

Sin, Isabelle; Stillman, Steven; Fabling, Richard

Working Paper

What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, and Discrimination

IZA Discussion Papers, No. 10975

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Sin, Isabelle; Stillman, Steven; Fabling, Richard (2017) : What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, and Discrimination, IZA Discussion Papers, No. 10975, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/170959>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 10975

**What Drives the Gender Wage Gap?
Examining the Roles of Sorting, Productivity
Differences, and Discrimination**

Isabelle Sin
Steven Stillman
Richard Fabling

AUGUST 2017

DISCUSSION PAPER SERIES

IZA DP No. 10975

What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, and Discrimination

Isabelle Sin

Motu Economic and Public Policy Research

Steven Stillman

Free University of Bozen-Bolzano and IZA

Richard Fabling

Independent Researcher

AUGUST 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

What Drives the Gender Wage Gap? Examining the Roles of Sorting, Productivity Differences, and Discrimination*

As in other OECD countries, women in New Zealand earn substantially less than men with similar observable characteristics. In this paper, we use a decade of annual wage and productivity data from New Zealand's Linked Employer-Employee Database to examine different explanations for this gender wage gap. Sorting by gender at either the industry or firm level explains less than one-fifth of the overall wage gap. Gender differences in productivity within firms also explain little of the difference seen in wages. The relationships between the gender wage-productivity gap and both age and tenure are inconsistent with statistical discrimination being an important explanatory factor for the remaining differences in wages. Relating across industry and over time variation in the gender wage-productivity gap to industry-year variation in worker skills, and product market and labor market competition, we find evidence that is consistent with taste discrimination being important for explaining the overall gender wage gap. Explanations based on gender differences in bargaining power are less consistent with our findings.

JEL Classification: J16, J31, J71

Keywords: gender wage gap, discrimination, sorting, productivity, competition

Corresponding author:

Steven Stillman
Department of Economics and Management
Free University of Bozen-Bolzano
Piazza Università 1
Bozen-Bolzano, 39100
Italy
E-mail: steven.stillman@unibz.it

* We thank Uwe Sunde, Michael Oberfichtner, Erich Battistin as well as seminar participants at Ludwig Maximilian University of Munich, University of Erlangen-Nuremberg, Aarhus University, Milan Labor Lunch Seminar, Motu Economic and Public Policy Research, NZ Pay Equity Commission, University of Canterbury, Victoria University Wellington, Auckland University of Technology, NZ State Services Commission, NZ Ministry for Women, NZ Productivity Commission and The Research Institute for the Evaluation of Public Policies (FBK-IRVAPP) and conference participants at the European Society for Population Economics Annual Conference. The results in this paper are not official statistics. They have been created for research purposes from the Integrated Data Infrastructure (IDI), managed by Statistics New Zealand. The opinions, findings, recommendations, and conclusions expressed in this paper are those of the author(s), not Statistics NZ, the Marsden Fund, Motu Economic and Public Policy Research, or the Free University of Bozen-Bolzano. Access to the anonymized data used in this study was provided by Statistics NZ under the security and confidentiality provisions of the Statistics Act 1975. Only people authorized by the Statistics Act 1975 are allowed to see data about a particular person, household, business, or organization, and the results in this paper have been confidentialized to protect these groups from identification and to keep their data safe. Careful consideration has been given to the privacy, security, and confidentiality issues associated with using administrative and survey data in the IDI. Further detail can be found in the Privacy impact assessment for the Integrated Data Infrastructure available from www.stats.govt.nz. The results are based in part on tax data supplied by In-land Revenue to Statistics NZ under the Tax Administration Act 1994. This tax data must be used only for statistical purposes, and no individual information may be published or disclosed in any other form, or provided to Inland Revenue for administrative or regulatory purposes. Any person who has had access to the unit record data has certified that they have been shown, have read, and have understood section 81 of the Tax Administration Act 1994, which relates to secrecy. Any discussion of data limitations or weaknesses is in the context of using the IDI for statistical purposes, and is not related to the data's ability to support Inland Revenue's core operational requirements.

1. Introduction

In the majority of OECD countries, women have made substantial progress in the labor market since the 1970s, with both wages and labor force participation increasing relative to that for men (Goldin 2014; Blau and Katz 2016).¹ Notwithstanding these large improvements, women still earn less than ‘comparable’ men in all OECD countries and, since the 2000s, progress for women in the labor market seems to have stalled (Kunze 2017; Olivetti and Petrongolo 2016; Blau and Kahn 2000). This has led to a resurgence in papers examining explanations for the remaining gender wage gap. Possible reasons can be broadly aggregated into four categories: i) productivity differences related to some combination of workforce commitment, fertility decisions, and social norms (Azmat and Ferrer 2016; Anderson et al. 2002; Angelov et al. 2016; Bertrand et al. 2015); ii) sorting of workers into different industries, occupations, and firms for similar reasons as above, as well as different preferences for competition (Gneezy et al. 2003; 2008); iii) differences in bargaining ability, especially in regards to rent sharing (Babcock and Laschever 2003; Card et al. 2016); and iv) discrimination either deriving from preferences (taste) or judgements about expected productivity (statistical).

In this paper, we use a decade of annual wage and productivity data from New Zealand’s Linked Employer-Employee Database (LEED) to evaluate the relative importance of each of the above explanations for the gender wage gap. We begin by examining the importance of gender differences in sorting between industries and firms in explaining the gender wage gap in a standard wage equation. In contrast to most previous related work, our data cover the universe of workers and we have an administrative match between workers and firms; hence measurement error is minimized. Next, we follow the methodology originated in Hellerstein et al. (1999) and Hellerstein and Neumark (1999) to jointly estimate translog firm production functions and wage bill equations that allow us to examine the relative marginal contribution of women to both firm output and wages. Comparing the relative productivity of women to their relative wages gives us a measure of whether women are paid less than men for work of equal value. This could come about because women are less good at bargaining for firm rents or because employers discriminate against women for either taste or statistical reasons.

¹ A large literature has focused on explaining this gender convergence and has highlighted a number of important contributing factors including increased education among women; increased selection of higher-skilled and older women in the labor force; changing social norms; increased control over fertility decision; and a shift in the economy towards more female-friendly service industries (Jacobsen et al. 2015; Olivetti and Petrongolo, 2016; Bailey et al. 2012).

We take two approaches to attempt to distinguish between these reasons. First, we examine how the relative wage-productivity gap varies by the employee's age and tenure with the firm. Canonical models of statistical discrimination show that the gap should get smaller with both age and tenure as women have more opportunities to reveal their true productivity (Aigner and Cain 1977). Second, we examine how the relative wage-productivity gap varies across industries and over time. Our data can be linked to an annual business survey that collects rich information on a variety of factors that should be related to both female bargaining power and the ability of firms to taste discriminate against female workers, in particular, the degree of competitiveness in both the product market that the firm faces and the labor market in which it hires, as well as the occupational, age, and skill distribution of workers in the industry. We use this data to examine how the relative wage-productivity gap at the industry-year level varies with these measures and relate our findings back to different models of taste discrimination and worker bargaining power.

We find that gender differences in sorting between either industries or firms explain less than one-fifth of the gender wage gap. Of the 17.5 percentage point gender wage gap that remains unexplained by sorting at most 2 percentage points is explained by gender productivity differences. In fact, with the caveat that it is difficult to accurately measure labor inputs on the intensive margin in our data, our preferred estimates suggest that women are statistically indistinguishable from men in terms of productivity, but have 18 percent lower relative wages. These results are robust to two approaches for dealing with the potential endogeneity of inputs in the production function. The estimated relative gender wage-productivity gap increases with both age and tenure, even though women's relative productivity is equal to that for men up to age 40 and does not vary by tenure. These results are inconsistent with simple models of statistical discrimination that argue that, as individuals reveal their true productivity, differences between wages and productivity should decline (Altonji and Pierret 2001).

Relating across-industry and over-time variation in the gender wage-productivity gap to industry-year variation in worker skills, and product market and labor market competition, we find that the gender wage-productivity gap is larger in industry-years with both more skilled workers and less product market competition. This is true even when controlling for industry fixed effects. These results are consistent with models of employer taste discrimination, especially those that allow for frictions in job search (Black 1995; Bowlus and Eckstein 2002; Flabbi 2010). To test the alternative explanation that gender differences in bargaining power

are important, we examine the interaction between labor market competitiveness, product market competitiveness, and workforce skill. Worker bargaining power should be more important for determining wages when it is difficult for firms to hire skilled workers; in contrast, it is cheaper for firms to discriminate when it is easier to hire workers. We find that, in less-competitive high-skilled industry-years, the gender wage-productivity gap is smaller when it is difficult for firms to hire workers. This is more consistent with taste discrimination being important for explaining the overall gender wage gap than with explanations based on gender differences in the willingness to bargain.

New Zealand is a small open economy with a gender wage gap similar to the US and many European countries (Dixon 2000; Kunze 2017). It was once a highly regulated economy but comprehensive market-oriented reforms were initiated in 1984. In less than a decade, the economy was opened to foreign capital and international trade, government assistance to industry was dramatically reduced, state-owned enterprises were privatized and the employer–employee bargaining process was decentralized (Evans et al. 1996). As a consequence, New Zealand now has a highly flexible labor market with low rates of unionization and centralized bargaining. Furthermore, female employment rates are high (around 80 percent) and nearly the same as men, on average, women are more educated than men, and the economy is strongly dominated by the service sector (Mercante and Mok 2014). Hence, in many ways, New Zealand can be seen as leader in the increased reliance on international trade and reduced employment in manufacturing and routine task-intensive activities that is currently happening in most OECD countries (Autor 2014). As discussed in Olivetti and Petrongolo (2016), these trends should favor women in the labor market and our results should be seen in that light.

Our paper makes a number of contributions to the literature. To our knowledge, we are the first to examine the relative importance of the four major explanations for why women earn less than men in one paper using one data source that covers the entire economy. Beyond that, we believe we are the first to use rich firm-level information to examine how the gender wage-productivity gap varies by industry and how this relates to important characteristics of those industries, and to use this data to examine different possible explanations for the overall gender wage gap. Importantly, we are also the first paper to examine the relative importance of sorting and discrimination in a country with flexible labor markets where few workers are covered by collective bargaining agreements, as is also the case in most Anglo-Saxon countries. Furthermore, our finding that taste discrimination is an important determinant of the gender

wage gap in a progressive country with high female employment, one of the most flexible labor markets in the world and a labor market that is strongest in sectors that favor women, suggests that this is also likely to be the case in other countries.

