SCHOOL OF
ECONOMICS AND
MANAGEMENT

LUND
UNIVERSITY

Convex Earnings, Childcare, and the Gender Pay Gap

An Exploratory Study of Canadian College Graduates

August Hedenstierna Jonson

Bachelor's Thesis

Lund University School of Economics and Management

Supervisor: Sarah Rosenberg, Lund University

Submitted January 8th, 2024

Abstract

In the last ten years labor economists have become increasingly interested in convex earnings structures in the labor market and their impact on the gender pay gap, an approach spearheaded by Claudia Goldin in her 2014 Presidential Address to the American Economic Association. The study has seen widespread acclaim, but very little replication. This study replicates parts of Claudia Goldin's work, applying it to college graduates in Canada and in five subnational provinces and regions. First, using 2015-2017 releases of the Canadian Labour Force Survey, it plots the gender pay gap residual across different age categories. Second, it investigates the correlation between the pay gap residual and convex earnings structures as proxied by the elasticity of income with respect to work hours. Third, it considers this correlation separately for individuals with and without children, and in light of provincial differences in childcare patterns, as recorded by Statistics Canada's General Social Survey on Time Use. The study does not attempt to establish statistical significance, but instead uses descriptive statistics to sketch an outline of patterns that would need further analysis to confirm. The results suggest that there is a gender pay gap present among college graduates in Canada that cannot be explained by age, profession, tenure or the raw number of work hours. There is mixed evidence that this can be explained by convex earnings structures for the population at large, though the evidence is stronger when only individuals with children are considered.

Key words: Gender pay gap, Canada, Convex earnings structures, Elasticity of income with respect to work hours, Childcare patterns

Table of Contents

Acknowledgement

I would like to thank Sarah Rosenberg of the Lund University School of Economics and Management for advising me throughout my bachelor's thesis project. Sarah's enthusiasm for the project kept it grounded while pushing it forward, and her advice turned the research process to one of the most educational experiences yet throughout my studies.

1. Introduction

When discussing the 20th century entry of women into the western labor force, the gender pay gap residual is often at the very center of discussion. It is understood as the portion of the pay gap that cannot be explained by labor market variables like differing education levels, career choices, et cetera. Its exact size and constitution are the primary parts of discussions on gender wage discrimination, and arguably at the core of labor economics as a field.

Claudia Goldin's seminal 2014 work *A Grand Gender Convergence: Its Last Chapter* can be considered one of the most important contributions to the discussion on the gender pay gap residual in decades. Among other things, it discusses the gender pay gap and convex earnings structures in the United States – how working many, consecutive work hours is disproportionately rewarded in the labor market, and how it disfavors women. Goldin's work constitutes the principal architecture of an entire convex earnings-pay gap framework that has been widely celebrated, not least by Goldin receiving the 2023 Sveriges Riksbank's Prize in Economic Sciences in memory of Alfred Nobel.

Even so, the framework has received surprisingly little replication. In a study replicating the method for the Italian labor market, economists Sergio Destefanis, Fernanda Mazzotta, and Lavinia Parisi of the University of Salerno remarked that the framework had until then seemingly not been replicated in other contexts than that of the United States (2023, p.2). Goldin's (2014) framework and arguments have been invoked in other cultural contexts (for instance Sweden in Angelov et al., 2016: p.571), without published attempts to find larger scale evidence of the theory. This emphasizes a need to replicate Goldin's (2014) findings in new contexts.

Furthermore, replication studies can serve to strengthen or disprove specific aspects of Goldin's (2014) overall framework. There is bound to be value in comparative study of convexity of earnings between subnational regions, so as to understand regional variations. Another avenue of study would be to compare subsections of a population that respond differently to convex earnings structures. Destefanis et al. (2023) did something like this by investigating separately college graduates and non-college graduates. For instance, Goldin (2014: p.1111) specifically mentions the role having children plays in reducing work hours. Considering this, studying how workers with and without children differ could both validate or disprove the larger Goldin framework, but also provide valuable insights for the gender pay gap at large.

Canada provides a natural next step for such a study, being culturally, economically, and geographically close to the US. Furthermore, Canadian childcare policy has historically been regulated largely individually by its subnational provinces (Pasolli, 2015). This provides an excellent opportunity for cross regional comparisons. Finally, the government statistical agency, Statistics Canada, has excellent data on both the labor market, and time use surveys for differences in childcare patterns (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b, 2018a).

This study explores the gender pay gap residual among college graduates in Canada. It does this using weekly figures for earnings and work hours, rather than monthly or annual. The gap is plotted across different age categories and regions within Canada. The paper then considers the correlation between the gender pay gap residuals for different professions and the presence of convex earnings structures, represented by elasticity of income with respect to work hours. This correlation (which I refer to as the "elasticity-pay gap correlation") is plotted for different subnational regions in Canada, as well as separately for populations with and without children. The paper does not attempt to establish statistical significance or draw firm conclusions – rather, it is an exploratory study that invites further research.

In the background, this paper starts by walking through the structure of convex earnings, and Goldin's (2014) argument how it impacts the gender pay gap. It then discusses the logic behind using elasticity of income with respect to work hours as a proxy for convex earnings structures. Finally, the background discusses the nation of Canada, and the reasons it makes for a good case study on the subject. In the methods section the quantitative work of this paper is carefully laid out. I generate information on the elasticity-pay gap correlation, the overall pay gap residuals in Canada, as well as differences in childcare trends within the nation. Finally, I discuss potential issues related to analyzing weekly data, and then compare the overall results and draw some general conclusions.

2. Background

2.1 – Literature Review

2.1.1 – Convex Earnings and Their Impact on the Gender Pay Gap

In their simplest form, convex earning structures encapsulate the phenomenon of work hours being increasingly productive (and by extension, increasingly rewarded in the workplace) when they are consecutive. This dynamic is shown in figure 1. The result is, among other things, that workers are rewarded for working long hours. There are other types of convexity tangentially relevant for the question at hand, one of which is elaborated upon in appendix I.

