide the result by usual hours worked per week.

B Further Details on Construction of Motivating Figures

Figure 1 uses a sample from the 1980 Census constructed as in our main analysis as described in Section A.1 above. In order to maintain consistency with the main empirical analysis of the paper, and especially to maintain comparability with Figure 5, we estimate industry premia in the same way that we obtain the ‘gender neutral’ industry premia used in the construction of our instrument as described in section 3.2 and footnote 26. That is, we first estimate the log wage regression

$$\ln W_{i,gnc} = X_i' \beta_g + D_{gnc} + \epsilon_i. \quad (\text{B-1})$$

where D_{gnc} is a gender×industry×city fixed effect and X_i includes controls for a quartic in potential experience; hispanic, black, and immigration dummies; an indicator for whether an individual is observed in a city located in one’s birth state – all interacted with education (four categories) – and four occupation dummies. All covariates have coefficients that vary by gender. We then decompose the three-way fixed effect into three two-way fixed effects by estimating

$$D_{gnc} = D_{gn} + D_{nc} + D_{gc} + \epsilon_{gnc} \quad (\text{B-2})$$

weighted by the cell size. We then take the fixed effects that vary by industry and define the premium as $\kappa_{gnc} \equiv D_{gn} + D_{nc}$. Finally, we extract the industry component of this by estimating

$$\kappa_{gnc} = d_n + d_g + d_c + \epsilon'_{gnc}, \quad (\text{B-3})$$

³⁸Links are http://www.nber.org/data/cps_may.html and <http://www.nber.org/data/morg.html>

and take d_n to be the industry premium, normalizing so that the average premium is zero. In constructing $\pi_{fn} - \pi_{mn}$, we take $\pi_{gn} \equiv \text{emp}_{gn} / \sum_j \text{emp}_{gj}$ where emp_{gn} is the sum of the sampling weights in the gender \times industry cell. The markers are weighted by $\text{emp}_n \equiv \text{emp}_{fn} + \text{emp}_{mn}$.

Figure 2 uses a sample from the CPS as described in Section A.2 above. We estimate the following log wage regression year-by-year, using provided sampling weights:

$$\ln W_i = X_i' \beta + d_n + \alpha \cdot \text{female}_i + \epsilon_i, \quad (\text{B-4})$$

where d_n are industry fixed effects (52 categories), female_i is a female indicator, and X_i contains an education indicator (5 categories), potential experience and its square, an occupation indicator (4 categories), and a black indicator. The within-industry gender wage gap is the estimate of α . The rent gap is $\sum_j (\pi_{fj} - \pi_{mj}) \cdot d_j$, where $\pi_{gj} \equiv \text{emp}_{gj} / \sum_j \text{emp}_{gj}$ where emp_{gj} is the sum of the sampling weights in the gender \times industry cell.

Figure 3 is constructed using CPS data from 1979-2011. The sample is restricted to full-time wage and salary workers between the ages of 22-54. Wage calculations in the figure hold the age distribution fixed at the 1990 distribution. For the non-college group, which includes high school drop-outs, graduates, and those with some post secondary, we additionally hold the the composition of these education sub-groups constant at their 1990 distribution. Real wage are deflated by to 2000 dollars using the ‘all item’ CPI obtained from <http://data.bls.gov/cgi-bin/surveymost?cu>.

Figure 4 uses a similar procedure to Figure 1, taking the female-male difference in D_{gc} from (B-2) to be the city-level within-industry gender gap. This is performed using 1980 Census data and then repeated using 2000 Census data.

C Further Details on Model and Empirical Implementation

C.1 Unemployment Flow Utility

If workers received a flow utility of b_{gc} while unemployed, so that (2) becomes

$$rU_{gc} = b_{gc} + \sum_n q_{gnc} \cdot \max \{E_{gnc} - U_{gc}, 0\}, \quad (\text{C-5})$$

then the wage equation (10) becomes

$$w_{gnc} = \psi_{gnc} + (1 - \delta_{gc}) \cdot b_{gc} + \delta_{gc} \cdot \bar{w}_{gc}, \quad (\text{C-6})$$

where ψ_{gnc} is re-defined to be $\psi_{gnc} = \phi \cdot (P_{gnc} - rk_{nc} - b_{gc})$. As such, the reduced form wage equation (11) becomes:

$$w_{gnc} = b_{gc} + \psi_{gnc} + \gamma(u_{gc}) \cdot \sum_j \pi_{gjc} \cdot \psi_{gjc}. \quad (\text{C-7})$$

The mechanism we highlight, and the parameter we attempt to estimate, is clearly unaffected. The re-definition of ψ_{gnc} has no impact since we already allow for a gender \times city component via ξ_{gc} . Similarly, the empirical approach is unaffected since b_{gc} would simply be absorbed into the gender \times city component of the error term. That is, $\nu_{gc} = (1 + \gamma) \cdot \xi_{gc} + b_{gc}$.