2. Data

We make use of two components of Statistics New Zealand's Integrated Data Infrastructure - the Longitudinal Business Database (LBD) and Linked Employer-Employee Data (LEED). LEED includes monthly earnings at each employer for all individuals paid a wage through the tax system since April 1999 (Carroll and Wood, 2003; Kelly 2003).² Individuals are identified by unique longitudinal identifiers derived from their tax numbers and are linked to their employers in the LBD, which has comprehensive information on each firm's financial performance and other business characteristics (Fabling 2009). The Longitudinal Business Frame forms a common backbone for both the LEED and LBD and provides physical locations for plants, as well as detailed industry information. Plants, workers and firms all have high quality longitudinal links once firm-level changes in legal form are accounted for, enabling the construction of a worker-firm panel dataset that covers all workers and firms (Fabling 2010).

We calculate average monthly earnings for each individual at the annual level starting with the year covered by April 2000 to March 2001 (the tax year in New Zealand) up to and including the year covered by April 2010 to March 2011.³ We calculate separate figures for each employer and in our analysis of individual wages we focus on the highest paying job in each year, as here we need to be able to match individual-year observations to one distinct employer. Limited demographic information on individuals is available for the entire sample period, specifically age, gender and location.

One significant limitation of the LEED data is that it contains no information on hours worked. Fabling and Maré (2015b) have created an algorithm that uses information on multiple jobs, minimum wage rates, job spells and notified job end dates to estimate relative labor input for each worker. It calculates an individual's full-time equivalent (FTE) labor supply at each employer by assuming that the statutory minimum wage is observed, multiple-job workers have

² LEED uses information from tax and statistical sources to construct a record of paid jobs. Each month all New Zealand employers file an Employer Monthly Schedule (EMS) record with Inland Revenue (IRD), which lists all employees at that firm in the month, the gross income they received and the tax that was deducted at source.

³ We exclude the first year of LEED data from all our estimates because the method used to derive firm capital stock requires a previous year of data. Most firm level information is collected for the tax year hence our use of this aggregation for individuals. We also use this data to create an uncensored measure of an individual's tenure with a firm.

the same total labor supply as single job workers and that hourly wage rates are likely to be constant over adjacent months at the start and end of jobs. This approach overestimates the true hours worked for a subset of workers, particularly part-time workers with high wages and only one job. We discuss below how we exploit this information in a variety of ways in our analyses in order to evaluate the robustness of our findings.

The LBD provides annual measures of firm-level (revenue-based) gross output, capital services and intermediate consumption, sourced from a mix of the Annual Enterprise Survey and tax-filed financial accounts data.⁴ This information, which is needed to estimate firm level production functions, is only available for for-profit firms in industries that are part of the ‘measured sector’, identified as “industries that mainly contain enterprises that are market producers. This means they sell their products for economically significant prices that affect the quantity that consumers are willing to purchase” (Statistics New Zealand, 2014).⁵ Among firms with employees, this includes 83% of firms covering 69% of overall employment. The main sectors excluded in terms of employment are government, education, healthcare and social services.

Our analysis uses three separate samples. The first is a 50 percent random sample of all worker-year observations in LEED in 2000/01 – 2010/11, excluding individuals who were ever a working proprietor during the sample window. As wage data for working proprietors may not accurately capture their labor input, it is not possible to identify their highest earning job or their FTE labor inputs. This sample includes 11.3 million observations on 3.1 million individuals. The second sample is a 50 percent random sample of all worker-year observations where an individual works for a firm with valid productivity data and at least five employees, providing a sample of 5.4 million observations on 2.1 million individuals. This second restriction is imposed because we identify the relative productivity of women by focusing on the share of female employment at each firm and how this relates to firm output. Hence, very

⁴ Gross output is measured as the value of sales of goods and services, less the value of purchases of goods for resale. Capital input is measured as the cost of capital services including depreciation costs; capital rental and leasing costs; and the user cost of capital. The inclusion of rental and leasing costs (including rates) ensures consistent treatment of capital input for firms that own their capital stock and firms that rent or lease their capital stock. The user cost of capital is calculated as the value of total assets, multiplied by an interest rate equal to 10 percent, to approximate the combined cost of interest and depreciation. Intermediate consumption is measured as the value of other inputs used up in the production process. Nominal measures of gross output and factor inputs are separately deflated using the Producer or Capital Goods Price Index, which is available separately for each industry grouping in the measured sector. See Fabling and Maré (2015a) for further details.

⁵ Excluded industries include ownership of owner-occupied dwellings, local and central government, education and training, and health care and social assistance. The sample is further restricted to exclude firms with missing or implausible production data, including those with unusually large changes (more than 400%) in any of gross output, total employment, capital services, or intermediate consumption.

small firms provide little variation to help identify the key parameters in the model. All our results are robust to instead restricting the sample to firms with at least ten employees and in our individual level analysis below we show that the gender wage gap, as well as the importance of sorting into different industries and firms, is similar in the two samples. Our third sample is a firm-year level data set that includes all firm-years where firms have valid productivity data and at least five employees. This sample, which is used to estimate production functions, provides 290 thousand observations on 67 thousand firms. All regressions run using firm data are weighted by the number of employees at each firm so that the result still reflect the situation for the average worker and are comparable to those from the individual level regressions.

Table 1 presents descriptive statistics for all three samples. Reflecting the high employment rate of women noted above, the overall sample has a perfect gender balance. New Zealand has a fairly young workforce by OECD standards with 25 percent of employees less than age 25 and 34 percent between 25 and 39. Based on the approximation of hours developed by Fabling and Maré (2015b), the average female works 0.725 FTE and the average male 0.816 FTE. The average male earns 3,655 and the average female 2,437 per month in nominal dollars, which is a gender wage gap of 33 percent unadjusted and 25 percent adjusting for gender difference in average FTEs. The average individual works for a firm with 93 employees and 47% are employed at firms with multiple plants.

Individuals working at firms with at least five employees and valid productivity data are more likely to be young and male. This reflects that the unmeasured sectors, especially government, education and health, employ more women and typically have higher returns to tenure. Unsurprisingly, firms are also on average larger in this sample and workers are more likely to work at firms that have multiple plants. The unadjusted earning gap is slightly larger in this sample at 37 percent with no adjustment and 28 percent adjusting for average FTEs.

Finally, the third panel shows the key variables used for estimating firm level production functions. The statistics here are weighted for firm size so reflect the situation for the average worker in our data. The average firm in this sample employees 32 workers, but the average individual works at a firm with 1,280 employees. Most individuals work at firms that do not have a working proprietor, but a small proportion of firms have many. We control for the number of working proprietors in our joint analysis of firm output and wage bill.

3. Sorting and the Gender Wage Gap

We begin by using the full sample of workers to examine the average difference in earnings between all men and women. This provides a measure of the gender pay gap comparable to that obtained using survey data. We next examine what proportion of the gender pay gap is explained by the sorting of workers into different firms. We do this by examining differences in earnings between men and women working in the same industries and then at the same firms (Heinze and Wolf 2010; Bayard et al. 2003).

We start by estimating the following regression model:

$$\text{LnEarn}_{ijt} = \alpha_t + \delta * \text{Female}_i + X_{ijt}\beta + \mu_{ijt}, \quad (1)$$

where LnEarn_{ijt} is the (log) monthly earnings for individual i who is employed at firm j at time t , Female_i is a dummy variable indicating whether or not an individual is female, X_{ijt} are dummies for an individual's age-group (<25, 25-39, 40-54, >54) and their log(FTE) in this job, α_t are year fixed effects, and μ_{ijt} is a stochastic error term.⁶ We include an individual's highest paying job in a particular year only. The coefficient of interest in this model is δ , which measures the log earnings gap for the average female worker. We control for log(FTE) instead of calculating a pseudo-hourly wage (e.g. monthly earnings * FTE / 167 hours worked per month) to allow a more flexible relationship between labor input and earnings, which we think could be especially important since the accuracy of our FTE measure is likely correlated with both age and gender. While we do not have information on education, previous work using survey data has found that gender differences in education are generally unimportant for explaining gender wage differences in New Zealand (Dixon 2000).

We then expand this model, first by adding four-digit ANZSIC industry fixed effects (this is the most disaggregated level available and includes around 500 distinct industries, e.g. clothing and footwear repair) and then firm/year fixed effects. In this final specification, the gender wage gap measured by δ reflects the average difference in wages for men and women with similar characteristics employed at the same firm in a particular year. This captures many of the unobserved differences in worker skills that are often excluded from traditional regression models of the gender wage gap. Comparing δ in these specifications to that estimated in the

⁶ Our results are unaffected when controlling for a quadratic in age as is more standard in a wage regression. We use this specification instead because discrete age groups are used when allowing for heterogeneity in labor inputs when we estimate firm production functions.

initial regression model reveals how much of the gender wage gap is explained by gender differences in the sorting of workers across industries and firms.

Table 2 presents the results from estimating these regression models. Standard errors are clustered at the individual level in each specification allowing for arbitrary correlation in earnings over time for each individual. Given our use of a 50 percent sample of the population, our estimates are extremely precise and all coefficients are significantly different from zero at the 1 percent level. We find that earnings increase with age up until age 40-54 and then decline. Turning to the main coefficient of interest, our baseline specification indicates that, controlling for only age and each worker's FTE, women earn 19.9 percent less than men on average.⁷

In the second column, we add controls for the four-digit industry associated with an individual's job. This reduces the unexplained gender wage gap to 18.8 percent; gender differences in industry sorting explain only 6.8 percent of the overall gender wage gap. This is lower than previous estimates in the literature, e.g. the US in 1990 (11.3 percent; Bayard et al. 2003) or New Zealand in 1997-98 (16.5 percent; Dixon 2010), which likely reflects the increasing flexibility of labor markets over time in Anglo countries. Next, we replace the industry controls with firm-year fixed effects. This further reduces the unexplained gender wage gap to 17.7 percent, with 12.2 percent of the overall wage gap explained by gender differences in the sorting of workers into different firms and the remaining attributable to gender differences in pay for workers at the same firms. This is also a lower estimate than that found in the literature, e.g. Portugal in 2002-09 (14.9-19.9 percent; Card et al. 2016) or the US in 1990 (15.6 percent; Bayard et al. 2003). Again, this likely reflects that more flexible labor market found in New Zealand in the 2000s, which allows for larger within-firm wage differences and also has opened up a wider variety of jobs for women.

Specifications (4) through (6) repeat the previous analysis for the subsample of individuals whose highest paying job is at a firm that employs at least five people and has valid productivity data. Even though women are a smaller share of employees in these firms and sectors of the economy that are considered to be especially female friendly (e.g. government, education, and health) are dropped from the analysis, the overall gender wage gap is quite similar, increasing only from 19.9 to 20.9 percent. In this sample, sorting explains a slightly larger share of the

⁷ Dixon (2000) estimated the log gender wage gap in New Zealand to be 17 percentage points using survey data from 1997-98. We get a similar estimate when we use a subset of the same survey data that can be linked to our administrative records. Hence, it appears that not being able to fully control for gender differences in hours worked leads us to overstate the overall gender wage gap by around 5 percentage points. However, our estimates of the importance of industry sorting are the same if we use the subset of data where we can control for hours worked.

overall gender wage gap, with industry sorting and firm sorting now explaining 13.6 and 18.3 percent of the gap, respectively. However, sorting still explains less than one-fifth of the overall gender wage gap.