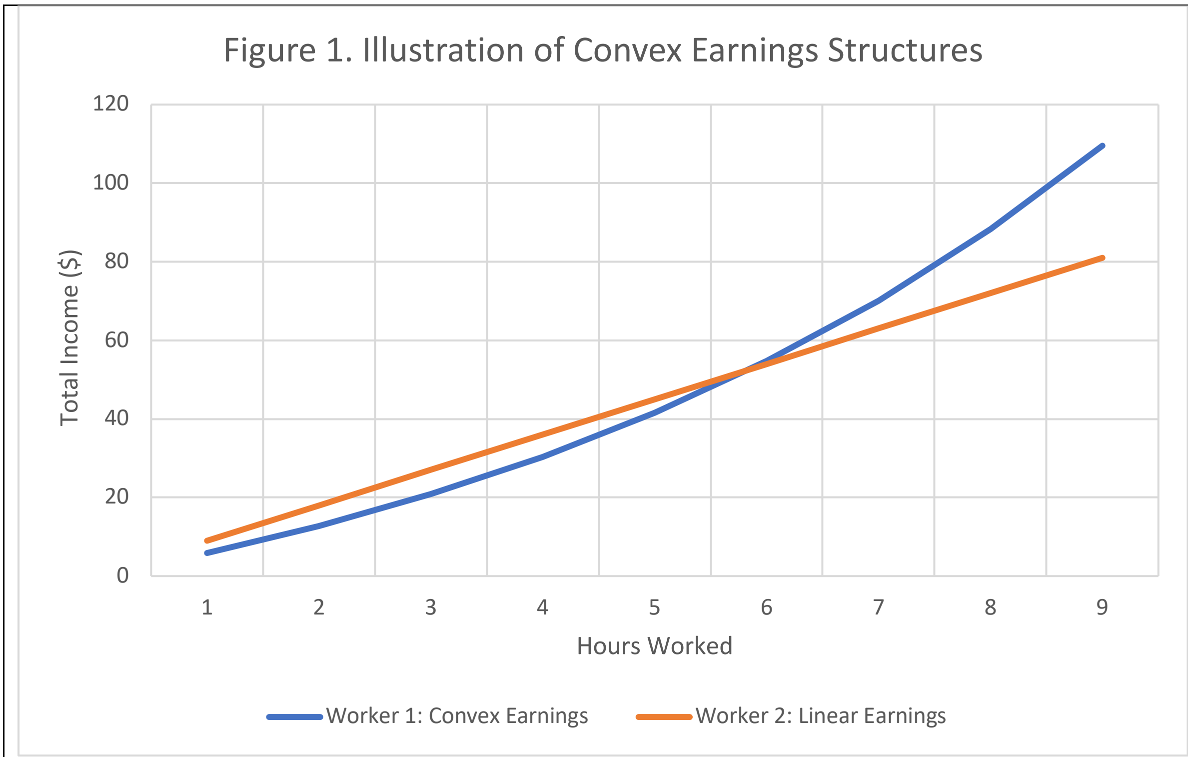

Notes: The above examples illustrates a simple form of convex earnings structures.

Claudia Goldin's 2014 *A Grand Gender Convergence: Its Last Chapter* was a landmark piece for analysis of convex earnings structures and their impact on the gender pay gap. It captures a significant part of the dynamic above. Its size is above all else dependent on the industry the individual works in, Goldin argues (2014: p.1098), and thus investigates the pay gap by employment sector and profession. She broadly finds that convex earnings structures, for instance earnings structures that reward working very long hours, are highly present in professions with high pay gap residuals. Conversely, professions where convexity of earnings structures are largely absent – for instance due to workers being easily substitutable (as in the case of for instance pharmacists (Goldin, 2014: pp.1115-1116)) – have a significantly smaller pay gap residuals.

The explanation for this phenomenon centers around women in the workforce generally demanding more flexible work hours than their male counterparts. While there are many factors to potentially consider in this dynamic – cultural norms, workplace discrimination, et cetera – one factor that Goldin (2014: p.1111) explicitly mentions is the role of childcare. Women are usually expected to take care of children to a larger extent than men, something that greatly limits their labor market flexibility, and their ability to work long hours. Consider again figure 1. A worker that needs to make time in their work schedule to attend the needs of their children – picking up children from school, attending parent-teacher conferences, et cetera – are more likely to have to cut their day short. In convex earnings structures, this parent will receive relatively less compensation. Since women more commonly take care of children (Statistics Canada, 2018b), this plays a significant part in the gender pay gap.

2.1.2 – Measuring Convex Earnings Structures

Convex earnings structures contribute to the pay gap residual in a way that would be hard to capture in large scale analysis. Ideally, one should consider the exact work hour structure of each worker. However, most labor force surveys usually only report raw earnings figures, and rarely if ever the exact structure of work hours. To investigate convex earnings structures in a quantitative manner, researchers need to use proxies.

Goldin (2014) approaches this problem by a variety of statistical means, but the most straight forward (if admittedly crude) method has to do with investigating the gender pay gap residual and its correlation with the elasticity of earnings with regards to work hours (see figure 3 in Goldin, 2014). When Goldin does this for different industries in the US, she finds a clear correlation between these variables – professions where this elasticity is high, tends to have larger pay gap residuals. Here the elasticity is used as proxy for convex earnings structures, and its correlation with the pay gap residual forms a first piece of evidence of their impact.

If wages paid within certain professions are more sensitive to the number of work hours than others, especially for overtime pay, then this would be visible in data as a higher elasticity of income with respect to work hours. Because of this, Goldin (2014), Bertrand

(2018) and Destefanis et al. (2023) interpret earnings elasticity with respect to work hours as a proxy for convex earning structures. The higher the elasticity, the more convex-like the earnings structures. In the light of figure 1 this correlation seems intuitive. Assuming women engaged in childcare are in aggregate are disfavored by convex earnings structures, we would expect this elasticity-pay gap correlation to be strong.

It is not a perfect proxy for convex earnings structures. It cannot capture nuances in different structure of work hours. A high elasticity controlling for overall work hours could be interpreted as a worker being rewarded by the labor marketplace for working very many hours, but not necessarily whether those work hours in a day are all consecutive (see Appendix I for an elaboration upon this). Furthermore, the elasticity does not itself convey work characteristics that would benefit from convexity of earnings. Goldin (2014) supplements this shortfall by analyzing workplace characteristics such as time pressure, freedom to make decisions, et cetera.

That being said, the crudeness of the measurement also provides several clear benefits. While the correlation should not itself be considered definitive evidence of convex pay structures, we would still expect it to be correlated with the gender pay gap residual, especially considering just how large a part of the gender pay gap is believed to be a result of convex earning structures (Goldin 2014, Destefanis 2023). Furthermore, the method is both transparent and intuitive. It does not contain as many subjective interpretations as Goldin's (2014) use of O*Net characteristics. Finally, since the data required for an elasticity-pay gap correlation study would be available in practically every national labor force survey, cross country comparison could be made very easily. In conclusion, elasticity of income with respect to work hours is a first step towards evidence of convex earnings structures, but it should be paired with supplemental analysis to form robust conclusions.

2.1.3 – Replication Studies

Goldin's work has spawned further academic discussion on long working hours and their impact on the gender pay gap. In a study on high earners, Bertrand's (2018) findings suggest something like Goldin's, suggesting that elasticity of hours worked increase for higher paying occupations. This implies that convex earnings structures become increasingly important for higher earners.

Goldin (2014) had a large focus on how the gender pay gap could be lessened by companies becoming better at dealing with flexible work hours. Cortés and Pan (2019), using the same datasets as Goldin (2014), turned the question to increasing the flexibility of highly educated female workers. In their words the punishment to being unable to work long hours would be lessened by "*supply of affordable and flexible substitutes for household production*" (2019: p.385). This focus on increased worker flexibility has also been argued by other authors, like Angelov et al. (2016). Overall, studies following Goldin (2014) have focused on the US and been broadly supportive of her views.

However, full scale replications of Goldin (2014) are exceedingly rare. The previously mentioned 2023 article by Destefanis et al. believes itself to be the first article replicating Goldin's (2014) findings outside of the United States, and despite extensive searching I was unable to find anything disproving this. The Destefanis et al. (2023) article does broadly confirm Goldin's findings, though with some important caveats. They find positive elasticity-pay gap correlations in the population at large, suggesting convex earnings structures are present and impactful. They do however find that even the lower segments of the Italian labor market are impacted by convex earnings structures. In terms of pure correlation scores, they are impacted more than high income earners. However, when replicating the workplace characteristics section of Goldin's (2014) paper, Destefanis et al. (2023) only find a correlation for graduate workers. This highlights that different population subsets might be impacted differently by convex earnings structures, and that these trends might differ between countries.