C.2 Deriving the wage equation

To derive the wage equation, use $E_{gnc} \geq U_{gc}$ and $J_{gnc} \geq V_{nc}$ in (1) and (2) to derive an expression for $E_{gnc} - U_{gc}$. Similarly, (3) and (6) can be used to derive an expression for $J_{gnc} - V_{nc}$. Using these expressions in the bargaining condition (5) gives us an expression connecting wages in a gender-industry-city w_{gnc} to average wages across industries within the gender-city $\bar{w}_{gc} \equiv \sum_j \pi_{gjc} w_{gjc}$, where π_{gnc} is the equilibrium employment share given by (9). Specifically:

$$w_{gnc} = \psi_{gnc} + \delta_{gc} \cdot \bar{w}_{gc}, \quad (\text{C-8})$$

where $\psi_{gnc} \equiv \phi \cdot (P_{gnc} - rk_{nc})$ captures job-specific rent, and $\delta_{gc} \equiv (1 - \phi) \cdot \sum_j q_{gjc} / [r + s + \sum_j q_{gjc}]$. To see that δ_{gc} varies by (g, c) because of differences in the equilibrium *quantity* of alternative employment opportunities, note that it can be expressed as a (decreasing) function of the unemployment rate by using the steady state condition, (8): $\delta_{gc} = \delta(u_{gc}) \equiv (1 - \phi) \cdot [1 - u_{gc}] / [1 + u_{gc} \cdot (r/s)]$. The reduced-form expression is therefore:

$$w_{gnc} = \psi_{gnc} + \gamma(u_{gc}) \cdot \sum_j \pi_{gjc} \cdot \psi_{gjc}, \quad (\text{C-9})$$

where $\gamma(u_{gc}) \equiv \delta(u_{gc}) / [1 - \delta(u_{gc})]$.

To account for heterogeneity in the quantity of alternative employment opportunities—i.e. heterogeneity in u_{gc} —we follow Beaudry et al. (2012) by taking a linear approximation of w_{gnc} around a ‘symmetric’ benchmark that ensures the unemployment rate is equal for all groups. Let $w_{gnc} = w_{gnc}(u_{gc}, \{\psi_{gnc}\}_n)$ denote the wage equation described in (C-9). Taking a linear approximation of this function around the point $(u_{gc} = \bar{u}, \{\psi_{gnc} = \bar{\psi}\}_{gnc})$ to get:

$$w_{gnc} \simeq \psi_{gnc} + \gamma \cdot \sum_j \pi_{gjc} \cdot \psi_{gjc} + \gamma_1 \cdot [u_{gc} - \bar{u}], \quad (\text{C-10})$$

where $\gamma \equiv \gamma(\bar{u})$ and $\gamma_1 \equiv \gamma'(\bar{u}) \cdot \bar{\psi}$.

To work in logs, fix an arbitrary value w_0 (e.g. the value for a specific omitted group or the average value across cells etc.), and make an approximation around this point:

$$(\ln w_{gnc} - \ln w_0) = \ln(w_{gnc}/w_0) \simeq \frac{w_{gnc} - w_0}{w_0}. \quad (\text{C-11})$$

Using this with (C-10) gives:

$$\ln w_{gnc} \simeq \text{const.} + \tilde{\psi}_{gnc} + \underbrace{\gamma \cdot \sum_j \pi_{gjc} \cdot \tilde{\psi}_{gjc}}_{\equiv \tilde{R}_{gc}^*} + \alpha \cdot u_{gc}, \quad (\text{C-12})$$

where $\tilde{\psi}_{gnc} \equiv \psi_{gnc}/w_0$ and $\alpha \equiv \gamma_1/w_0$. The constant, which equals $\ln w_0 - 1 - \alpha \cdot \bar{u}$, can be ignored (it will be part of the $X_i' \beta^g$ term in (21)).

D Selection Correction

D.1 Mulligan and Rubinstein (2008)

We implement the Heckman-two-step procedure in the same manner as described in Mulligan and Rubinstein (2008), except for small changes due to the fact that we use a different data source and have a different outcome equation.