4. Estimating Gender Difference in Productivity

Econometric Model

The estimation approach used in the previous section is informative about the extent to which men are paid more because they sort into higher-paying industries and firms. However, it does not allow us to account for within-firm differences in marginal productivity. Such differences could result from men and women performing different roles within the firm, expending different amounts of effort, or various other non-discriminatory reasons, such as women having less effective work experience. In this section, we follow the approach pioneered in Hellerstein et al. (1999) and Hellerstein and Neumark (1999) and jointly estimate augmented versions of translog production and firm wage bill functions that allow us examine how much of the within-firm gender wage gap is explained by gender differences in productivity.

A firm's gross output (Y) is modelled as a second-order translog approximation to an arbitrary production function that combines intermediate inputs (M), capital (K) and effective labor (\tilde{L}):

$$\ln Y = \alpha_1^Y \ln M + \alpha_2^Y \ln K + \alpha_3^Y \ln \tilde{L} + \alpha_4^Y (\ln M)^2 + \alpha_5^Y (\ln K)^2 + \alpha_6^Y (\ln \tilde{L})^2 + \alpha_7^Y (\ln M)(\ln K) + \alpha_8^Y (\ln M)(\ln \tilde{L}) + \alpha_9^Y (\ln K)(\ln \tilde{L}) + X\delta^Y + \varepsilon^Y \quad (2)$$

where X includes two-digit industry and year fixed effects, as well as controls for the number of working proprietors at the firm and a dummy for the firm having multiple plants, allowing for overall output to vary by these characteristics. Effective labor is modelled as a function of the quantity of labor supplied by workers differentiated into four age groups (<25, 25-39, 40-54, >54) and by gender. Men and women are assumed to supply different amounts of perfectly substitutable labor input, as are workers in different age groups. The marginal product of each type of workers is measured relative to a common base group, e.g. the relative productivity of

women is $\phi_f^Y = \frac{\partial Y / \partial L_f}{\partial Y / \partial L_m}$ where f indexes female labor and m indexes male labor.

Hence, (log) effective labor can be written as:

$$\ln \tilde{L} = \ln L + \ln(1 + (\phi_f^Y - 1) * s_{-f}) + \ln(1 + (\phi_{<25}^Y - 1) * s_{-lt25} + (\phi_{40-54}^Y - 1) * s_{-40-54} + (\phi_{>54}^Y - 1) * s_{-gt54}) \quad (3)$$

where, for each firm/year, s_f = share of employment that is female, s_{lt25} = share of employment that is age <25, s_{40-54} = share of employment that is age 40-54, s_{gt64} = share of employment that is age > 64, and each ϕ represents the marginal productivity of women relative to men or individuals of different age groups relative to 25 to 39 year-old workers. Here, the relative marginal products of women in each age group compared with men in the same group are restricted to be equal. This allows us to estimate the proportion of the overall gender wage gap that is explained by differences in productivity. We subsequently relax this assumption and allow gender differences in wages and productivity to vary by age and alternatively by tenure.

The marginal contribution of each type of worker to the total firm wage bill (W) can be estimated from an analogous model where all coefficients are allowed to differ from those estimated in the production function:

$$\ln W = \alpha_1^W \ln M + \alpha_2^W \ln K + \alpha_3^W \ln \tilde{L} + \alpha_4^W (\ln M)^2 + \alpha_5^W (\ln K)^2 + \alpha_6^W (\ln \tilde{L})^2 + \alpha_7^W (\ln M)(\ln K) + \alpha_8^W (\ln M)(\ln \tilde{L}) + \alpha_9^W (\ln K)(\ln \tilde{L}) + X\delta^W + \varepsilon^W \quad (4)$$

where
$$\ln \tilde{L} = \ln L + \ln(1 + (\phi_f^W - 1) * s_f) + \ln(1 + (\phi_{<25}^W - 1) * s_{lt25} + (\phi_{40-54}^W - 1) * s_{40-54} + (\phi_{>54}^W - 1) * s_{gt54}) \quad (5)$$

This model can be thought of as a flexible aggregation of individual wage regressions where the relative wages of different types of workers are allowed to depend on detailed firm structure.

Main Results

Table 3 presents the results from substituting (3) into (2) and (5) into (4) and jointly estimating the resulting equations using non-linear seemingly unrelated regressions (NLSUR). The main production function parameters estimated in a translog specification are non-invariant to the units of the data (Hellerstein et al. 1999). While there are normalization approaches that can be used to estimate the returns to scale parameters, since these are not the focus of our paper, we instead only present our estimates of the key ϕ parameters which measure the relative productivity (Panel A) of female workers and workers in different age-groups as well their corresponding relative wages (Panel B). Note that the way \tilde{L} is defined, productivity and wage differentials between groups are indicated by the estimate of the relevant ϕ being significantly different from one (rather than zero). For example, $\phi_f^Y = 0.69$ in the first row of the first specification implies that female workers are 31% less productive than male workers.

We also present in Panel C the percent difference in the relative contribution of different groups of workers to firm's wages and productivity, and test whether this is significantly different from zero. A positive coefficient here indicates that a type of worker is paid relatively less than their marginal product of labor and a negative coefficient indicates the opposite. Standard errors are clustered at the firm level in each regression allowing for arbitrary correlation in output and wages over time for each firm, and all regressions are weighted by the number of employees at each firm so that the results are representative for all workers in the measured sector working for firms with at least five employees.

One important issue we face in estimating both the production function and wage bill equation is how to account for gender and age difference in labor inputs on the intensive margin. Recall from the previous section that we have only a rough measure of hours worked for different workers and that this will be overstated for particular types of workers, e.g. those working part-time at one non-low paid job. We suspect that this is more likely to be the case for female than male workers, and for younger and older workers. If all workers are paid an equal percentage of their marginal product, then it should not matter how we measure labor inputs, since both relative wages and relative productivity will be underestimated by the same proportion (because workers are assumed to be working more than they are actually). However, if part-time workers are treated differently in the wage setting process, as we might expect, then the results could be sensitive to how we measure labor inputs.

For this reason, we initially take four different approaches for measuring labor inputs and judge whether our results are robust to the chosen method. In our first specification, we take the simplest approach and use a head-count measure of labor with all workers counting equally. In the second specification, we instead measure labor and labor shares by summing worker FTEs as measured using the Fabling and Maré (2015b) algorithm. In the third specification, instead of summing workers FTEs, we use the derived FTE measure to identify whether employees are likely to work part time (FTE<1) or full time. We then include a separate ϕ for part-time workers relative to full-time workers of the same gender and age group, which does not differ by gender or age group. In our final specification, instead of using the FTE measure to classify individual workers, we calculate the average FTE of male and female workers in each firm/year and estimate an adjustment factor that modifies labor inputs to account for gender differences in average FTEs. Specifically, the labor input of a particular type of worker i is calculated as $\tilde{L}_i = \# \text{Workers}_i * \text{AvgFTE}_i^\delta$, with δ invariant across types. The values of δ in the production function and wage bill equation are estimated as part of the NLSUR estimation. This final

approach is the most flexible as it allows the data to speak to how the estimated FTEs translate to different levels of output and wages.

Turning to the results, when labor inputs are measured only based on the number of employees at a firm, we find that women are, on average, 31% less productive than men, young and older workers are 80-82% less productive than 25-39 year-olds and 40-54 year-olds are 22% less productive than 25-39 year-olds. In each case, these productivity differences are significant. Wage differences follow the same pattern for workers. Comparing productivity differences to wage differences, we find that women are paid 11.6 percent less than men for the same contribution to firm output and older workers are paid more than double relative to their contribution. For the other age groups, relative wages are not significantly different to relative contributions to productivity.

The remaining results indicate that there are clear differences in how full-time and part-time workers are compensated and that this is important for evaluating the relative productivity of women, young and older workers, which are all groups with a high propensity to engage in part-time work. All of our methods for adjusting labor inputs to account for differences along the intensive margin of work produce remarkably similar results. In each case, the relative contribution of women to firm output is not significantly different to that for men. The age-productivity gradient is also much less steep once we control for intensive margin differences. Overall, we find the same pattern for relative productivity-wage differences, but now we estimate that women are paid 15.7 to 16.9 percent less than men relative to their marginal product, and older workers 41.5 to 51.6 percent more than prime-aged workers relative to their marginal product. Again, none of the other age differences are significant. These results imply that of the 17.5 percentage point gender wage gap that remains unexplained by the sorting of workers into firms at most 2 percentage points is explained by gender productivity differences.

It is important to note that the method used here identifies relative differences in the productivity of different types of workers by relating cross-sectional variation in the demographic characteristics (age and gender) of workers to cross-sectional variation in output across firms. Controlling for industry fixed effects means that these comparisons are only made within industry and using a translog production function allows for flexible correlations between the demographic characteristics of the firm's workforce and the firm's production technology embedded in their use of intermediate goods and capital. Hence, the key assumption is that, within 2-digit industries, the gender mix of employees is not correlated with the ratio

of the average female wage to the average male wage, and that, equivalently, high-productivity and low-productivity workers do not sort into firms with particular gender ratios.

That said, a number of our findings suggest that issue is unlikely to be important. First, the fact that we find that across-firm sorting explains at most one-fifth of the overall gender wage gap suggests that across-firm sorting based on productivity differences is also unlikely to be an important phenomena. Second, tightly related to this finding, our finding that the relative contribution of female workers to the overall wage bill is similar to the within-firm gender wage gap also suggests that across-firm sorting based on wages is unimportant. Finally, our finding that, controlling for differences in intensive margin labor inputs, women are equally productive to men, also suggests that across firm sorting based on productivity differences is unlikely to bias our estimates of how relative wages compare to relative marginal products for women.

Robustness

In Table 4, we examine the robustness of our results to a number of sampling and modelling decisions and, in particular, to the potential endogeneity of inputs. While all three specifications above that adjust labor inputs on the intensive margin produce similar results, we focus on the final specification, which uses a data driven approach to estimate how FTEs translate to output and wages, here and in all our further robustness checks and extensions. This approach is the most flexible and also estimates the relative contribution of women to the firm wage bill to be nearly the same as the gender wage gap estimated using individual wage regressions, which provides a type of external validity for the approach.