2.2 – Elaborating on the Literature

2.2.1 – Canada, Childcare, and Elaborating Upon Goldin

Goldin's (2014) framework enjoys a certain prestige, while remaining comparatively untested. Goldin claims in her article that the "*themes developed here are more broadly applicable*" (2014: p.1091), meaning outside of the US. It is therefore important to consider her framework in new cultural contexts, including countries with differing labor market structures. The framework could also be tested by subjecting it to different segments of the labor market. Destefanis et al. (2023) implicitly does something like this, by comparing their results between college graduates and non-college graduates. This method could be expanded upon to consider other population subsets, that are impacted differently by convex earnings structures.

One such population subdivision is whether workers have children or not. The roles of children in the gender pay gap are well documented (Angelov et al. 2016; Waldfogel 1997) and explicitly mentioned as influencing female labor supply by Goldin (2014: p.1111). If women are much more likely to take care of children and this encroaches on their temporal flexibility to the point that they cannot work long hours, this would be of importance when testing the Goldin framework. It could also have major policy ramifications. Considering more sharply the Goldin framework through the lens of childcare would likely yield both interesting and useful results.

To this end, this paper investigates Canada and subnational regions within it. This provides a new context for Goldin's (2014) model, in terms of a new labor market, culture, and economy. Canada as a country is also suitable for studying childcare, due to how it is structured within the country. It has historically been regulated by provinces (Pasolli, 2015) and not by the federal government until 2021 (Government of Canada, 2023). The provinces have historically had an overall market liberal framework supplemented by some public subsidies (Pasolli, 2015). The distinct exception to this is Québec, which in 1997 became the first North American jurisdiction to introduce a publicly funded childcare system, through their $5- (and later $7) dollar a day kindergarten program (Kohen et al., 2008). Consider this in light of Cortés and Pan (2019). Though it could be argued that the highly educated workers in their study could probably afford childcare, universal access to kindergarten services is a clear example of a substitute for household production, which would increase the flexibility of female workers. In brief, Canada provides excellent opportunities for study of the topic at hand.

2.2.2 – Formulating a Hypothesis

This purpose of this paper is threefold. First it seeks to map the overall gender pay gap residual in Canada across different regions and age categories. Second, it follows in Goldin's (2014) footsteps by plotting the elasticity-pay gap correlation for different professions, to approximate the importance of convex earnings structures for the gender pay gap. Instead of the very specific professions used in Goldin (2014), this study uses 43 broader categories of main task at work, as reported in the underlying dataset (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b). Thirdly, it breaks the elasticity-pay gap correlation down into workers with and without children, to approximate differences for the two groups. The paper then puts the regional results in the context of differing childcare trends in the different regions.

The analysis is done for Canada nationally, as well as for five regions within Canada: the provinces of British Columbia, Ontario and Québec, and the larger regions of the Prairies (Alberta, Saskatchewan, and Manitoba), and Atlantic Canada (Newfoundland and Labrador, New Brunswick, Prince Edward Island, and Nova Scotia). I will be referring to these as regions as opposed to provinces, as some are agglomerations of provinces. In line with Goldin's (2014) reasoning and parts of her paper, the analysis is restricted to individuals with a bachelor's degree or more, since in Goldin's own words "*…idiosyncratic temporal demands are generally more important for the highly-educated workers…*" (Goldin, 2014: p.1094).

The paper limits itself to just investigating the elasticity-pay gap correlation, and does not go into depths on workplace characteristics, such as O*NET in the case of Goldin (2014) and Destefanis et al. (2023). Nor does it delve deep into whether convex earnings structures are present, but rather maintains focus on specifically how it interfaces with the gender pay gap. A study on how convex earnings structures varies between Canadian provinces is a study that could easily be performed with the data sets presented in section "3. Data and Methods", but it is beyond the scope of the paper.

Rather, focusing on the elasticity-pay gap correlation by itself gives more space to explore how it interacts with childcare. To this end, this paper considers recorded amounts of time spent on childcare in each region, as mapped in the 2015 version of the Canadian time use survey (Statisics Canada, 2018) – at time of writing, the latest release. The elasticity-pay gap correlations are then placed in the context of these differing childcare patterns, and some conclusions are drawn.

This paper does not seek to draw definitive conclusions surrounding its questions, for two main reasons. Firstly, singular points of data, like the elasticity-pay gap correlations, would need supplemental analyses (like workplace characteristics analysis) to draw firm conclusions from. Secondly, as in the study that inspired it, the income figures calculated as the gender pay gap control for number of work hours. As such the paper makes no attempt at establishing statistical significance in the data used. Instead it uses descriptive statistics to sketch an outline of the general trends, from which qualitative reasoning is used to invite more precise analysis.

Overall, we would expect a pay gap residual to be generally existent and favoring men, as has been shown in literature on the Canadian gender pay gap (Fortin, 2019). Considering Goldin's (2014) framework we would consider the elasticity-pay gap correlation to be generally positive. Furthermore, considering the importance of children as a determinant of the gender pay gap (Angelov et al., 2016; Fortin, 2019), we would expect people with children to have significantly higher residuals than people without. Finally, as childcare limits long work hours, we would expect the elasticity-pay gap correlation to be lower for individuals that spend less time on childcare, such as people without children, or people living in regions where childcare is more egalitarian.

2.2.3 – Using Weekly Data

This study makes heavy use of six different releases of the Canadian Labour Force Survey (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b), which reports figures for singular weeks. Performing this survey on weekly figures as opposed to monthly could in theory carry some issues. For instance, individuals who did not work during the sample periods are omitted. Because of this, there is a risk that results from individual subsamples might be skewed, purely by bad luck. To account for this possibility, the study accounts for potential asymmetries in absence by region. This is elaborated upon in the methods section. Other issues with using weekly data are elaborated upon in the methods section.

3. Data and Methods

3.1 – Data

3.1.1 – General Social Survey on Time Use, Statistics Canada, 2015 Wave

The General Social Survey on Time Use is compiled and released by Statistics Canada, with the latest release being in 2018 for the survey period 2015. The survey publication includes aggregated figures on time spent on everyday activities including "Care of household children under 18 years". The public release contains compiled data for six different regions – Canada as a whole, the provinces of British Columbia, Ontario, and Québec, and two aggregates of a number of provinces: first being The Prairies (Alberta, Saskatchewan, and Manitoba), and the

Atlantic Provinces (New Brunswick, Prince Edward Island, Nova Scotia, and Newfoundland and Labrador). The survey excludes the sparsely populated territories of Yukon, The Northwest Territories, and Nunavut.

In each province, variables are released compiled in age categories for individuals 15 years old and up, in intervals of ten years (15-24 years old, 25-34 years old, et cetera) up to the age of 64, with a final category for people of 65 years or older. For the activity group *care of household children under 18 years*, data for each region was collected on the average number of hours spent on childcare per day, as well as participation rates in the activity reported as percent of the overall population. Each figure was gathered for men and women separately in all 6 regions (including Canada as a whole). Figures on these variables were only consistently reported on for age categories 25 to 34 years, 35 to 44 years, and 44 to 54 years, so the analysis will be constricted to these categories.

3.1.2 – Labour Force Survey 2015-2017

The Labour Force Survey (LFS), specifically the public use microdata file (PUMF) is compiled through household interviews and released by Statistics Canada every month. The survey collects a wide variety of datapoints through interviews, with respondents from all ten provinces of Canada (the three territories are excluded). The data collected regards a single week in the month, usually the week containing the 15th day of the month (Statistics Canada, 2020: Section 3).