1. Following (Mulligan and Rubinstein, 2008), we assume that men’s inverse Mill’s ratio (IMR) is zero, or that men have no selection bias.
2. We form an variable for the number of children in the household under 5 years old. The categories are zero, 1, 2, 3, and 4 or more.
3. We restrict age to be between 25-44 following Machado (2012). The rational for this restriction is that the exclusion restriction uses information on fertility and nearly all the variation comes from women under the age of 45.
4. We estimate a first-stage probit equation where the dependent variable is an indicator for a women working with a valid wage. The right hand side of the probit contains all of the demographic variables we use in our baseline procedure described in the text, but additionally includes an indicator for marriage and a marriage-number of children interaction. This latter interaction becomes the exclusion restriction, as in Mulligan and Rubinstein (2008). Additionally, we include a full set of city dummy variables.
5. To be consistent with our second stage, we allow the effects marriage and the marriage-children interaction on working to vary by education level (4 groups) by fully interacting these variables with education. This is a slight departure from Mulligan and Rubinstein (2008).
6. Having estimated the probit, we construct an estimated IMR for women, and include this in our first-stage regression adjustment procedure described in the text. When including the estimated IMR and a marriage indicator, we allow the coefficients on these variables to vary by education. This is a slight generalization of Mulligan and Rubinstein (2008), who restrict the selection effect to be the same for education groups, but is consistent with the arguments of Machado (2012) who suggest that imposing the same selection function for all women is overly restrictive.

D.2 Dahl (2000)

As described in the paper, our main approach to addressing the issue of selection on unobservables of workers across cities follows Dahl (2002). Consider a model where each worker has a potential or latent wage that she’d earn in each possible city, and chooses to live in the city with the highest wage net of moving costs. To model this explicitly, consider the log wage equation, (21), conditional to worker i choosing city c :

$$\mathbb{E}[\ln W_{i,gnc} | d_{i,gnc} = 1] = X_i' \beta_g + \mu_{gnc} + \mathbb{E}[\epsilon_{i,gnc} | d_{i,gnc} = 1], \quad (\text{D-13})$$

where, as before, i indexes individuals, g , n , and c index gender, industry and city. The variable $d_{i,gnc}$ is a dummy variable equalling one if i chooses to live and work in city c . The

error mean term, $\mathbb{E}[\epsilon_{i,gnc}|d_{i,gnc} = 1]$, is non-zero if worker city selection is not independent of the unobserved component of wages. In this case, estimating (D-13) via OLS will lead to coefficients that are, in general, inconsistent.

Dahl argues that the error mean term in (D-13) can be modelled as a full set of probabilities that a person born in a given state would choose to live in each possible city in the Census year. Given the number of states and cities, this would lead to difficulties implementing an estimator that models the error mean term in such a flexible way. Instead, Dahl presents a sufficiency assumption under which the error mean term is restricted to be only a function of the first-best choice, or the choice that was actually made by i . This assumption suggests that one need only model the error mean term as a function of the probability that an individual born in state s is observed in city c . The sufficiency condition essentially says that two people with the same probability of choosing to live in a given city have the same error mean term in their regression: knowing the differences in their probabilities of choosing other options is not relevant for the size of the selection effect in the process determining the wage where they actually live. In Dahl’s paper, he relaxes this assumption somewhat, and models the error mean term as a function of the first-best probability, and the retention probability (the probability that i is born in state s and chooses to stay).

Dahl proposes a non-parametric estimator in which he divides individuals up into cells defined by discrete categories for education, age, gender, race and family status. He then uses the proportion of people within the cell that is relevant for person i who actually made the move from i ’s birth state to his destination and the proportion who stayed in his birth state as the estimates of the two relevant probabilities. This is a flexible estimator which does not impose any assumptions about the distribution of the errors in the processes determining the migration choice. Dahl models the mean error term as a low-order polynomial of each of the selection probabilities.

We implement Dahl’s approach in the same manner apart from several small changes to adapt his approach to our context:

1. While Dahl focuses on selection of individuals across states, we focus on cities. However, we only know an individual’s state of birth, not the birth city. Thus, we calculate the probability an individual chooses city c who is born in state s .³⁹
2. Individuals observed in a city that is outside their birth state are called “movers.” Movers are divided into demographic groups based on age (5 groups), education (4 groups), gender, ethnicity (2 groups) and state of birth. The first-best probability is calculated as the fraction of each demographic group who chose c . We call this probability $m(x_i)$, where x_i denotes i ’s demographic group.
3. Mover’s retention probabilities are calculated as the fraction of each demographic group observed living in their state of birth. We call this $r(x_i)$.
4. Dahl excludes immigrants from his analysis, whereas we included them. We treat immigrants the same as “movers”, and divide non-US countries into 17 ‘sending’ states. We calculate their first-best choice by the fraction of each immigrant demographic group

³⁹For cities that span more than one state, we categorize person who is observed in a city that is at least partly in their birth state as a non-mover.