Here, we only present the main results for gender differences with the age results available in Appendix Table 1. In the first column, we repeat the results from our preferred specification from Table 3 for ease of comparison. In the second column, we re-estimate this specification omitting the two-digit ANZSIC industry fixed effects. Perhaps surprisingly, women appear to sort into more productive industries that also have higher wages but larger gender differences in wages relative to marginal productivity. Hence, failing to control for industry fixed effects leads to an 8 percent increase in the gender wage-productivity gap. In the third column, we drop firms with more than one working proprietor. These individuals contribute to firm output but not to the wage bill so including firms with large numbers of working proprietors might bias the results if this is also correlated with gender composition. Dropping nearly one-third of the firms from our sample has little impact on the results.

Another issue with estimating production functions is the potential endogeneity of inputs. In our application, this will only bias our estimates of relative wage-productivity differences if firm demographic characteristics are correlated with production decisions. We first examine whether the potential endogeneity of intermediate inputs affects our estimates. We do this by transforming our production function into a value-added model where the outcomes is now $\ln(Y - M)$ and intermediate inputs are no longer included is an independent variable. This sidesteps the endogeneity issue by avoiding estimation of a coefficient for this variable.⁸ As can be seen in the fourth column of Table 4, the results using this approach are nearly identical to our original estimates.

Next, we use the control function approach introduced by Olley and Pakes (1996) to account for the potential endogeneity of both intermediate inputs and capital. In particular, we extend the production function estimated above to include a third degree polynomial in capital and investment (capital next period minus 0.9 of capital this period). As described in Olley and Pakes, under the assumption that labor is a non-dynamic input and productivity shocks follow a first-order Markov process, this approach controls for any unobserved shocks to firm productivity that are correlated with both the dynamic inputs (capital and intermediate inputs) and output. One issue with this approach is that investment calculated from the data can be zero or negative, which in theory should not be possible. Hence, we estimate two versions of this model, one with the full sample of firm-years for which we can calculate investment and a second for the subsample of firms with positive investment. In both cases, our estimates for the relative productivity of women are nearly identical to our original estimate.

Discussion

Our main finding from estimating firm level production functions and wage bill equations is that of the 17.5 percentage point gender wage gap that remains unexplained by the sorting of workers into different firms at most 2 percentage points is explained by gender productivity differences. As discussed in Hellerstein et al. (1999), this result is inconsistent with the assumption that the labor market is characterized by profit-maximizing or cost-minimizing firms in a competitive spot labor market. Most previous papers using this approach have argued that any differences here are indicative of gender discrimination. A second possibility, as

⁸ As discussed in Hellerstein et al. (1999), there are other potential benefits of a value-added specification. For example, net output is typically more comparable across industries and firms that differ in their degree of vertical integration. Interestingly, this specification can be derived from quite polar assumptions: either the elasticity of substitution between intermediate inputs and value added is zero or it is infinite.

discussed in Card et al. (2016), is that a subset of firms are profitable and share rents with their employees and that women are not as good as men at getting a share of these rents (which arguably could also be because of discrimination).

In the remaining sections of the paper, we take two approaches to attempt to distinguish between these different explanations for the estimated gender wage-productivity gap, as well as to evaluate the relative importance of statistical versus taste discrimination. In the next section, we examine how the relative wage-productivity gap varies by age and tenure within the firm, and relate this to models of statistical discrimination. In the following section, we examine how the relative wage-productivity gap varies across industries and over time. We then link this information to data on a variety of factors that should be related to both female bargaining power and the ability of firms to taste discriminate against female workers and evaluate different explanations for the relative wage-productivity gap.

5. Heterogeneity in Gender Difference in Productivity

In Table 5, we present results from estimating extended versions of the production functions and wage bill equations discussed in the previous section. In particular, we first extend each to allow for an interaction between gender and age in the augmented labor input function. In particular, we now allow the productivity of workers, as well as their contribution to wages, to vary along both dimensions simultaneously. The estimated ϕ s are now relative to the base group of 25-39 year-old men. However, the results as presented show the age-productivity gradient for men and then the relative productivity of women in each age group. The model is the same in all other dimensions as what we previously estimated.

Women are found to be equally productive as similarly aged men up to age 40. In the 40-54 year-olds group, women are substantially less productive than men. This could reflect some combination of the selection of women who remain in the workforce, skill depreciation that occurs after the main childbearing years, and birth cohort effects. It is outside the scope of this paper, but examining this more closely would be an interesting extension. More relevant for our paper, we find that the gender difference in wages relative to productivity increases with age. We find no evidence of a gender wage-productivity gap for young women, a 16 percent gap for women aged 25-39, a 21 percent gap for those aged 40-54 and a 49 percent gap for older women.

We also extend our analysis to instead allow gender differences in wages relative to productivity to vary by an individual's tenure with the firm. We define three tenure groups;

individuals who have been at the firm for less than one year (23% of employees), those that have been at the firm for one to two years (54% of employees) and those who have been at the firm for two or more years (23% of employees).⁹ New Zealand has a generally high turnover of employees, so this analysis is best for examining productivity and wage growth early in one's job. The estimated ϕ s are now relative to the base group of men who have been at the firm for less than one year. However, the results as presented show the tenure-productivity gradient for men and then the relative productivity of women in each tenure group. Again, the model is the same in all other dimensions as what we previously estimated and we do not interact age and tenure or age and gender.

For both genders, there is suggestive evidence that productivity increases with tenure, but for women the tenure-productivity profile is much steeper. Men's wages also increase with tenure, but women's decline in relative terms. Taken together, while there is no initial gender wage-productivity gap, women who remain at the same firm for more than one year end up getting paid over 20 percent less than men for the same marginal product.

Statistical discrimination arises when workers are unable to signal their true productivity and employers use aggregate information to estimate expected productivity. Even if there are no differences in average productivity between groups of workers, if productivity varies more among one particular group, then it will make sense for employers to pay lower wages to workers in that group (Aigner and Cain 1977). One might expect this to be the case for women given their historical lower levels of commitment. However, as workers spend more time in the labor market and with a particular firm, they increase their ability to signal their true productivity through both direct observation, and a better resume and references. Hence, if statistical discrimination is important for explaining differences between the wages and productivity of a particular group of workers, then this gap should get smaller with both age and tenure (Altonji and Pierret 2001).

In fact, we find the opposite with the gender wage-productivity gap increasing monotonically with both age and tenure. This finding taken together with the fact that productivity does not seem to decline with age (except among women aged 40-54) and tenure, suggests that statistical discrimination is not important for explaining why women in New Zealand are paid less relative to their marginal product than are men.

⁹ We define tenure based on the number of months employed at a firm as is standard in the literature. In the time period we focus on, New Zealand did not have paid parental leave so months at the firm should represent a similar commitment for male and female workers.

6. Wage Gap-Productivity Differences Across Industries

In our final analysis, we examine how the relative gender wage-productivity gap varies across industries and over time. This is done by estimating equations (2) and (4) separately by industry and then by industry-year. Figure 1 shows how our estimates of ϕ vary across 37 industries.¹⁰ Here, we pool all years for each industry to more clearly illustrate the variation in our data. In general, there is much less variation in ϕ over time within industries than there is across industries. For each industry, we present ϕ_f^Y on the y-axis and ϕ_f^W on the x-axis. Industries fall into one of four quadrants based on whether each ϕ is above or below 1.

There are no industries where women are both relatively more productive and paid more than men. There are two industries, Agricultural Support Services (AA32) and Building Construction (EE11) where women are paid relatively more than men but are relatively less productive. This likely reflects the fact that women in these industries are almost entirely in administrative jobs. The remaining industries are split roughly into one-third where women are relatively more productive than men but paid relatively less and two-third where women are both relatively less productive and paid less than men. Finally, the 45-degree line splits the sample into around one-third of industries where women are paid more than or equal to their marginal product and around two-thirds where women are paid less than it.

Figure 2 presents for each industry the estimated difference between the relative marginal product of women and their relative contribution to the wage bill. As in the previous tables, a positive estimate means that women are paid less than men relative to their marginal product and a negative estimate means the opposite. Ninety-five percent confidence intervals for each estimate are also presented. There is only one industry, Other Services, where women are paid relatively more than their relative marginal product and the difference is significantly different from zero. On the other hand, there are twelve industries where women are paid relatively less than their marginal product and the gap is over 40 percent in five of these: Finance and Insurance, Telecommunications, Transport Equipment Manufacturing, Electricity, Gas and

¹⁰ There are 50 3-digit NZSIOC industries in the measured sector of the economy (these are consolidated 3-digit ANZSIC industries which are aggregated to match with macroeconomic data). Examples include Wood Product Manufacturing and Finance and Insurance. We aggregate these further up to 39 industries by combining some smaller industries (in terms of employment) within the same one-digit category. This figure omits two industries with very imprecisely omitted estimates. From the following figure we drop the five industries with the least precisely estimated results. When estimating the model at the industry-year level, we drop all the industry-years with fewer than 40 observations (all of which come from the same industry), where our non-linear SUR model struggles to converge.

Water, and Rail, Water and Air Transport. It is worth noting that these are all sectors that are typically thought to be non-competitive with large amounts of rent sharing.

We next model this variation in the gender wage-productivity gap across industry-years. In particular, we use the data in LEED along with linked annual household survey data (the Household Labor Force Survey) and annual business survey data (the Business Operations Survey - BOS) to measure different factors that should be related to both female bargaining power and the ability of firms to taste discriminate against female workers.¹¹ These surveys are both representative surveys and the BOS is specifically designed to be representative at the industry level. Unfortunately, one limitation of the BOS is that the survey only started in 2005 and hence our remaining analysis is limited to the 2005-2011 period. However, all of our main results are qualitatively the same for this reduced time period.

Starting with Becker (1957), the argument has been made that taste discrimination cannot persist in a perfectly competitive product market because firms that discriminate will lose money compared to those that do not and will be driven out of the market. This has led a number of papers to focus on the relationship between product market competition and discrimination.¹² An alternative strand of the literature has focused on the importance of competition in the labor market, with the intuition that employers forgo profit whenever a vacancy remains unfilled and hence will find it more expensive to discriminate if labor markets are tight.¹³ A final related literature is that on rent sharing. An important insight from this literature is that rent sharing is most likely to occur where workers have specialized skills or training and search is costly. Quoting from McLaughlin (1994), “Unskilled workers are not excluded from the model, but rents for these workers are not predicted to be quantitatively important. The analysis points to managers, professionals, scientists, and skilled technicians and craftsmen as the workers on jobs with potentially important rents.”¹⁴

¹¹ These surveys can both be directly linked to individuals and firms in our data. However, because they are only for representative samples of individuals and firms, and the BOS oversamples large firms, relying on the linked data would severely restrict our previous analyzes. We are also worried that firm level reports on areas like product market competitiveness and difficulty hiring might be endogenously determined with productivity shocks. This is much less of a concern when focusing on industry level variation.