Respondents are kept in the surveys for six months, with a sixth of the respondents replaced each month, resulting in a complete change of respondents every six months. Respondent identifiers are not included in the PUMF, so this survey uses the results from March and September for the years of 2015-2017, for a total of six survey periods. Survey months and years are later controlled for in regressions using dummy variables, to counter seasonality. The years are chosen to match the survey period for the latest release of the above-mentioned time use survey. The LFS changed sampling methodology in 2015 (Statistics Canada, 2020: Section 4), hence why the period 2015-2017 was chosen rather than 2014-2016.

This paper uses data from a total of 66,819 respondents with a bachelor's degree and a reported non-zero hourly salary, and non-zero work hours during the last week. The data is subdivided by province, and this study aggregates provinces into five regions (British

Columbia, The Prairies, Ontario, Québec, and the Atlantic Provinces), according to how the time use survey labels provinces.

I have used the following variables from the LFS PUMF (March and September 2015-2017), which are further explained in Statistics Canada (2020).

sex – *Sex of respondent* given as 1 – [male] and 2 – [female] in original data set. These were changed to 0 – [female] and 1 – [male] to make the dummy variable which represents the pay gap residual a more intuitive positive figure.

hrlyearn – Hourly earnings at main job reported in Canadian cents, so the figure is divided by 100 to gain salary figures in Canadian dollars.

uhrsmain – *Usual hours worked per week at main job*, as set out by for instance an employment contract. Reported as multiplied by ten, so uhrsmain was divided by ten to get the figure in a standard hour form.

ahrsmain – *Actual hours worked per week at main job* as reported by the respondent. Reported as multiplied by ten, so ahrsmain was divided by ten to get the figure in a standard hour form.

tenure – *Job tenure with current employer* reported in number of months

educ – *Highest educational attainment*

age_12 – *Five-year age group of respondent* (15 to 19 years, 20 to 24 years … 65 to 69 years, 70 or above). Three groups were dropped to restrict the sample to individuals of working age, 20 to 64 years.

efamtype – *Type of economic family* including information of number of children

lfsstat – *Labour force status* including being employed, employed but absent, unemployed or not in labour force

survyear – *Survey year*, reported as 2015, 2016, or 2017

survmnth – *Survey month*, reported as either 3 or 9 for March or September respectively

prov – *Province* as documented by numerical code. Provinces Newfoundland and Labrador, Prince Edward Island, Nova Scotia, and New Brunswick agglomerated to The Atlantic Provinces, and Manitoba, Saskatchewan, and Alberta agglomerated to The Prairies.

noc_43 – *Occupation at main job*, categorized into 43 classified work tasks, including management, occupations in finance, health, general trades, law, etc.

I created the following binary categorical variables:

child_stat – Whether an individual has a child in their household. Defined as by *efamtype*, where responses 1, 2, 5, 8 and 11 indicate no child in the family

bachelor – whether an individual has received a bachelor's degree. Defined by *educ*, where responses 5 and 6 indicate respondent having a bachelor's degree.

Finally, I calculated a total income variable (**total_income**) by multiplying each individual's usual hourly wages at their main job (**hrlyearn**) by usual weekly work hours at their main job (**uhrsmain**).

3.2 – Methods

3.2.1 – Absence from LFS by Province

From the Labor Force Survey (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b) the amount of individuals absent from work at the time of the survey (defined as respondents with an lfsstat labour force status "employed but absent") are shown as a percentage of the total labor force (defined as the sum of respondents with a labor force status "employed" and "employed but absent"), by province.

3.2.2 – Illustrating Childcare Differences

From the time use survey for 2015 (released 2018), two simple differences for average hours spent per day and participation rate were calculated by subtracting the female value with the male, by age category and region (including Canada as a whole). This yielded two simple measures that in conjunction will be used to map regional differences in childcare. Results are laid out in table 1.

3.2.3 – Pay Gap Residual

A number of linear regression models were run, with the natural logarithm of total income as the dependent variable, controlling for sex as a binary dummy variable, the natural logarithm of tenure with employer, natural logarithm of number of actual hours worked in the last week, dummy variables for each age category, and dummy variables for survey periods. The gender pay gap residual was calculated as the coefficient on the sex dummy variable, where a higher value would indicate a gender pay gap residual more skewed in favor of men. These regressions were run for each province or region individually, by each subsample (pooled, with children, without children) individually.

When mapping changes in the pay gap residual over different age categories, the ten five-year intervals were merged into five ten-year intervals, both for ease of comprehension and to match it to the similar intervals in the time use survey. For all other regressions age categories were calculated as ten five-year intervals.

The LFS PUMF does not contain information on respondents' ethnicities, so it cannot be included in analysis. Since age is reported in intervals of five years, I cannot replicate Goldin's (2014) method of using age as quartic, but instead use a separate dummy variable for each age category. The analysis is restricted to individuals with a bachelor's degree. I do not separately control for educations above a bachelor's degree.

This method means that actual work hours is used as a controlling variable for a dependent variable constructed from usual work hours. This might seem odd at first, but this is explained by how the figures are reported in the Labour Force Survey. Usual work hours are usually reported by a reference document, like an employment contract (Statistics Canada, 2020: Section 3). For people with for instance a standard full time labor contract, this figure would not cover overtime hours, which actual work hours would. Finally, a method very similar to this one was used in Borjas (1980) to deal with the problem of division bias, which will be elaborated upon in the discussion.

3.2.4 – Elasticity of Income with Respect to Work Hours

The elasticity of income with respect to work hours was estimated separately for each province/region and profession by running a linear regression on the natural logarithm of total income on the natural logarithm of number of actual hours worked, controlling for the natural logarithm of tenure with employer and dummy variables for each age category and survey period. A constant elasticity for the sector is then estimated as the coefficient on the natural logarithm of actual hours worked. Unlike the calculation of the pay gap residual, the regressions did not control for sex, nor were the elasticities calculated separately for people with and without children.

3.2.5 – Elasticity-Pay Gap Correlation

The elasticity-pay gap correlation is calculated as the standard Pearson correlation score between elasticity of income with respect to work hours for each sector, with the pay gap residual of the same sector. Calculations are done for each region separately. In the calculation, professions with fewer than 30 employees were removed from the sample.

4. Findings

Table 1. Sample Sizes, LFS March and September 2015-2017

	Percent of Employed Absent from Survey	Number of people in study with children	Number of people in study without children	Total
Canada	10.4%	37691	29128	66819
British Columbia	12.3%	3794	3641	7435
Prairies	7.4%	11236	8562	19798
Ontario	13.0%	11067	8107	19174
Québec	9.0%	6415	4470	10885
Atlantic	10.5%	5179	4348	9527
Source: Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b; author's calculations				

Table 2. Summary Statistics, LFS March and September 2015-2017

Variable description	Hourly Wage	Usual Work Hours	Actual Work Hours	Tenure with current employer
Unit	Canadian Dollars	Hours	Hours	Months
Name of variable	hrlyearn	uhrsmain	ahrsmain	tenure
Number of obs	66819	66819	66819	66819
mean	26.62569	36.56898	36.91416	90.05773
median	23	40	40	61
standard deviation	14.10801	9.769497	12.23248	81.68719
min	3.56	0.2	0.1	1
max	173.08	99	99	240