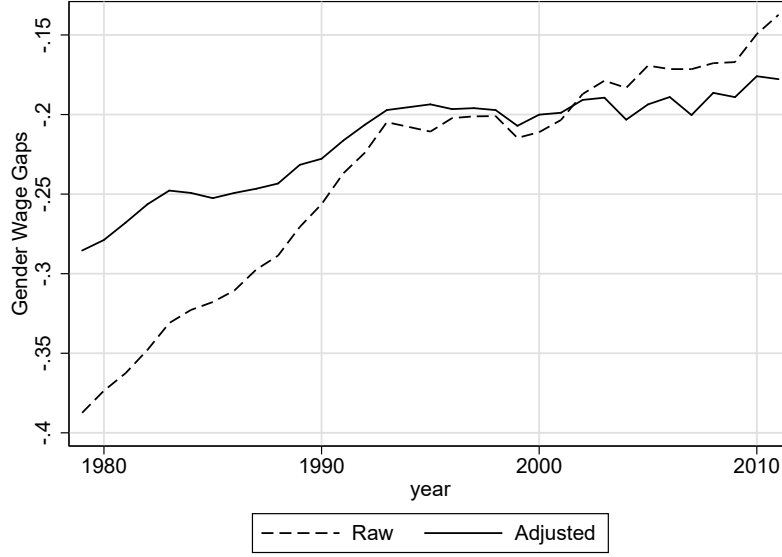


Figure 10: Raw and Adjusted Gender Pay Gaps Over Time

observed in city c born in country s . We call this $i(x_i)$. We ignore immigrant retention probabilities.

5. People who live in a city at least partly in their birth state are called “stayers”. Stayers are divided into demographic groups based on age (5 groups), education (4 groups), gender and ethnicity (2 group). Their staying probability is calculated as the fraction of each group born in s observed living in a city at least partly in s . We call this $p(x_i)$.

Having estimated $m(x_i)$, $r(x_i)$, $i(x_i)$, and $p(x_i)$, we include these estimated probabilities (and their squares) in our first-stage regression adjustment procedure, allowing, as with the other controls, their coefficients to vary by gender. These terms come in significant in our first-stage regressions, suggesting that selection across cities is potentially important. Thus, we estimate all of our specifications with and without this first-stage adjustment to assess sensitivity.

E Additional Results

E.1 Additional Motivating Figures

Figure 10 plots the within-industry gender pay gap alongside the ‘raw’ gender gap. The latter is simply the coefficient on a female dummy from a regression of log wages that does not control for anything else (using the same CPS data). The two series share the qualitative feature that there is a relatively rapid gap narrowing until the mid-1990s, after which the narrowing continues at a slower pace.

Figure 11 explores the price component of the change in the rent gap during 1980-1990 and 1990-2000 by providing an industry-by-industry breakdown. Similarly, Figure 12 explores the shares component and Figure 13 explores the interaction component.