¹² For example, Black and Strahan (2001) show that removing regulations on the banking sector in US led to a reduction in the gender wage gap in this industry while Hirsch et al. (2014) find that the gender wage gap is smaller at firms in Germany that face more product market competition.

¹³ This intuition is formalized in Biddle and Hamermesh (2013) which develops an equilibrium search model that shows that employers discriminate less in a tight labor market and is tested in recent work by Baert et al. (2015), who find that interview rates for immigrants in a correspondence study in Belgium are higher in occupations for which vacancies are difficult to fill.

¹⁴ There is limited empirical evidence for this assertion. For example, Macis and Schivardi (2016) find that increased exports brought about by currency depreciation lead to more rent sharing among experienced workers.

The literature on employee bargaining is less developed. Ellingsen and Rosén (2003) model the choice of whether to bargain with workers in a standard sequential search model with heterogeneous workers. They show that firms are more likely to bargain with higher skilled and more productive workers, and particularly in occupations and jobs with a high dispersion in worker's' skills. Bargaining is also more likely to occur in tight labor markets since in such market workers have higher reservation wages. This last insight is useful for distinguishing between taste discrimination and gender differences in workers' bargaining skills as explanations for industry differences in the gender wage-productivity gap, as taste discrimination should be more prevalent when it is easier to hire workers while bargaining should be more important when labor markets are tight.¹⁵

In order to test whether either taste discrimination or gender differences in bargaining explain across industry-year variation in the gender wage-productivity gap, we develop indexes that measure for each industry-year: i) employee skills; ii) product market competition and profitability; and iii) difficulty in hiring skilled workers. Our data contains a number of variables that proxy for these underlying concepts. Instead of specification mining to choose which variables to include in the regression, we use a data driven approach to develop an index for each area. Specifically, for each concept we run a principal component analysis across all industry/year combinations for a group of variables that fits into the category. For each, we then use the first principal component, normalized to mean zero and variance one, to measure the underlying concept.

Table 6 presents the results from this exercise. To measure employee skills, we include variables on the occupational distribution of workers in a particular industry-year, as well as the age distribution of workers and the fraction of employees in each of the four quartiles of the skill distribution based on the worker fixed effects estimated by Maré and Hyslop (2006). Overall, the first principal component explains 43 percent of the variation in the data and loads positively on high-skilled occupations, prime-age workers and workers in the top skill quartile and negatively on low-skilled occupations, young workers and workers in the low half of the skill distribution. We call this component "Higher Skill". To measure competition and profitability, we include the distribution of reported competitiveness in the product market in a particular industry-year, as well as the average capital to labor ratio. Here, the first principal component explains 48 percent of the variation in the data and loads positively on firms with

¹⁵ Supporting this model, Brenzel et al. (2014) examine wage bargaining in the Germany labor market and find that bargaining is more likely for more-educated job applicants and in tighter labor markets.

no competition or 1-2 competitors and on the capital to labor ratio, and negatively on firms with many competitors, some dominant. We call this component “Less Competition”. To measure difficulty in hiring workers, we include the distribution of reported difficulty in hiring managers and professionals and in hiring technicians and associate professionals. Here, the first principal component explains 48 percent of the variation in the data and loads positively on firms with severe and some difficulty hiring both type of employees and negatively on firms with no difficulty hiring either type. We call this component “More Difficulty Hiring”.

We now regress the gender wage-productivity gap in each industry-year on the three indexes created in the previous step, which measure worker skills, product market competition and labor market competition, along with a full set of interactions between these variables. We also control for year fixed effects to allow for any unrelated time-trends in the gender wage-productivity gap (there are none, so we do not present the coefficients from these variables). We drop ten industry-year observations that have fewer than 40 firms and weight all estimates by the inverse of the standard error of the previously-estimated dependent variable.¹⁶ Results are presented in Table 7. Standard errors are clustered at the industry level to allow for arbitrary correlation in the error terms over time within industries.

There are a number of key findings. First, industry-years with a one standard deviation more skilled workforce have a gender wage-productivity gap that is 19.2 percentage points higher if they have the mean level of product market competition and difficulty hiring. Second, this gap is doubled if the industry-year is one standard deviation less competitive, or is eliminated if the industry-year is one standard deviation more competitive than average. Third, this additional effect of lower levels of competition is eliminated if the industry-year has a one standard deviation higher difficulty in hiring. Overall, we find that the gender wage-productivity gap is larger in industry-years with higher skilled workers, lower levels of product market competition, and more competitive hiring markets.

The second column shows that these results are robust to including 26 two-digit NZSIOC industry fixed effects; this implies that differences over time and across very similar industries in product market competitiveness, workers skills, and difficulty hiring are correlated with differences in the gender wage-productivity gap in these industries. In the third column, we add three-digit NZSIOC industry fixed effects and hence rely on only time variation to identify our model. While none of our coefficient estimates are significantly different than those in the

¹⁶ Our results are qualitatively unaffected by instead weighting by the inverse of the standard error of the previously-estimated dependent variable multiplied by the number of workers in each industry-year.

previous specification, none are now significantly different from zero either. Moving from two-digit to three-digit industry fixed effects only increases the R-squared of the model by 0.007 suggesting that two-digit industry fixed effects capture most of the across-industry variation in the gender wage-productivity gap.

These results are consistent with models of taste discrimination, especially those that allow for frictions in job search (Black 1995; Bowlus and Eckstein 2002; Flabbi 2010), and strongly suggest that taste discrimination persists in these industries and is important for explaining the overall gender wage gap. We provide new evidence that discrimination is greatest when product markets are non-competitive for an industry but firms find it easy to hire. One can imagine this to be a commonplace situation as jobs in profitable industries that share rents with their employees are likely to be in demand. We are the first paper to our knowledge to show that discrimination is concentrated in industries that employ more high-skilled workers. Our finding that the gender wage-productivity gap is smaller in less competitive industries with a more skilled workforce when there is it is more difficult for firms to hire workers is less consistent with gender differences in bargaining power being important for explaining the overall gender wage gap.

7. Conclusions

In this paper, we use a decade of annual wage and productivity data from New Zealand's Linked Employer-Employee Database (LEED) to evaluate the relative importance of each of the following explanations for the gender wage gap: i) sorting of workers into different industries and firms; ii) productivity differences; iii) discrimination either deriving from preferences (taste) or judgements about expected productivity (statistical); and iv) differences in bargaining ability, especially in regards to rent sharing.

We begin by examining the importance of gender differences in sorting between industries and firms in explaining the gender wage gap in a standard wage equation. We find that gender differences in sorting between either industries or firms explain less than one-fifth of the gender wage gap. Next, we jointly estimate translog firm production functions and wage bill equations that allow us to examine the relative marginal contribution of women to both firm output and wages. Comparing the relative productivity of women to their relative wages gives us a measure of whether women are paid less than their marginal product. Of the 17.5 percentage point gender wage gap that remains unexplained by sorting at most 2 percentage points is explained by gender productivity differences. The remaining gender wage gap could come

about because women are less good at bargaining for firm rents or because employers discriminate against women for either taste or statistical motives.

We take two approaches to attempt to distinguish between these reasons. First, we examine how the relative wage-productivity gap varies by age and firm tenure. The estimated relative gender wage-productivity gap increases with both age and tenure, even though women's relative productivity is equal to that for men up to age 40 and does not vary by tenure. These results are inconsistent with simple models of statistical discrimination that argue that, as individuals reveal their true productivity, differences between wages and productivity should decline. Second, we examine how the relative wage-productivity gap varies across industries and over time and how this relates to both female bargaining power and the ability of firms to taste discriminate against female workers. We find that the gap is larger in industry-years with higher skilled workers, lower levels of product market competition, and less difficulty in hiring workers. These results are consistent with models of taste discrimination and strongly suggests that taste discrimination persists in these industries and is important for explaining the overall gender wage gap. Our findings are less consistent with gender differences in bargaining power being important for explaining the overall gender wage gap.

New Zealand is a small open economy with a gender wage gap similar to the US and many European countries. It has a highly flexible labor market with low rates of unionization and centralized bargaining, high female employment rates, and the economy is strongly dominated by the service sector. In many ways, New Zealand can be seen as leader in the increased reliance on international trade and reduced employment in manufacturing and routine task-intensive activities that is currently happening in most OECD countries. These characteristics of the New Zealand labor market should favor women in many ways. Hence, our finding that taste discrimination against women persists in many sectors of the economy is discouraging and suggests that stronger enforcement of equal pay regulations could be beneficial for many women. It also seems quite likely that similar situations exist in the majority of OECD countries, which have less flexible labor markets or a slower shift towards the service sector.

References

- Aigner, Dennis J. and Glen G. Cain, 1977. "Statistical Theories of Discrimination in Labor Markets," *Industrial and Labor Relations Review* 30 (2), 175-187.
- Altonji, Joseph G. and Charles R. Pierret. 2001. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116 (1): 313-350.
- Anderson, Deborah J., Melissa Binder and Kate Krause, 2002. "The Motherhood Wage Penalty: Which Mothers Pay it and Why?" *American Economic Review*, 92 (2), 354-358.
- Angelov, Nikolay, Per Johansson and Erica Lindahl, 2016. "Parenthood and the gender gap in pay." *Journal of Labor Economics*, 34(3), 545-579.
- Autor, David H. 2014. "Skills, Education, and the Rise of Earnings Inequality Among the "Other 99 Percent" *Science*, 23 May 2014: 344 (6186), 843-851.
- Azmat, Ghazala and Rosa Ferrer (2016) "Gender Gaps in Performance: Evidence from Young Lawyers." forthcoming *Journal of Political Economy*.
- Babcock, Linda and Sarah Laschever. 2003. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton, NJ: Princeton University Press.
- Bailey, Martha J., Brad Hershbein, and Amalia R. Miller. 2012 "The Opt-In Revolution? Contraception and the Gender Gap in Wages." *American Economic Journal: Applied Economics*, 4(3): 225-254.
- Bayard, Kimberly, Judith Hellerstein, David Neumark and Kenneth Troske (2003) "New evidence on sex segregation and sex differences in wages from matched employee-employer data." *Journal of Labor Economics* 21(4): 887-922
- Baert, Stijn, Bart Cockx, Niels Gheyle and Cora Vandamme. 2015. "Is there Less Discrimination in Occupations where Recruitment is Difficult." *Industrial Labor Relations Review* 68(3): 467-500.
- Becker, Gary S (1957) *The economics of discrimination*. Chicago, IL: The University of Chicago Press.
- Bertrand, Marianne, Emir Kamenica and Jessica Pan. 2015. "Gender Identity and Relative Income Within Household." *Quarterly Journal of Economics*, 130 (2): 571-614.
- Black, Sandra E and Philip E Strahan (2001) "The division of spoils: rent-sharing and discrimination in a regulated industry." *American Economic Review* 91(4): 814-830.
- Black, Dan A., 1995. "Discrimination in an equilibrium search model." *Journal of Labor Economics*, 13(2), 309-333.
- Blau, Francine D. and Lawrence M. Kahn (2000), 'Gender Differences in Pay', *Journal of Economic Perspectives*, 14(4): 75--99.
- Blau, Francine D. and Lawrence M. Kahn (2016). "The gender-wage gap: Extent, trends, and explanations," forthcoming *Journal of Economic Literature*.
- Bowlus, Audra J., and Zvi Eckstein (2002) "Discrimination and skill differences in an equilibrium search model." *International Economic Review* 43 (4), 1309-1345.
- Brenzel, Hanna, Hermann Gartner and Claus Schnabel. 2014. "Wage bargaining or wage posting? Evidence from the employers' side." *Labour Economics* 29: 41-48.