Table 3. Trends in Childcare Differences, 2015

Region	Age Category	Participation Rate Male	Participation Rate Female	Average hours per day Male	Average hours per day Female	Differences Ptcpt. Rate	Differences Daily Hours
Canada	25 to 34 years	19.8%	41.1%	2.2	3	-21.3%	-0.8
Canada	35 to 44 years	39.6%	51.7%	1.9	2.6	-12.1%	-0.7
Canada	45 to 54 years	14.9%	17.9%	1.6	1.5	-3.0%	0.1
Atlantic	25 to 34 years	23.8%	43.0%	1.8	3.3	-19.2%	-1.5
Atlantic	35 to 44 years	44.1%	44.7%	1.9	2.1	-0.6%	-0.2
Atlantic	45 to 54 years	13.0%	12.7%	1.9	1.7	0.3%	0.2
Québec	25 to 34 years	26.2%	49.5%	2.2	2.8	-23.3%	-0.6
Québec	35 to 44 years	38.4%	54.8%	1.8	2.1	-16.4%	-0.3
Québec	45 to 54 years	14.4%	15.6%	1.4	1.4	-1.2%	0
Ontario	25 to 34 years	16.3%	35.8%	2.6	3.3	-19.5%	-0.7
Ontario	35 to 44 years	39.1%	48.3%	2	2.8	-9.2%	-0.8
Ontario	45 to 54 years	17.3%	18.6%	1.8	1.4	-1.3%	0.4
Prairies	25 to 34 years	18.9%	48.1%	1.7	2.9	-29.2%	-1.2
Prairies	35 to 44 years	37.2%	57.0%	1.8	2.9	-19.8%	-1.1
Prairies	45 to 54 years	12.8%	20.6%	1.6	1.5	-7.8%	0.1
British Columbia	25 to 34 years	19.0%	30.0%	2.1	2.7	-11.0%	-0.6
British Columbia	35 to 44 years	44.6%	51.9%	2	2.2	-7.3%	-0.2
British Columbia	45 to 54 years	12.3%	19.1%	1.4	1.8	-6.8%	-0.4

Source: Statistics Canada, 2018a; author's calculations

Figure 2. Gender Pay Gap Residual by Age Category and Province

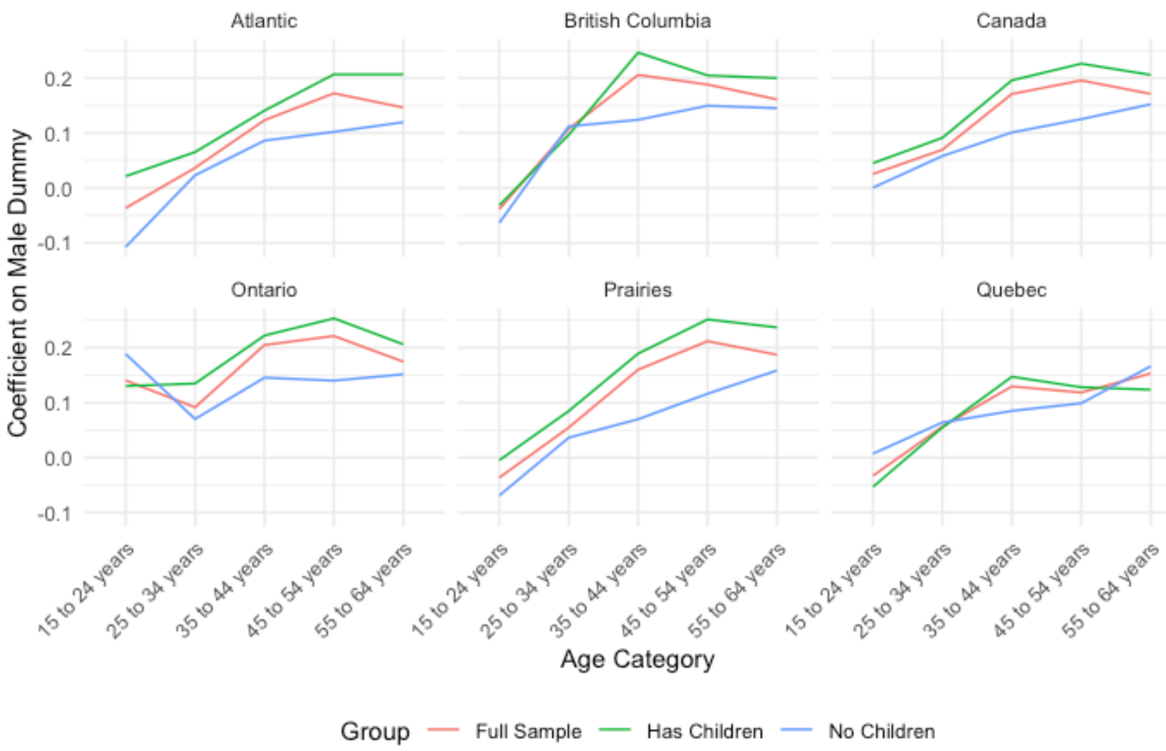

Notes: Sample consists of individuals with a bachelor's degree reporting a positive income and work hours during the sample week between ages 20 to 64 years old. Points in graph are calculations by main task at work. The gender pay gap residual is represented as the coefficient on a male dummy variable in a linear regression with the natural logarithm of total income as the dependent variable. Other controls include natural logarithms of work hours during last week and number of months at current job, and dummy variables for age categories of ten years width, and survey period.

Data source: Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b; author's calculations

Figure 3. National Gender Pay Gap and Elasticity by Occupation at Main Job

Bachelors degree or above, With and Without Children, 2015-2017

Notes: Sample consists of individuals with a bachelor's degree reporting a positive income and work hours during the sample week between ages 20 to 64 years old. Points are calculations made separately by main task at work. Elasticity is calculated as coefficient on the natural logarithm of hours worked last week in a regression with natural logarithm of total income as dependent variable, controlling for the natural logarithm of number of months at current work, age categories of five years width, and survey period. The gender pay gap residual is represented as the coefficient on a male dummy variable in a linear regression with the natural logarithm of total income as the dependent variable, with same controls as for the elasticity, with natural logarithm of hours worked last week added.

Data source: Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b; author's calculations

Figure 4. National Gender Pay Gap and Elasticity by Occupation at Main Job

Bachelors degree or above, With Children, 2015-2017

Figure 5. National Gender Pay Gap and Elasticity by Occupation at Main Job

Bachelors degree or above, Without Children, 2015-2017

Notes for figures 4 and 5:

Sample consists of individuals with a bachelor’s degree reporting a positive income and work hours during the sample week between ages 20 to 64 years old. Points are calculations made separately by main task at work. Elasticity is calculated as coefficient on the natural logarithm of hours worked last week in a regression with natural logarithm of total income as dependent variable, controlling for the natural logarithm of number of months at current work, age categories of five years width, and survey period. The gender pay gap residual is represented as the coefficient on a male dummy variable in a linear regression with the natural logarithm of total income as the dependent variable, with same controls as for the elasticity, with natural logarithm of hours worked last week added. Sample is split by child status in calculations of the gender pay gap, but the overall elasticity is calculated by the pooled sample.

Data source: Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b; author's calculations

Table 4. Correlations Between Elasticity and Pay Gap Residuals by Region, LFS 2015-2017

Region	Canada	British Columbia	Prairies	Ontario	Québec	Atlantic
Has Children	0.193	0.118	0.158	0.331	-0.163	0.071
Pooled Sample	0.041	0.142	0.054	0.23	0.002	0.146
No Children	-0.027	0.046	-0.061	-0.125	0.290	-0.243

Source: Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b; author's calculations

5. Discussion

This section discusses the results found in the graphs and tables of the previous section. It starts with considering the eventual issues with analyzing a dataset constructed from singular weeks. It then discusses three questions in order. First, whether there is a pay gap residual present in Canada and its subnational regions. Second, whether there is any evidence of convex earnings structures impacting the gender pay gap. Third, it investigates whether this correlation differs between people with and without children. Finally, it highlights some potential avenues for improving upon the study.