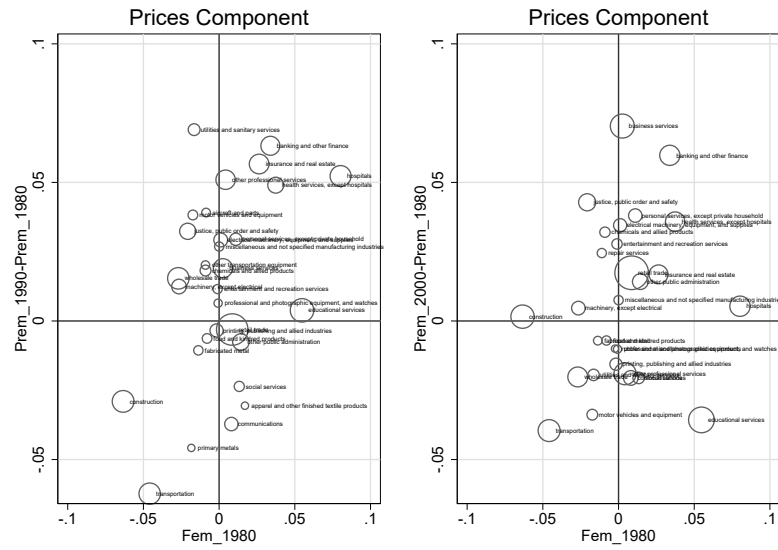


Figure 11: The Price Component

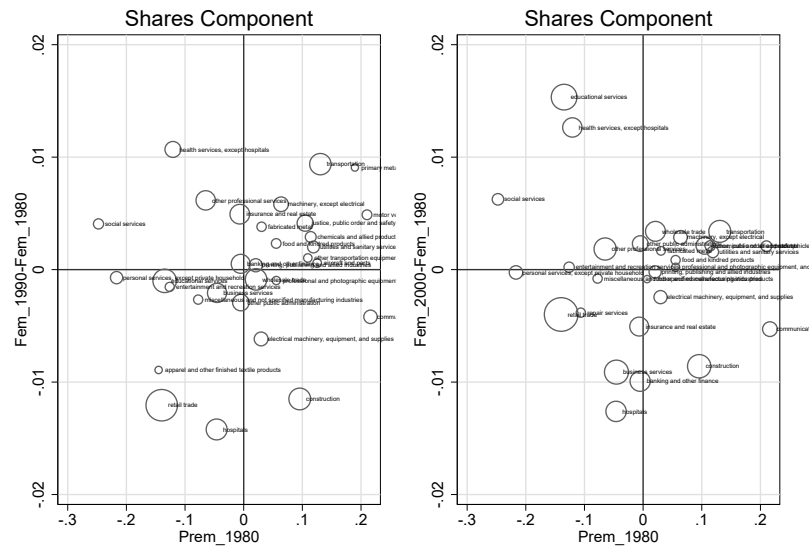


Figure 12: The Shares Component

E.2 Good Jobs and Bad Jobs

Industries classified as being in the ‘good’ sector are indicated with an asterisk. These are defined as industries with a 1980 industry premium above the (employment-weighted) median. In 1980, 49% of workers are in the good sector, and this falls to 43% in 1990, to 38% in 2000, and to 33% in 2010.

<i>Industry</i>	<i>Prem. 1980</i>	<i>Prem. 1990</i>	<i>Prem. 2000</i>	<i>Prem. 2010</i>	<i>Emp. 1980</i>	<i>Emp. 1990</i>	<i>Emp. 2000</i>	<i>Emp. 2010</i>
private household services	-.448	-.428	-.325	-.24	.72	1.08	.66	.95
social services	-.247	-.271	-.292	-.292	2.93	5.26	7.08	10.7
personal services, except private household	-.217	-.187	-.149	-.152	4.47	7.39	9.25	10.97
leather and leather products	-.15	-.142	-.011	.078	.62	.36	.26	.2
apparel and other finished textile products	-.145	-.176	-.192	-.197	3.43	2.98	2.19	1.38
retail trade	-.14	-.143	-.125	-.148	32.49	47.96	57.58	65
educational services	-.135	-.131	-.167	-.195	21.01	26.81	33.57	41.35
entertainment and recreation services	-.126	-.115	-.087	-.098	2.17	4.26	5.57	5.76
health services, except hospitals	-.121	-.072	-.036	-.006	6.32	12.4	19.04	26.5
toys, amusements, and sporting goods	-.116	-.009	-.011	.142	.42	.41	.5	.48
repair services	-.106	-.085	-.061	-.077	2.62	3.97	5.73	5.27
furniture and fixtures	-.098	-.053	-.045	-.05	1.61	1.89	2.17	1.37
miscellaneous and not specified manufacturing industries	-.078	-.051	-.043	-.052	2.55	4.29	4.57	3.3
textile mill products	-.066	-.055	-.047	-.063	2.62	2.21	1.33	.61
other professional services	-.065	-.014	-.033	.018	10.66	18.13	24.96	26.9
hospitals	-.046	.006	.012	.074	16.25	21.19	21	24.94
business services	-.045	-.026	.044	.045	9.22	17.82	29.29	33.16
lumber and wood products, except furniture	-.014	-.033	.002	-.008	1.06	1.18	1.2	1.03
insurance and real estate	-.007	.05	.067	.075	12.55	18.76	18.13	18.45
banking and other finance	-.005	.058	.118	.15	10.78	17.47	20.75	21.94
other public administration	-.005	-.012	.002	.05	12.27	13.19	12.64	12.25
rubber and miscellaneous plastics products	.007	*	.003	-.007	.003	2.32	2.55	3.07
printing, publishing and allied industries	.02	*	.017	.001	-.024	5.93	8.38	7.61
wholesale trade	.021	*	.036	.016	.042	16.28	22.41	20.03
electrical machinery, equipment, and supplies	.03	*	.059	.093	.13	9.94	9.4	8.91
fabricated metal	.03	*	.02	.013	.01	5.66	4.74	4.28
food and kindred products	.055	*	.049	.042	.031	5.02	4.69	4.5
professional and photographic equipment, and watches	.056	*	.062	.052	.114	3.39	3.46	3.53
paper and allied products	.057	*	.091	.073	.112	2.37	2.24	1.95
stone clay, glass, and concrete product	.059	*	.042	.038	.016	2.27	1.99	1.79
machinery, except electrical	.064	*	.076	.081	.096	11.45	10.73	9.77
construction	.096	*	.067	.068	.089	15.74	23.02	27.79
justice, public order and safety	.105	*	.137	.18	.243	10.22	12.79	14.99
other transportation equipment	.109	*	.13	.134	.182	3.01	3.96	2.35
chemicals and allied products	.114	*	.133	.165	.18	5.86	5.98	5.74
utilities and sanitary services	.119	*	.188	.169	.198	5.71	6.96	6.74
aircraft and parts	.125	*	.164	.161	.206	3.66	4.1	2.59
transportation	.131	*	.068	.029	.013	17.83	22.12	24.52
mining	.185	*	.213	.218	.329	2.3	1.92	1.17
primary metals	.19	*	.144	.122	.108	5.39	3.27	2.61
motor vehicles and equipment	.21	*	.248	.215	.142	5.83	5.19	6.35
communications	.216	*	.179	.158	.135	7.39	8.75	11.03
petroleum and coal products	.224	*	.253	.273	.396	.92	.9	.8
tobacco manufactures	.263	*	.47	.43	.44	.41	.33	.21
not specified metal industries					.13	.2	.12	.08