- Card, David, Ana Rute Cardoso and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics*, 131(2), 633-686.
- Carroll, Nicholas, and Julian Wood. 2003. "Preliminary Research into Sustainable Employment Measures Using the Linked Employer-Employee Database (LEED)." *Statistics New Zealand Report*, September.
- Dixon, Sylvia (2000) Pay Inequality between Men and Women in New Zealand, *Department of Labour, Occasional Paper 2000/1*, www.dol.govt.nz/
- Ellingsen, Tore and Åsa Rosén. 2003. "Fixed or Flexible? Wage-Setting in Search Equilibrium." *Economica* 70(278), 233-250
- Evans, Lewis, Grimes, Arthur, Wilkinson, Bryce, Teece, David, 1996. Economic reform in New Zealand 1984–95: the pursuit of efficiency. *Journal of Economic Literature* 34 (4), 1856–1902
- Fabling, Richard (2009) "A rough guide to New Zealand's Longitudinal Business Database", *Global COE Hi-Stat Discussion Paper GD09-103*.
- Fabling, Richard. 2010. "Keeping it Together: Tracking Firms in New Zealand's Longitudinal Business Database," *Motu Working Paper* 11-01.
- Fabling, R., & Maré, D. C. (2015a). "Production function estimation using New Zealand's Longitudinal Business Database." *Motu Economic and Public Policy Research Working Paper* 15-15 (September).
- Fabling, R., & Maré, D. C. (2015b). "Addressing the absence of hours information in linked employer-employee data." *Motu Economic and Public Policy Research Working Paper* 15-17 (October).
- Flabbi, Luca, 2010. "Gender discrimination estimation in a search model with matching and bargaining." *International Economic Review*, 51(3), 745-783.
- Gneezy, Uri, Muriel Niederle, and Aldo Rustichini, Aldo. 2003. "Performance in Competitive Environments: Gender Differences." *Quarterly Journal of Economics* 118: 1049-1074.
- Gneezy, Uri, Kenneth L. Leonard and John A. List. 2008. "Gender Differences in Competition: Evidence from a Matrilineal and a Patriarchal Society." *Econometrica* 77: 909-931.
- Goldin, Claudia, 2014. "A grand gender convergence: Its last chapter." *American Economic Review*, 104(4), 1091-1119.
- Heinze, Anja and Elke Wolf, 2010. "The intra-firm gender wage gap: a new view on wage differentials based on linked employer–employee data," *Journal of Population Economics* 23(3), pages 851-879.
- Hellerstein, J. K. and Neumark, D. (1999), 'Sex, Wages and Productivity: An Empirical Analysis of Israel Firm-level Data', *International Economic Review*, 40(1): 95--123.
- Hellerstein, J.; Neumark, D.; Troske, K. (1999), 'Wages, Productivity, and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations', *Journal of Labor Economics*, 17(3): 409-446.
- Hirsch, Boris, Michael Oberfichtner and Claus Schnabel. 2014. "The levelling effect of product market competition on gender wage discrimination." *IZA Journal of Labor Economics* 3:19 (December).

- Jacobsen, Joyce, Melanie Khamis and Mutlu Yuksel, 2015. "Convergence in men's and women's life patterns: Lifetime work, lifetime earnings, and human capital investment," *Research in Labor Economics*, Gender Convergence in the Labor Market, Emerald Group Publishing Limited, 41, 1-33.
- Kelly, Neil. 2003. "Prototype Outputs Using Linked Employer-Employee Data." *Statistics New Zealand Report*, September.
- Kunze, Astrid (2017) "The Gender Wage Gap in Developed Countries." *IZA Discussion Paper* 10826 (June).
- Macis, Mario and Fabiano Schivardi. 2016. "Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills?" *Journal of Labor Economics* 34(4): 945-978.
- Maré, David C, and Dean R Hyslop. 2006. "Worker-Firm Heterogeneity and Matching: An Analysis Using Worker and Firm Fixed Effects Estimated from LEED." Wellington: Statistics New Zealand.
- Maré, David C, Dean R Hyslop and Richard Fabling. 2015. "Firm Productivity Growth and Skill." *Motu Economic and Public Policy Research Working Paper* 15-18 (October).
- McLaughlin, Kenneth J. 1994. "Rent Sharing in an Equilibrium Model of Matching and Turnover." *Journal of Labor Economics* 12(4): 499-523.
- Mercante, Joseph and Penny Mok. 2014. "Estimation of Labour Supply in New Zealand" *New Zealand Treasury Working Paper* 14/08.
- Olivetti, Claudia and Barbara Petrongolo, 2016. "The evolution of gender gaps in industrialized countries." *Annual Review of Economics* 8, 405-434.
- Olley, S. and Pakes, A. (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry", *Econometrica*, 64(6), 1263-1298.
- Statistics New Zealand. (2014). *Productivity Statistics (1978-2013) - Hot Off The Press*. Wellington: Statistics New Zealand.

Table 1: Descriptive Statistics for Individuals and Firms

	Mean	Standard Deviation
Individual Data - All Individuals		
Female	0.496	0.500
Aged <25	0.247	0.431
Aged 25 to 39	0.337	0.473
Aged 40 to 54	0.289	0.453
Aged 55+	0.127	0.333
Average monthly earnings - Men	3,655	3,958
Average monthly earnings - Women	2,437	2,327
Average FTEs - Men	0.816	0.267
Average FTEs - Women	0.725	0.303
Average firm head count (geometric mean)	92.8	12.6
Firm had multiple plants in 2+ months	0.469	0.499
Individuals		3,132,700
Individuals*Years		11,265,400
Individual Data - Individuals Employed at Firms with 5+ Workers and in Productivity Data		
Female	0.425	0.494
Aged <25	0.290	0.454
Aged 25 to 39	0.347	0.476
Aged 40 to 54	0.261	0.439
Aged 55+	0.102	0.302
Average monthly earnings - Men	3,774	4,015
Average monthly earnings - Women	2,379	2,397
Average FTEs - Men	0.829	0.258
Average FTEs - Women	0.724	0.301
Average firm head count (geometric mean)	152.9	9.0
Firm had multiple plants in 2+ months	0.547	0.498
Individuals		2,050,200
Individuals*Years		5,425,400
Firm Data - Firms with 5+ Workers and in Productivity Data		
Gross output (\$000,000)	288	943
Total wage bill (\$000,000)	52.4	118.3
Intermediate consumption (\$000,000)	164	693
Capital (\$000,000)	37.3	123.8
Number of Employee (unweighted)	32	199.7
Number of Employee	1280	2856
Share female	0.422	0.259
Share with age <25	0.251	0.203
Share with age 25-39	0.353	0.150
Share with age 40-54	0.287	0.147
Share with age 55+	0.107	0.105
Share with no working proprietors	0.721	0.448
Share with one working proprietor	0.111	0.314
Share with 2-5 working proprietors	0.159	0.365
Share with 6-10 working proprietors	0.003	0.058
Share with 11+ working proprietors	0.005	0.072
Firms		66,765
Firms*Years		290,490

Note: All employment variables refer to the firm where in individual received the most pay during the year. Earnings are in nominal dollars. Unless noted all firm level variables are weighted by the number of employees at each firm so that the figures are representative for the average workers. The paper provides more information on how each sample is defined.

Table 2: Industry and Firm Sorting and the Gender Wage Gap

Dependent variable: ln(earnings)						
Variable	<u>All Individuals</u>			<u>Individuals Working in Firms in the Restricted Analysis Sample</u>		
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-0.222** (0.000)	-0.207** (0.001)	-0.195** (0.001)	-0.235** (0.001)	-0.203** (0.001)	-0.192** (0.001)
Aged 25 to 39	0.346** (0.000)	0.277** (0.000)	0.263** (0.000)	0.387** (0.001)	0.304** (0.001)	0.283** (0.001)
Aged 40 to 54	0.393** (0.001)	0.323** (0.001)	0.328** (0.001)	0.446** (0.001)	0.364** (0.001)	0.356** (0.001)
Aged 55+	0.288** (0.001)	0.229** (0.001)	0.255** (0.001)	0.318** (0.001)	0.256** (0.001)	0.266** (0.001)
FTEs at highest paying job (ln)	1.220** (0.000)	1.172** (0.000)	1.131** (0.001)	1.200** (0.001)	1.135** (0.001)	1.089** (0.001)
4-Digit Industry Fixed Effects		Yes			Yes	
Firm by Year Fixed Effects			Yes			Yes
R-squared	0.802	0.829	0.872	0.798	0.833	0.864
Observations	11,265,400	11,265,400	11,265,400	5,425,400	5,425,400	5,425,400

Notes: Asterisks denote: * p<0.05, ** p<0.01. Standard errors, in parentheses, are robust and clustered at the individual level. This table presents the results of OLS regressions at the individual-year level where the dependent variable is (log) annual earnings from an individual's highest paying job. The sample in columns (1) to (3) is employees in all firms; the sample in columns (4) to (6) is employees in firms with at least 5 employees and valid productivity data. All columns use a random 50% sample for empirical tractability. All regressions include year fixed effects.

