

UNIVERSITY OF OTTAWA

DOCTORAL THESIS

Essays on Labour Economics

Author:
Jing CUI

Supervisors:
Dr. Pierre BROCHU
Dr. Jonathan CRÉCHET

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy*

in the

Department of Economics
Faculty of Social Science
University of Ottawa

September 26, 2025

© Jing CUI, Ottawa, Canada, 2025

Declaration of Authorship

I, Jing CUI, declare that this thesis titled, “Essays on Labour Economics” and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: Jing CUI

Date: September 26, 2025

Acknowledgements

Completing this PhD thesis has been a challenging yet rewarding journey, and I could not have done it without the support, guidance, and encouragement of many people.

First and foremost, I am deeply grateful to my supervisor, Professor Pierre Brochu, who guided me from the very beginning and supported me unwaveringly with every detail of my PhD path. His mentorship has been instrumental in shaping both my research and my development as a researcher.

I extend my heartfelt gratitude to my co-supervisor, Professor Jonathan Créchet, for his endless patience, encouragement, and thoughtful guidance. He has been a role model for the kind of person I aspire to become.

I am also deeply thankful to Professor Francesca Rondina, whose door was always open to me. Her constructive suggestions, insightful feedback, and technical support have been invaluable throughout this journey.

My sincere thanks to Professor Minjoon Lee, who first ignited my interest in macroeconomics as my first-year professor. His valuable feedback, technical guidance, and encouragement have profoundly impacted my academic growth.

I am deeply grateful to Professor Eliza Forsythe, whose work has been a major inspiration for my research. Her structured approach to economic questions shaped much of this thesis, and her thoughtful feedback as external examiner was immensely helpful.

I am also grateful to Professors David Gray, Jason Garred and Louis-Philippe Morin for their discussions and insightful feedback, as well as to Professors Yazid Dissou and Sergio Lago Alves, from whom I have learned so much.

A special mention goes to Ruben Hipp for showing me how to communicate effectively with readers and audiences, as well as for the positive energy he always brings.

To my friends and colleagues, Winggie Tian, Paola Jaimes Santamaria, Punardeep Sikka, Pu Sun, Jintao Fu, Yani Zhang, Olivia Jiang, Derek Mikola, Brian Deng, Grace Li, Blurette Arcady Mongoue, Mansoureh Abbasi, and Ekaterina Alekhanova: thank you for the stimulating discussions, collaborative spirit, and encouragement that kept me motivated. I am especially grateful to Gabrielle Li for her constant companionship and support.

I am profoundly thankful to my grandparents, who, even in their eighties, spoke to me every day and listened to my thesis chapters with love and curiosity. To my parents, I owe my deepest gratitude for their unconditional love and support, which have been my foundation through every step of this journey.

Finally, special thanks to my cats, Bee and PrPr, who brought me joy and helped me unwind during stressful times.

This thesis is a testament to the collective effort of all those who have supported me along the way. To each of you, I am forever grateful.

Abstract

The first chapter studies the labour force participation of older individuals during COVID-19. COVID-19 significantly changed the labour participation rates of older Canadians, leading to substantial flows among employment, unemployment, marginal attachment, and non-attachment. Using the Canadian Labour Force Survey (LFS), this paper examines the impact of these flows on the participation rates of older individuals and explores whether COVID-19 prompted early retirements. Unlike the Great Recession, the pandemic caused significant direct separations from employment to non-participation. Additionally, older women experienced slower participation rate recovery than men due to higher outflows and lower inflows. Notably, many individuals who initially became non-attached to the labour force in early 2020 transitioned back to employment in the following months of the same year. Generally, the pandemic did not increase older individuals' self-reported retirement transitions and reduced their probability of staying non-attached to the labour market.

The second chapter examines the cyclicity of worker flows across experience levels in Canada. Using the LFS, I estimate individual monthly transition probabilities over business cycles conditional on labour-market experience and job tenure. The job-finding rate and separation rate are relatively more cyclical for the youth. I find that experience is a major contributor to the cyclical fluctuations in employer-to-employer probabilities, whereas tenure is a major contributor to the cyclicity of employment-to-nonemployment.