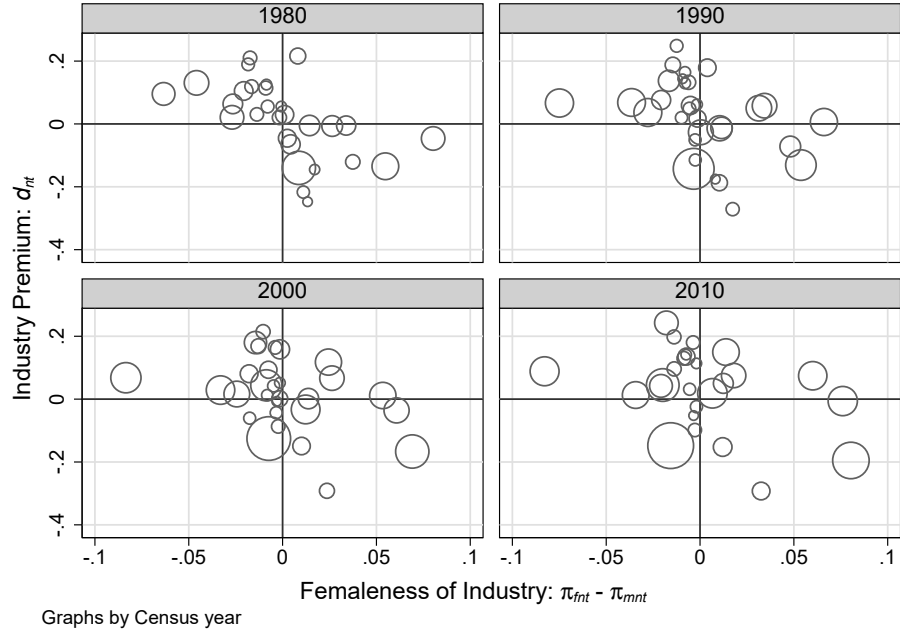


Figure 14: Male Jobs are Good Jobs

E.3 ‘Average Wages’ Specification

As mentioned in the text, we can estimate (10) directly. Table E1 reports the results of this exercise, using Z_{gct} as an instrument for \bar{w}_{gc} .