Table 3: Joint NLSUR Estimation of Augmented Translog Production Function

Measure of Labour Inputs	Head Count	FTEs	FT and PT as Separate Inputs	Head Count w/ FTE Adjustment
Panel A: Production Function				
Phi female	0.688** (0.036)	0.930 (0.041)	1.077 (0.054)	0.977 (0.045)
Phi part-time			0.458*** (0.018)	
FTE adjustment coef				1.325** (0.063)
Phi age < 25	0.204** (0.052)	0.534** (0.039)	0.676** (0.046)	0.624** (0.041)
Phi age 40-54	0.778** (0.051)	0.838** (0.044)	0.861** (0.050)	0.858** (0.045)
Phi age 55+	0.179** (0.052)	0.448** (0.051)	0.473** (0.060)	0.490** (0.054)
Panel B: Wage Bill Equation				
Phi female	0.608** (0.013)	0.773** (0.010)	0.908** (0.011)	0.817** (0.011)
Phi part-time			0.444** (0.005)	
FTE Adjustment coef				1.388** (0.019)
Phi age < 25	0.220** (0.026)	0.499** (0.011)	0.637** (0.011)	0.590** (0.010)
Phi age 40-54	0.848** (0.016)	0.922** (0.011)	0.949** (0.011)	0.950** (0.011)
Phi age 55+	0.372** (0.025)	0.638** (0.013)	0.716** (0.013)	0.694** (0.012)
Panel C: Percent Difference in Contribution to the Wage Bill and Productivity (1 - phi_wb/phi_pf)				
Female	11.6%** (0.041)	16.9%** (0.036)	15.7%** (0.042)	16.4%** (0.038)
Age < 25	-7.9% (0.165)	6.6% (0.067)	5.7% (0.062)	5.4% (0.061)
Age 40-54	-9.0% (0.064)	-10.0% (0.057)	-10.3% (0.062)	-10.6% (0.056)
Age 55+	-108.2%* (0.518)	-42.4%** (0.152)	-51.6%** (0.183)	-41.5%** (0.148)
Number of Firms	290,490	290,490	290,490	290,490

Notes: Asterisks on phis indicate difference from 1, on the comparison terms in Panel C difference from 0: * p<0.05, ** p<0.01. Standard errors in parentheses account for clustering at the firm level. Results are from non-linear estimation of translog production function and wage bill equation, allowing labour contribution to vary by gender and age separately. All regressions are weighted by either firm head count (cols 1 and 3) or firm FTE (cols 2 and 4). The phis represent the relative contribution of a particular group to the production or wage bill compared to either men or individuals aged 25-39. All models include 2-digit ANZSIC06 industry and year fixed effects, as well as controls for (log) labour, (log) capital, (log) intermediate consumption, the squares and interactions between these inputs and controls for the number of working proprietors at the firm.

Table 4: Production Function and Wage Bill Robustness Estimates (Head Count w/ FTE Adjustment)

Robustness Check	Main Estimates	Without Industry Fixed Effects	Only Firms with <2 WPs	Value Added Specification	Endogenous Inputs	Endogenous Inputs (No Neg Investment)
Panel A: Production Function						
Phi female	0.977 (0.045)	1.215** (0.068)	0.960 (0.051)	0.972 (0.035)	0.975 (0.045)	0.993 (0.049)
Panel B: Wage Bill Equation						
Phi female	0.817** (0.011)	0.915** (0.015)	0.822** (0.013)	0.815** (0.011)	0.814** (0.012)	0.815** (0.012)
Number of Firms	290,490	290,490	196,164	286,383	219,996	163,725
Panel C: Percent Difference in Contribution to the Wage Bill and Productivity (1 - phi_wb/phi_pf)						
Female	16.4%** (0.038)	24.7%** (0.035)	14.4%** (0.046)	16.1%** (0.029)	16.5%** (0.038)	17.9%** (0.041)
Number of Firms	290,490	290,490	196,164	286,383	219,996	163,725

Notes: Asterisks on phis indicate difference from 1, on the comparison terms in Panel C difference from 0: * p<0.05, ** p<0.01. Standard errors in parentheses account for clustering at the firm level. Results are from non-linear estimation of translog production function and wage bill equation, allowing labour contribution to vary by gender and age. Labour is specified as head count * (average FTEs)^{delta} in all regressions and all regressions are weighted by firm FTE. The phis represent the relative contribution of a particular group to the production or wage bill compared to men or individuals aged 25-39. All models unless noted include 2-digit ANZSIC06 industry and year fixed effects, as well as controls for (log) labour, (log) capital, (log) intermediate consumption, the squares and interactions between these inputs and controls for the number of working proprietors at the firm. Results for age-group differences are available in Appendix Table 1.

Table 5: Gender Differences in Wages and Productivity by Age and Tenure

	Panel A: Production Function		Panel B: Wage Bill Equation		Panel C: % Diff in Contribution to the Wage Bill and Productivity ($1 - \phi_{wb}/\phi_{pf}$)	
	Main	Female interaction	Main	Female interaction	Main	Female interaction
Age < 25	0.638** (0.060)	1.031 (0.121)	0.505** (0.012)	1.154** (0.032)	20.8%** (0.073)	-12.0% (0.127)
Age 25-39	base	1.064 (0.060)	base	0.893** (0.015)	base	16.0%** (0.046)
Age 40-54	0.950 (0.064)	0.732** (0.060)	1.078** (0.019)	0.575** (0.013)	-13.5% (0.073)	21.4%** (0.061)
Age 55+	0.360** (0.078)	1.672 (0.417)	0.670** (0.018)	0.849** (0.028)	-86.3%* (0.383)	49.2%** (0.121)
	Main	Female interaction	Main	Female interaction	Main	Female interaction
<1 Year Tenure	base	0.881 (0.128)	base	0.985 (0.023)	base	-11.8% (0.167)
1-2 Years Tenure	1.162 (0.127)	0.948 (0.086)	1.102** (0.019)	0.745** (0.018)	5.2% (0.101)	21.4%** (0.067)
2+ Years Tenure	1.145 (0.114)	1.027 (0.054)	1.020 (0.015)	0.795** (0.011)	10.9% (0.087)	22.6%** (0.040)

Notes: Asterisks on phis indicate difference from 1, on the comparison terms in Panel C difference from 0: * $p < 0.05$, ** $p < 0.01$. Standard errors in parentheses account for clustering at the firm level. Results are from non-linear estimation of translog production function and wage bill equation, allowing labour contribution to vary by the intersection of gender and age or tenure and age. Labour is specified as head count * (average FTEs)^{delta} in all regressions and all regressions are weighted by firm FTE. The phis represent the relative contribution of a particular group to the production or wage bill compared to men aged 25-39 or men with less than one year of tenure. All models include 2-digit ANZSIC06 industry and year fixed effects, as well as controls for (log) labour, (log) capital, (log) intermediate consumption, the squares and interactions between these inputs and controls for the number of working proprietors at the firm. The number of observations is 290,490.

Table 6: Workers Skills, Competition and Difficulty Hiring by Industry/Year

Variables	Mean	Corr w/ 1st PC
Employee Skills		
Fraction of employees managers and professionals	0.145	0.320
Fraction of employees technicians and associate professionals	0.083	0.320
Fraction of employees tradespeople and related workers	0.158	0.067
Fraction of employees other occupations	0.614	-0.343
Average fraction of employees aged <25	0.216	-0.279
Average fraction of employees aged 25-39	0.348	0.304
Average fraction of employees aged 40-54	0.310	0.188
Average fraction of employees aged >55	0.126	-0.014
Average fraction of employees in worker FE quartile 1	0.199	-0.344
Average fraction of employees in worker FE quartile 2	0.271	-0.410
Average fraction of employees in worker FE quartile 3	0.280	-0.041
Average fraction of employees in worker FE quartile 4	0.250	0.423
Proportion of variation explained by 1st principal component		0.426
Competition and Profitability		
Fraction of firms w/ no competition	0.047	0.529
Fraction of firms w/ 1-2 competitors	0.235	0.519
Fraction of firms w/ many competitor, some dominant	0.578	-0.579
Fraction of firms w/ many competitor, none dominant	0.186	0.180
Average log capital-labor ratio	9.83	0.289
Proportion of variation explained by 1st principal component		0.476
Difficulty Hiring Workers		
Fraction of firms w/ severe difficulty hiring managers and profs	0.159	0.164
Fraction of firms w/ some difficulty hiring managers and profs	0.515	0.450
Fraction of firms w/ no difficulty hiring managers and profs	0.326	-0.550
Fraction of firms w/ severe difficulty hiring technicians and asc profs	0.196	0.222
Fraction of firms w/ some difficulty hiring technicians and asc profs	0.499	0.357
Fraction of firms w/ no difficulty hiring technicians and asc profs	0.305	-0.540
Proportion of variation explained by 1st principal component		0.479

Note: The results above are on three separate principal component analyses of the groups of variables listed for 273 industry-years. The sample starts in 2005 as this is the first year where information on competition and difficulty hiring workers is available.

Table 7: Explaining Variation in the Gender Wage-Productivity Gap

Variable	(1)	(2)	(3)
Higher skill (1st PCA component)	0.192*** (0.033)	0.192*** (0.061)	0.066 (0.142)
Less competition (1st PCA Component)	0.059* (0.030)	0.011 (0.027)	0.059 (0.071)
More difficulty hiring (1st PCA component)	0.076* (0.041)	0.063* (0.036)	0.058 (0.048)
Higher skill * Less competition	0.196*** (0.053)	0.148*** (0.051)	0.103 (0.075)
Higher skills * More difficulty hiring	-0.009 (0.032)	-0.038 (0.028)	-0.017 (0.035)
Less competition * More difficulty hiring	0.004 (0.030)	0.034 (0.030)	-0.001 (0.054)
Higher skill * Less competition * More difficulty hiring	-0.139*** (0.045)	-0.118** (0.045)	-0.079 (0.061)
2-Digit Industry Fixed Effects	No	Yes	No
3-Digit Industry Fixed Effects	No	No	Yes
R-squared	0.076	0.129	0.136
Observations	266	266	266

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors, in parentheses, are clustered at the 3-digit industry level. This table presents the results of OLS regressions at the 3 digit industry-year level where the dependent variable is the estimated percent difference in the contribution of female employment to the wage bill and to productivity in each industry-year. Seven industry-years with fewer than 40 firms are dropped. All independent variables are the first principal component described in table 6 and are standardized to have mean 0 and standard deviation 1. Only data from 2005 to 2011 is included and all regressions are weighted by the inverse of the standard error of the estimated dependent variable. All regressions include year fixed effects as well as the listed variables.