5.1 – Potential Issues with Weekly Data

The underlying data is being reported for single weeks, not annually or monthly. Individual identification codes of the LFS are not released with the public use microdata files (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b). As a result, this analysis is done for six individual weeks, each six months apart, across three different years. This can be compared to other studies that use data from months (Destefanis et al. 2023) or full years (Goldin, 2014). This results in two main branches of criticism that must be understood when interpreting the results.

Firstly, since the analysis only contains individuals with positive reported work hours, individuals that happened to take that work week off, whether due to maternity leave, illness, or similar, would not be included. The underlying sample could be skewed by bad luck in a specific sample period. However, table 1 shows that while there is some variation between regions, absence from work in all regions hovers fairly evenly around 10% (table 1). Within this area it does not seem that the survey periods used are affected by "bad luck" to the point that it precludes analysis.

The second type of criticism regards size of data values. An analysis based on weekly data will certainly result in smaller figures for both gender pay gap and work hours-income elasticity than for monthly or annual data. There might be a risk that the effects we wish to study do not become visible in such a short time. The pay gap figures do not seem impacted by this in a major way. They tend to be in the range of 0.0-0.2, for a logged variable. This makes the size of the numbers comparable to Goldin (2014) with values roughly ranging between 0 and 0.4, and Destefanis et al. (2023) with values around 0 and 0.25 (though both these studies calculated the gap as coefficients on a female dummy rather than male, making them 0 to negative 0.4 and 0 to negative 0.25 respectively). As regards the elasticities, this study's figures were between 0 and 1, compared to Destefanis et al.'s 0 to 2. In practice, this study's figures were comparable to those of other studies.

However, the use of weekly figures in the underlying dataset could be subject to a division bias. The underlying dataset calculated usual hourly earnings by dividing reported earnings by usual reported work hours (Statistics Canada 2020, Section 3). In his paper 1980 paper, George Borjas spelled out the issues with estimating earnings-work time elasticities specifically, as the appearance of weekly hours both in the dependent and an independent variable could put downward pressure on the elasticity, skewing the results.

Borjas' solution to this included using an alternative, independent measures for work hours. He suggested hours worked last week as an option to usual work hours, which is what this paper has used as an independent variable in elasticity calculations. However, there is a risk that the overlap between an individual's usual and reported hours are not sufficiently distinct from each other. In the underlying sample, of the 66819 respondents (which does not include individuals listing 0 hours worked last week), 45.4% reported a different actual number of hours worked last week than their usual hours (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b, author's calculations). There is still an overlap, but it is far from universal, and likely does not preclude analysis.

It does however mean that the results from this analysis are not directly comparable to those of other studies. Furthermore, beyond just having weekly data, there are other factors at play. Canadian dollars being consistently valued less than a US dollar (Yahoo Finance, 2024) would also skew the figures somewhat. With currency conversion, alongside access to the individual identifiers PUMFs for every month between 2015-2017 to construct figures for the longer term, the analysis would likely be comparable to others like it.

In conclusion, while the data being limited to a single week does limit the analysis, and while it would benefit from more long-term data, weekly data is likely still sufficient for this analysis. There are some caveats to mind, such as potential division biases in elasticity and the potential for some asymmetric non-respondence not considered here. As it stands, the analysis can be taken at face value and its conclusions can be relied upon, but exact values cannot be directly compared to other studies using similar methods.

5.2 – Canadian Pay Gaps and Childcare

Overall there is evidence of gender pay gap residuals being almost universally positive across Canada and its provinces, among age categories and occupations. This goes in line with existing literature on the subject (Fortin, 2019). In this study, gaps are usually in the range of 0.0 - 0.2 (figures 2–5). Except for Québec, gaps follow similar trends across different regions. They start low or even negative, spiking between the age categories 25-34 years and 35-44 years (figure 2), and remain high throughout older age categories. The spike is noteworthy because it roughly coincides with the 2016 average age of first child in Canada of 29.2 (Statistics Canada, 2018c). This spike is also smaller for subsamples without children (figure 2).

Gender pay gaps tend to be lower among people without children, as seen in figure 2 and trendlines in figures 4-5. Considering the crucial spike among people with children between age groups 25-34 years and 35-44 years, four of the regions see an immediate spike to almost at-or-above 0.2. These regions are The Prairies, British Columbia, Ontario, and for Canada at the national level. In terms of childcare these regions have moderately or very high discrepancies in what percent of population partake in childcare by gender, as well as in how time spent on childcare differ across gender (table 3). Conversely, the two regions with a smaller pay gap spike – namely Québec and the Atlantic Provinces (figure 2) – also have overall lower childcare gender discrepancies (table 3).

Québec is an outlier in multiple ways. It almost universally has lower gender pay gap residuals than any other region for any age category (figure 2). Yet, in early stages it has very high childcare gender discrepancy (table 3). In fact, in the crucial 35-44 years age bracket, when the average parents would raise their first child, Québec has both a lower male participation rate and higher female participation rate than any other province, save for the Prairies. Québec's deviation is discussed further in section 5.4.

Overall, we can conclude that there is evidence of a gender pay gap in Canada, and that it is lower for people without children. This suggests that childcare requirements are a major component of the gender pay gap.

5.3 – Evidence of Convex Earnings Structures

There is overall evidence for elasticity-pay gap correlations being positive in roughly half of the regions (table 4). This can be read as evidence for convex earnings structures having a major impact on the gender pay gap residual being limited.

Figure 3 illustrates the elasticity-pay gap relationship for Canada as a whole. There does not seem to be a correlation found either in the figure, or in the Pearson correlation score (table 4). Within individual regions, however, there is more evidence of a correlation. Three out of five subnational regions see a correlation above 0.1, with Ontario having the largest correlation at 0.23. Two regions stand out as having a very small correlation – Québec and the Prairies. Although the scores should not be directly compared to Goldin (2014) or Destefanis, Mazotta and Parisi (2023), it should be noted that these correlations are overall smaller than the ones found in the other studies. The small size of the correlations might not be considered compelling evidence of convex earnings structures having a major impact on the gender pay gap.

One possible cause of this result is our sample arguably being too broad. Goldin (2014) limits her correlation analysis to the professions where men earn the most. Similarly, Bertrand (2018) successively narrows down her analysis to higher and higher paying jobs, and generally finds successively higher elasticities of income with respect to work hours. If a small number of high earning professions have significant gender pay gaps, this could translate into an increased gender pay gap overall. By this logic, being able to show a strong elasticity-pay gap correlation among top earners, as does Goldin (2014), would make a case for convex earnings structures having a crucial impact on professions that would make a disproportionately large impact on the overall gender pay gap.

However, this logic might be a bit too simplistic. Destefanis, Mazotta and Parisi (2023) find a strong elasticity-pay gap correlation for the lower income end of the Italian labor force as well as for the upper end. Furthermore, this paper has restricted itself to college graduates, who can both be generally expected to earn more, and who in Goldin's (2014) view are more susceptible to specific temporal demands. As such, the broadness of the sample should not be overly relied upon as an explanation of the small correlations.

Some individual regions' pooled samples could support the importance of convex earnings structures in the gender pay gap. The two regions that can generally be said to have the highest and lowest gender pay gaps according to figure 2 – Ontario and Québec, respectively, also have the highest and lowest elasticity-pay gap correlation respectively (0.23 and 0.002). British Columbia, a province with comparatively high pay gap residuals (figure 2) and high childcare gender discrepancies (table 3) also has a high elasticity-pay gap correlation (0.142). If this correlation is read as a proxy for convex earnings, these figures support the idea that convex earnings structures are a major factor of the gender pay gap residual.