Table E1: Estimates: Year 1980-2000

	OLS		2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)
R_{gct}	0.72** (0.050)	0.70** (0.043)	0.45** (0.13)	0.62** (0.082)	0.59** (0.071)	0.58** (0.073)
u_{gct}	-0.0024 (0.0065)	-0.00053 (0.0058)	-0.0064 (0.0072)	-0.0048 (0.0059)	-0.0018 (0.0060)	-0.0013 (0.0060)
Observations	16476	16476	16476	16476	16476	16476
R^2	0.991	0.991	0.991	0.991	0.991	0.991
Fixed Effects:						
City \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times Year	Yes	Yes	Yes	Yes	Yes	Yes
Female \times Ind. \times City	Yes	Yes	Yes	Yes	Yes	Yes
Controls						
Education		Yes		Yes	Yes	Yes
Size		Yes			Yes	Yes
Participation		Yes				Yes
First-Stage:						
Z_{gct}			3.55** (0.67)	4.27** (0.67)	4.28** (0.64)	4.18** (0.68)
F -Stat.			28.30	40.68	44.44	37.75
p -val			0.00	0.00	0.00	0.00

Notes: Standard errors, in parentheses, are clustered at the city level. (**) and (*) denotes significance at the 5% and 10% level, respectively. All models estimated on a sample of 100 large U.S cities using the Census 1980-2000. The dependent variable is the regression adjusted city-industry-gender log wage. Each regression is weighted by the size of the city-industry-gender cell, and cells with less than 20 men and women are excluded. Columns 1-2 are estimated via Weighted Least Squares and Columns 3-6 are estimated via Two Stage Least Squares. The bottom panel of the table shows the results of the first-stage for the excluded variable of the 2SLS procedure for columns 3-6.

References

- Acemoglu, D. (2001, January). Good jobs versus bad jobs. *Journal of Labor Economics*, 1–22.
- Aizer, A. (2010). The gender wage gap and domestic violence. *The American Economic Review* 100(4), 1847–1859.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The Quarterly Journal of Economics* 115(3), 715–753.
- Albrecht, J. and S. Vroman (2002, February). A Matching Model with Endogenous Skill Requirements. *International Economic Review* 43(1), 283–305.
- Altonji, J. G. and R. M. Blank (1999, October). Race and gender in the labor market. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 3 of *Handbook of Labor Economics*, Chapter 48, pp. 3143–3259. Elsevier.
- Ashenfelter, O. and T. Hannan (1986). Sex discrimination and product market competition: The case of the banking industry. *The Quarterly Journal of Economics* 101(1), 149–173.
- Bar, M., S. Kim, and O. Leukhina (2015, October). Gender Wage Gap Accounting: The Role of Selection Bias. *Demography* 52(5), 1729–1750.
- Bayer, P., S. L. Ross, and G. Topa (2008, December). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy* 116(6), 1150–1196.
- Beaudry, P., D. A. Green, and B. Sand (2012). Does industrial composition matter for wages? a test of search and bargaining theory. *Econometrica* 80(3), 1063–1104.
- Beaudry, P. and E. Lewis (2014). Do male-female wage differentials reflect differences in the return to skill? cross-city evidence from 1980-2000. *American Economic Journal: Applied Economics* 6(2), 178–94.
- Bertrand, M. (2011). Chapter 17 - new perspectives on gender. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, Part B, pp. 1543 – 1590. Elsevier.
- Bertrand, M., E. Kamenica, and J. Pan (2015). Gender identity and relative income within households. *The Quarterly Journal of Economics* 130(2), 571–614.
- Black, S. E. and E. Brainerd (2004, July). Importing Equality? The Impact of Globalization on Gender Discrimination. *ILR Review* 57(4), 540–559.
- Black, S. E. and P. E. Strahan (2001, September). The division of spoils: Rent-sharing and discrimination in a regulated industry. *American Economic Review* 91(4), 814–831.
- Blau, F. D. and L. M. Kahn (1997, January). Swimming upstream: Trends in the gender wage differential in 1980s. *Journal of Labor Economics* 15(1), 1–42.
- Blau, F. D. and L. M. Kahn (2000). Gender differences in pay. *Journal of Economic Perspectives* 14(4), 75–99.