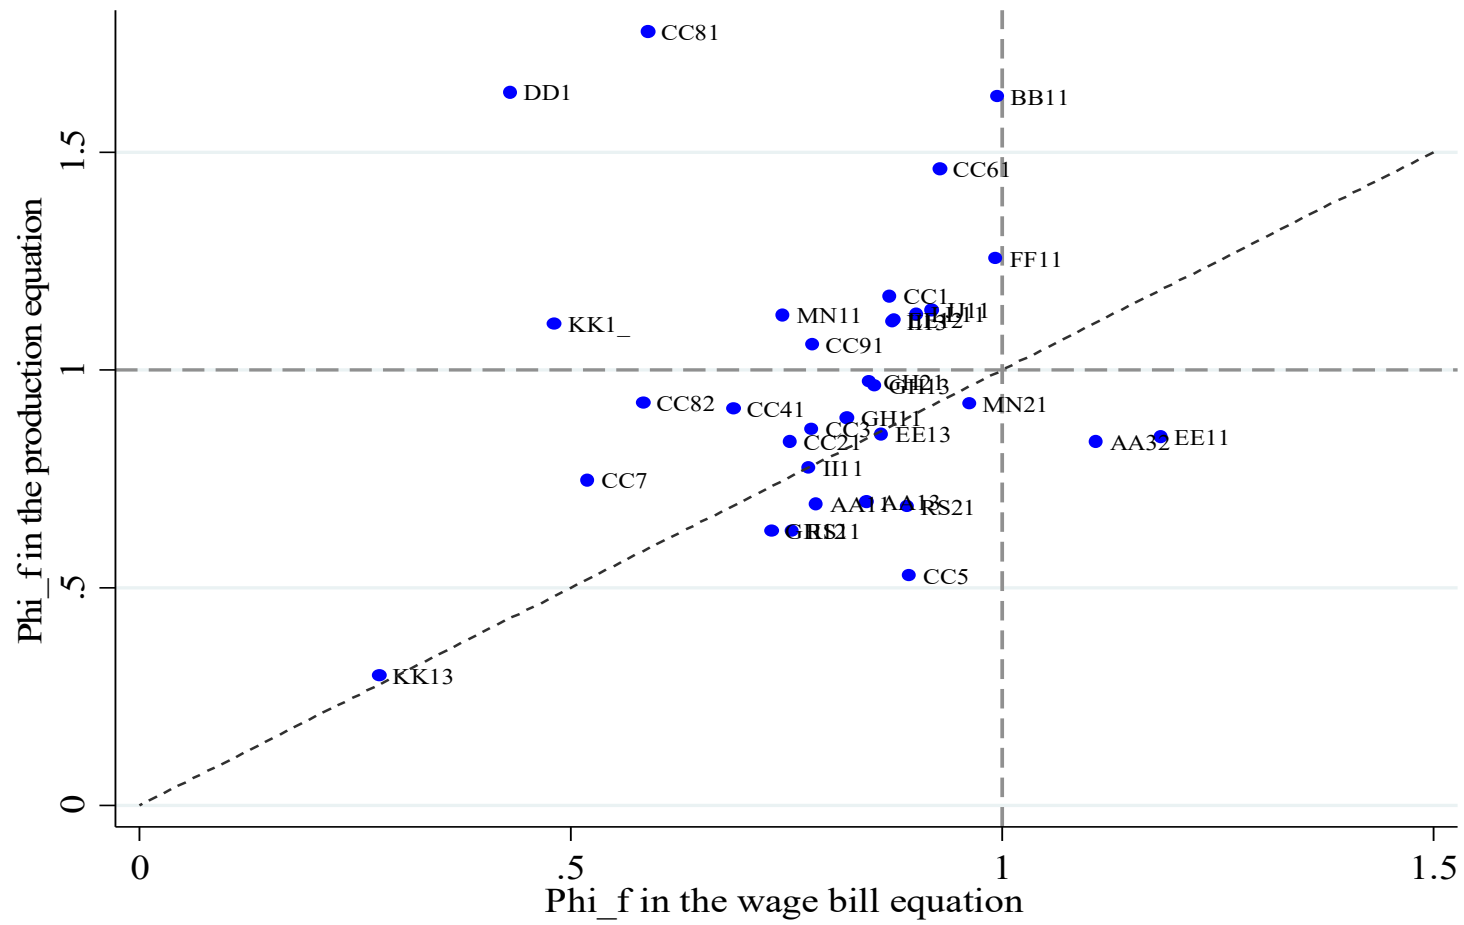


Figure 1: Variation in Female Contribution to Production and Wages across Industries

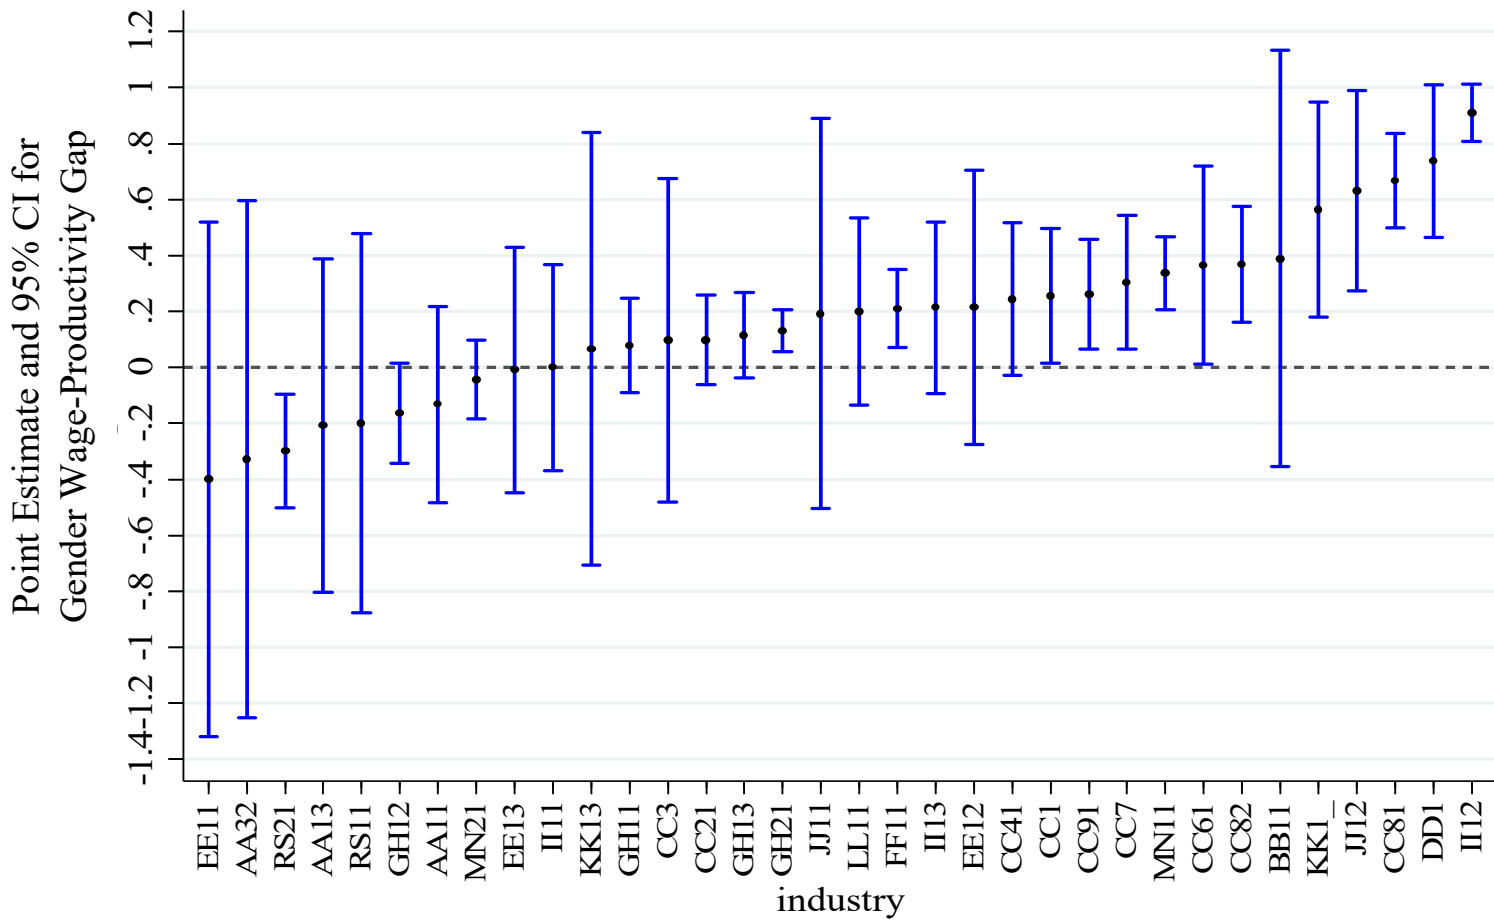


Figure 2: Variation in the Gender Wage-Productivity Gap across Industries

Appendix Table 1: Production Function and Wage Bill Robustness Estimates (Head Count w/ FTE Adjustment)

Robustness Check	Main Estimates	Without Industry Fixed Effects	Only Firms with <2 WPs	Value Added Specification	Endogenous Inputs	Endogenous Inputs (No Neg Investment)
Panel A: Production Function						
Phi age < 25	0.624** (0.041)	0.282** (0.042)	0.618** (0.052)	0.519** (0.026)	0.569** (0.039)	0.546** (0.041)
Phi age 40-54	0.858** (0.045)	0.640** (0.044)	0.863* (0.054)	0.830** (0.032)	0.802** (0.041)	0.766** (0.043)
Phi age 55+	0.490** (0.054)	0.264** (0.053)	0.429** (0.062)	0.461** (0.042)	0.427** (0.052)	0.398** (0.057)
Panel B: Wage Bill Equation						
Phi age < 25	0.590** (0.010)	0.410** (0.012)	0.546** (0.012)	0.577** (0.010)	0.581** (0.011)	0.576** (0.012)
Phi age 40-54	0.950** (0.011)	0.797** (0.016)	0.958** (0.014)	0.946** (0.011)	0.942** (0.012)	0.940** (0.013)
Phi age 55+	0.694** (0.012)	0.514** (0.019)	0.670** (0.015)	0.671** (0.013)	0.685** (0.013)	0.678** (0.014)
Number of Firms	290,490	290,490	196,164	286,383	219,996	163,725
Panel C: Percent Difference in Contribution to Productivity and Wage Bill (1 - phi_wb/phi_pf)						
Age < 25	5.4% (0.061)	-45.7%* (0.195)	11.6% (0.074)	-11.1%* (0.054)	-2.0% (0.070)	-5.5% (0.081)
Age 40-54	-10.6%* (0.056)	-24.4%** (0.075)	-11.0%* (0.067)	-14.1%** (0.041)	-17.3%** (0.059)	-22.8%** (0.068)
Age 55+	-41.5%** (0.148)	-94.4%** (0.349)	-56.0%** (0.214)	-45.4%** (0.123)	-60.2%** (0.185)	-70.3%** (0.231)
Number of Firms	290,490	290,490	196,164	286,383	219,996	163,725

Notes: Asterisks on phis indicate difference from 1, on the discrimination terms difference from 0: * p<0.05, ** p<0.01. Standard errors in parentheses account for clustering at the firm level. Results are from non-linear estimation of translog production function and wage bill equation, allowing labour contribution to vary by gender and age. Labour is specified as head count * (avg FTEs)^{delta} in all regressions and all regressions are weighted by firm FTE. The phis represent the relative contribution of a particular group to the production or wage bill compared to men or individuals aged 25-39. All models unless noted include 2-digit ANZSIC06 industry and year fixed effects, as well as controls for (log) labour, (log) capital, (log) intermediate consumption, the squares and interactions between these inputs and controls for the number of working proprietors at the firm.