On the other hand, the remaining three regions offer evidence to the contrary. Canada on a national level, which had higher gender pay gap residual than most of its subnational regions (figure 2) has a very weak elasticity-pay gap correlation (0.041) (figure 3). The Atlantic provinces which had a marginally smaller residual than most other regions (figure 2) had the second highest correlation (0.146) (table 3). Finally, the Prairies, which had very high gender pay gap residuals (figure 2) had practically no elasticity-pay gap correlation whatsoever (0.054) (table 3). This should not be read as convex earnings structures overall not being present, but rather that they do not disproportionately impact high earners to the point of fueling the gender pay gap residual.

5.4 – Convex Earnings Structures and Children

Figures 4 and 5 illustrate the elasticity-pay gap relationship for Canada subdivided by two populations, people without children, and people with. There is a slight but visible correlation between the two variables for people with children in figure 4, shown to be 0.193 in table 4. For people without children the score turns slightly negative at -0.027, practically nonexistent. This suggests that convex earnings structures seem to play a role in the residual gender pay gap for individuals with children, though it does not for people without children. Although the figures on the elasticity-pay gap correlation for the pooled samples were inconclusive, the figures for subpopulations with different child statuses are behaving more according to our hypothesis. We would indeed expect women without children to not be subject to the temporal disturbances associated with childcare.

Four out of five subnational regions had higher correlations for subsamples with children than without, as per expectations (table 4). Furthermore, for the population with children, correlation sizes seem to roughly follow the size of the gender pay gaps found in figure 2. The Atlantic provinces have a lower correlation for people with children, and a comparatively low gender pay gap. British Columbia has a rather low elasticity-pay gap correlation among people with children (table 4), especially considering their rather high gender pay gap residual (figure 2). This could possibly be explained by the relatively high childcare differences visible in table 3. The Prairies, where the differences in pay gap between people with and without children are very high (figure 2), also has one of the largest differences between the groups' elasticity-pay gap correlations, with the correlation scores differing roughly 0.2 (table 3). Finally, Ontario has a very high elasticity-pay gap correlation for individuals with children, a very low one for individuals without (table 3), and high pay gap residuals (figure 2).

Québec consistently acts contrary to almost all other data points for all other regions, opening a discussion on how the province differs. Québec's universal kindergarten system (Kohen et al., 2008) might play a role here, if the policy for instance improved the labor market flexibility of mothers. The data from the time use survey presented in table 3 provides limited support for this, where Québécois women in the crucial 35-44 age bracket spend comparatively little time on childcare. This could support the Cortés and Pan's (2019) arguments about flexibility of the labor force, as the Québécois also consistently have lower pay gap residuals (table 2). This could have associated policy implications. However, it must be considered that the Atlantic provinces achieved a comparable gender childcare difference (table 3) and a 35-44 age pay gap residual similar to Quebéc (figure 2), without its kindergarten system.

This could suggest that the key to reducing the gender pay gap might hinge on the flexibility of employers, as highlighted by Goldin (2014). To make any conclusions about this, a study like Goldin (2014) and Destefanis et al. (2023) on workplace characteristics in Québec, and how they differ from Canada at large, would need to be performed.

5.5 – Limitations of the Study

This study, like those it replicates, should not be read as conclusive statistical proof. For instance, given the importance of overall number of work hours, linear regressions analyzing the impact of sex must include number of work hours when considering the gender pay gap. Barring some instrumental variable, attempts to make statistical inference would thus be subject to considerable endogeneity issues. As such, this paper has not attempted to establish statistical significance, instead relying on descriptive statistics.

I highlighted one of the benefits of a crude elasticity-pay gap analysis being that the method is transparent and is easily compared across countries and regions. To this end, another study might also convert Canadian dollars used in this study, to the US dollar. In that case, however, the survey periods of the underlying datasets must match. This study uses weekly data, so the figures might not be comparable to with studies done on annual data.

Unlike Goldin (2014) and Destefanis, Mazotta, and Parisi (2023), this study considers the main professional tasks of the respondent as opposed to their profession. For instance, an individual is not listed as a mailman, but (presumably) as "Mail and message distribution, other transport equipment operators and related maintenance workers". When considering the number of professions, the broadness of this small amount of categories is almost certainly a step back from the many, more precise categories of Goldin (2014) and Destefanis, Mazotta and Parisi (2023).

That being said, there might be value in studying gender pay gaps as a function of main tasks at work, as opposed to broad profession. The temporal demands of a senior barista at a coffee shop with managerial duties might have more in common with the manager of a small post office, than he does with the newly employed dishwasher. Such nuance could be captured by use of professional tasks, rather than place of occupation. Testing for this possibility would be easy, as Goldin (2014) already did it for professions in her study (p.1098). The same procedure would just need be repeated for main occupations at work, and then be compared to the procedure with professions. Such an analysis would however require a dataset with information on both points, which the Canadian Labour Force Survey PUMF does not provide (Statistics Canada, 2015a, 2015b, 2016a, 2016b, 2017a, 2017b).

6. Conclusion

This paper finds indications that a gender pay gap residual is present among Canadian college graduates across all regions and provinces. This is in spite of controlling for age, tenure, and work hours. In three out of five subnational regions the residual appears correlated with elasticity of income with regards to hours worked, providing mixed evidence suggesting that convex earnings play a role in the overall Canadian gender pay gap residual. For all regions except Quebéc, the elasticity-pay gap correlation was stronger among individuals with children than those without. Furthermore, the sizes of the correlations in different regions among individuals with children loosely correspond to the overall pay gap residual of that region. This suggests that having children plays a role in widening the gender pay gap, due to the temporal constraint it places on mothers.

This study can be considered the first, exploratory half of trying to establish the exact impact of convex earnings in Canada. The conclusions would need to be validated or disproven, either by using a method like Goldin's (2014) analysis of workplace characteristics, or quantitative work to establish statistical inference. This could also help identify whether the impact of convex work hours on the pay gap can chiefly be attributed to a lack of flexibility on behalf of employers, or on behalf of workers, which could have policy ramifications. If the former is true, a policy maker seeking to reduce the pay gap residual should not strive for policies that further promote rewarding long work hours, such as for instance overtime pay. If the latter is true, the focus should instead be on policies that promote flexibility among female workers, such as increased access to kindergarten.