- Blau, F. D. and L. M. Kahn (2006, October). The u.s. gender pay gap in the 1990s: slowing convergence. *Industrial and Labor Relations Review* 60(1), 45–66.
- Blau, F. D. and L. M. Kahn (2016, January). The gender wage gap: Extent, trends, and explanations. Working Paper 21913, National Bureau of Economic Research.
- Borjas, G. J. and V. A. Ramey (1995, November). Foreign competition, market power, and wage inequality: Theory and practice. *Quarterly Journal of Economics*, 1075–110.
- Bowlus, A. J. (1997, October). A Search Interpretation of Male-Female Wage Differentials. *Journal of Labor Economics* 15(4), 625–57.
- Bowlus, A. J. and L. Grogan (2009, April). Gender wage differentials, job search, and part-time employment in the UK. *Oxford Economic Papers* 61(2), 275–303.
- Brashears, M. E. (2008). Gender and homophily: Differences in male and female association in blau space. *Social Science Research* 37(2), 400 – 415.
- Card, D. (2009, May). Immigration and inequality. *American Economic Review* 99(2), 1–21.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2016, November). Firms and labor market inequality: Evidence and some theory. Working Paper 22850, National Bureau of Economic Research.
- Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics forthcoming*.
- Chandrasekhar, A., V. Chernozhukov, F. Molinari, and P. Schrimpf (2012, December). Inference for best linear approximations to set identified functions. CeMMAP working papers CWP43/12, Centre for Microdata Methods and Practice, Institute for Fiscal Studies.
- Cortes, G. M., N. Jaimovich, and H. E. Siu (2016). The end of men and rise of women in the high-skilled labor market. Working paper, University of British Columbia.
- Dahl, G. B. (2002). Mobility and the return to education: Testing a roy model with multiple markets. *Econometrica* (6), 2367–2420.
- Deaton, A. and D. Lubotsky (2003, mar). Mortality, inequality and race in American cities and states. *Social science & medicine* 56(6), 1139–53.
- Flabbi, L. (2010, 08). Gender Discrimination Estimation In A Search Model With Matching And Bargaining. *International Economic Review* 51(3), 745–783.
- Flabbi, L. and A. Moro (2012). The effect of job flexibility on female labor market outcomes: Estimates from a search and bargaining model. *Journal of Econometrics* 168(1), 81 – 95. The Econometrics of Auctions and Games.
- Fortin, N., T. Lemieux, and S. Firpo (2011). *Decomposition Methods in Economics*, Volume 4 of *Handbook of Labor Economics*, Chapter 1, pp. 1–102. Elsevier.

- Fortin, N. M. and T. Lemieux (2000, May). Are women's wage gains men's losses? a distributional test. *American Economic Review* 90(2), 456–460.
- Goldin, C. (2014). A grand gender convergence: Its last chapter. *The American Economic Review* 104(4), 1091–1119.
- Green, D. A. (2015, November). Chasing after good jobs. do they exist and does it matter if they do? *Canadian Journal of Economics* 48, 1215–1265.
- Herrmann, M. and C. Machado (2012). Patterns of selection in labor market participation. In *11th IZA/SOLE Transatlantic Meeting of Labor Economists*.
- Lemieux, T. (2006). Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill? *The American Economic Review* 96(3), pp. 461–498.
- Machado, C. (2012). Selection, Heterogeneity and the Gender Wage Gap. Institute for the Study of Labor (IZA) Working Paper 7005, Institute for the Study of Labor (IZA).
- Manski, C. F. (1993, July). Identification of endogenous social effects: The reflection problem. *Review of Economic Studies* 60(3), 531–42.
- Mulligan, C. B. and Y. Rubinstein (2008). Selection, investment, and women's relative wages over time. *The Quarterly Journal of Economics* 123(3), 1061–1110.
- Munshi, K. (2003, May). Networks in the modern economy: Mexican migrants in the u.s. labor market. *The Quarterly Journal of Economics* 118(2), 549–599.
- Olivetti, C. and B. Petrongolo (2016, January). The Evolution of Gender Gaps in Industrialized Countries. IZA Discussion Papers 9659, Institute for the Study of Labor (IZA).
- Oreopoulos, P., T. von Wachter, and A. Heisz (2012, January). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics* 4(1), 1–29.
- Park, J. H. (1994, August). Estimation of sheepskin effects and returns to schooling using the old and the new cps measures of educational attainment. Papers 338, Princeton, Department of Economics - Industrial Relations Sections.
- Pissarides, C. (2000). *Equilibrium Unemployment Theory* (Second ed.). Cambridge, MA: MIT press.
- Ruggles, S., T. J. Alexander, K. Genadek, R. Goeken, and M. Schroeder, Matthew B. Sobek (2010). Integrated Public Use Microdata Series: Version 5.0 [Machine-readable database]. Technical report, University of Minnesota.
- Schirle, T. (2015). The gender wage gap in the canadian provinces, 1997-2014. working paper, Wilfrid Laurier University.
- Weinberg, B. A. (2000, January). Computer use and the demand for female workers. *Industrial and Labor Relations Review* 53(2), 290–308.
- Welch, F. (2000, May). Growth in women's relative wages and in inequality among men: One phenomenon or two? *American Economic Review* 90(2), 444–449.