7. References

Angelov, N., Johansson, P., & Lindahl, E. (2016) Parenthood and the Gender Gap in Pay, *Journal of Labor Economics*, 2016, vol. 34, no. 3, pp.545-579, available online: https://www.jstor.org/stable/26553218

Bertrand, M. (2018) Coase Lecture: The Glass Ceiling, *Economica*, 2018 Apr 01. 85(338), pp. 205-231, doi: https://doi.org/10.1111/ecca.12264

Cortés, P. & Pan, J. (2019) When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Skilled Gender Wage Gap, *Journal of Labor Economics*, vol. 37, no. 2 April 2019, pp. 351-398, doi: https://doi.org/10.1086/700185

Destefanis, S., Mazzotta, F., & Parisi, L. (2023). Goldin's Last Chapter on the Gender Pay Gap: An Exploratory Analysis Using Italian Data, *Work, Employment and Society*, 0(0), doi: https://doi.org/10.1177/09500170221143724

Goldin, C. (2014) A Grand Gender Convergence: Its Last Chapter, *American Economic Review*, vol 104, no. 4, pp. 1091–1119, doi: http://dx.doi.org/10.1257/aer.104.4.1091

Pasolli, K. (2015) Comparing Child Care Policy in the Canadian Provinces, *Canadian Political Science Review*, vol. 9, no. 2, pp. 63–78, doi: https://doi.org/10.24124/c677/2015541

Kohen, D., Dahinten, S., Khan, S., & Hertzman, C. (2008) Child Care in Quebec: Access to a Universal Program, *Canadian Journal of Public Health*, vol. 99, no. 6, pp. 451-455, available online: https://www.jstor.org/stable/41995150

Government of Canada, 2023 Early Learning and Child Care Agreements, available online: https://www.canada.ca/en/early-learning-child-care-agreement/agreements-provinces-territories.html [accessed 8 January 2024]

Statistics Canada (2015a), Labour Force Survey: Public Use Microdata File, March 2015, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001xeng [accessed 9 November 2023] (note: all releases of the Labour Force Survey's Public Use Microdata File are found by the same doi link)

Statistics Canada (2015b), Labour Force Survey: Public Use Microdata File, September 2015, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001x-eng [accessed 9 November 2023]

Statistics Canada (2016a), Labour Force Survey: Public Use Microdata File, March 2016, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001x-eng [accessed 9 November 2023]

Statistics Canada (2016b), Labour Force Survey: Public Use Microdata File, September 2016, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001x-eng [accessed 9 November 2023]

Statistics Canada (2017a), Labour Force Survey: Public Use Microdata File, March 2017, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001x-eng [accessed 15 November 2023]

Statistics Canada (2017b), Labour Force Survey: Public Use Microdata File, September 2017, *Statistics Canada*, Public Use Microdata 71M0001X, doi: https://doi.org/10.25318/71m0001x-eng [accessed 15 November 2023]

Statistics Canada (2018a) Daily average time spent in hours on various activities by age group and sex, 15 years and over, Canada and provinces, *Statistics Canada Open Government*, Table 45100014, available online: https://open.canada.ca/data/en/dataset/f3b35173-0cff-4986-9f50-b93118530cc5 [accessed 7 November 2023]

Statistics Canada (2018b) Time use: Total work burden, unpaid work, and leisure, available online: https://www150.statcan.gc.ca/n1/en/daily-quotidien/180730/dq180730a-eng.pdf?st=oA42opuQ [accessed 3 January 2024]

Statistics Canada (2018c) Report on the Demographic Situation in Canada: Fertility: Overview, 2012 to 2016, available online: https://www150.statcan.gc.ca/n1/en/daily-quotidien/180730/dq180730a-eng.pdf?st=oA42opuQ [accessed 21 December 2023]

Statistics Canada (2020) Guide to the Labour Force Survey: 2020, *Surveys and statistical programs – Documentation:* 71-543-G, available online: https://www150.statcan.gc.ca/n1/pub/71-543-g/71-543-g2020001-eng.htm [accessed 9 November 2023]

Yahoo Finance (2024) CAD/USD (CADUSD = X): Historical Data, available online: https://ca.finance.yahoo.com/quote/CADUSD%3DX/history?period1=1063756800&period2=1704672000&interval=1d&filter=history&frequency=1d&includeAdjustedClose=true [accessed 8 January 2024]

8. Appendix

Appendix I. Flexibility in Work Days

Convexity of earnings can also impact workers if they need to leave work for parts of the work day. Consider figures 6 and 7 for a simple example. In figure 6, total earnings for a workday are a convex function of the number of consecutive work hours. There are two workers: worker one works consecutively for nine hours (1-9), and worker two has their day split into two halves, hours 1-4 and 8-12. At the start of their three hour interruption, worker two starts out at their first consecutive work hour, and can never reach the earnings rate that worker one has enjoyed in the second half of his day. The result is a major earnings difference, despite both workers the same number of hours. Compare this to figure 7, where earnings are completely linear. This time, since earnings are purely a function of the number of work hours, worker two does indeed catch up to worker one in their last work hour.

This basic example illustrates that in convex earnings structures, there can be a major punishment associated with taking a mid-day break from work irrespective of the total number of hours. This is one nuance from the convexity mentioned in the main text that would be difficult to capture with elasticity of income with respect to work hours. A precise study of this convexity would likely need to find a method to plot workers' daily schedule hour-to-hour on a larger scale.

Figure 6. Total Earnings With and Without Break with Convex Reward Structures

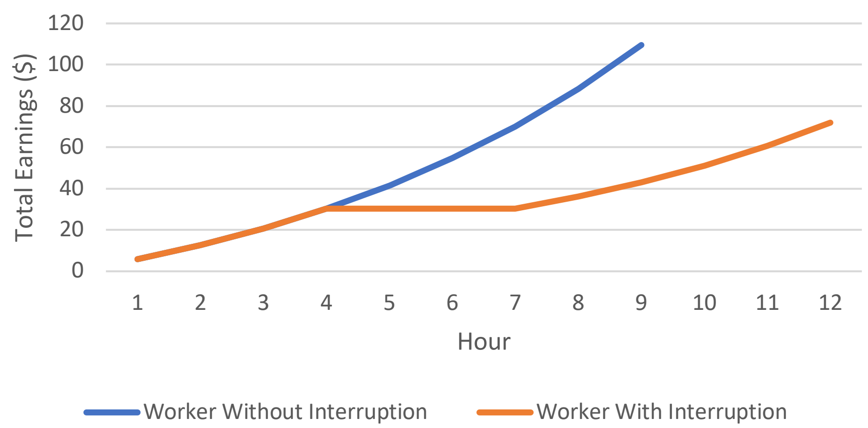

Notes: Total Earnings (E) are the sum of the hourly wage (e), which is an exponential function of number of consecutive hours worked (starting at 1), starting at a base line wage \$5 such that $E = \sum_{i=1}^{n} e_i = \sum_{i=1}^{n} 5^{(1+\frac{i}{10})}$

Figure 7. Total Earnings With and Without Breaks with Linear Reward Structures

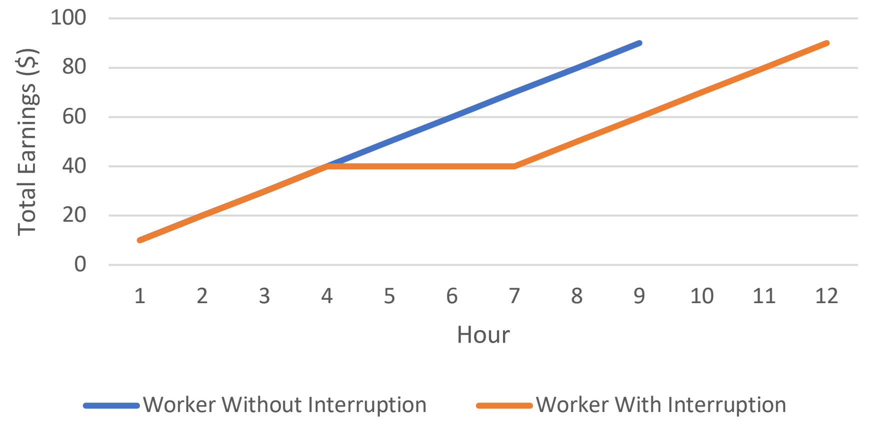

Notes: Total Earnings (E) are the sum of the hourly wage (e), which is a linear function of number of consecutive hours worked (starting at 1), with a fixed hourly wage $10 such that $E = \sum_{i=1}^{n} e_i = \sum_{i=1}^{n} 10n$