The third chapter studies the evolution of the gender unemployment gap in Canada. The gender unemployment gap—defined as women's unemployment rates minus men's unemployment rates—was positive before 1990 but has remained negative since then. I decompose the gender unemployment gap into contributions from gender differences in transition flows between employment, unemployment, and non-participation. The results show that gender differences in flows between employment and non-participation have been positive contributors to the gap over time, while gender differences in employment-to-unemployment flows have been a significant negative contributor. Over the decades, the contribution of flows between employment and non-participation has been decreasing. As employment-to-unemployment flows continue to contribute negatively to the gap, the diminishing contribution of flows between employment and non-participation explains the flip of the gender unemployment gap from positive to negative. Furthermore, I find that differences in industry and occupation composition play a significant role in explaining the gender difference in employment-to-unemployment transition rates.

General Introduction

This dissertation includes three essays on labour economics, particularly focusing on labour market transitions and their implications on participation rate and unemployment rate.

The first chapter, published at the *Journal of the Economics of Ageing*, examines labour force participation and retirement transitions of older individuals in Canada during COVID-19. Existing literature (e.g., [Davis \(2021\)](#); [Montes, Smith and Dajon \(2022\)](#); [Quinby, Rutledge and Wettstein \(2021\)](#)) has found that the pandemic led to a decline in labour force participation and an increase in retirement among older workers in the U.S. In Canada, labour force participation rates for older individuals also dropped sharply at the onset of the pandemic but showed a notable recovery later in 2020. To analyze the participation rate change of older individuals during the pandemic, I map the relationship between labour market stocks and transition rates and illustrate how changes in inflows and outflows influence the labour force participation rate.

Using data from the Labour Force Survey (LFS), I calculate transition rates among different labour market states during the pandemic and construct counterfactual participation rates based on pre-pandemic transition rates. While previous studies on labour market dynamics often rely on steady-state decomposition to assess the effects of inflows and outflows on unemployment rates, this method is less suitable for analyzing participation rate variations. Transitions at the participation margin are relatively low, thus convergence to a steady state can take a long time. To address this limitation, I use a counterfactual approach to disentangle the contributions of different transition flows to participation rate variations.

My findings reveal that the sharp decline in labour force participation in early 2020 was primarily driven by a substantial increase in outflows from the labour force. Moreover, many older individuals exited the labour force directly during COVID-19 without first becoming unemployed and then discouraged. However, beginning in May 2020, substantial inflows from non-participation into the labour force contributed to the subsequent recovery in participation rates among older workers. This pattern contrasts with the Great Recession, when older individuals largely remained in the labour force and many non-participants began actively looking for a job. This chapter shows how different economic shocks can affect transition rates in distinct ways, which leads to different labour force participation outcomes for older individuals. Understanding these dynamics can inform public policy discussions on how to mitigate the impact of future recessions on older workers.

The second chapter examines the cyclicity of worker flows across experience and tenure levels in Canada. [Forsythe \(2022\)](#) shows that in the U.S., young workers' hiring and separation probabilities are more sensitive to business cycles than those of more experienced workers. Several studies, including [Clark and Summers \(1981\)](#), [Jaimovich and Siu \(2009\)](#), [Hoyne, Miller and Schaller \(2012\)](#), [Bredemeier and Winkler \(2017\)](#), [Xu and Couch \(2017\)](#), and [Forsythe and Wu \(2021\)](#), have also found that economic downturns disproportionately impact younger workers. At the same time, younger workers tend to hold lower-tenure jobs, and tenure is

positively correlated with age. [Jung and Kuhn \(2019\)](#) show that tenure distribution explains age-related patterns in labor market outcomes, but the extent to which tenure accounts for the cyclicity of worker flows across age groups remains unclear.

To address this question, I first replicate the empirical findings of [Forsythe \(2022\)](#) in the Canadian context and find that young workers' hiring and separation probabilities are more sensitive to economic downturns in Canada as well. Leveraging the detailed monthly tenure information available in the LFS—which is not collected at a monthly frequency in the U.S. Current Population Survey (CPS)—I then extend the analysis by incorporating tenure into the regression to assess its relative importance compared to experience in explaining worker flow cyclicity.

My findings indicate that experience remains the dominant factor in cyclical fluctuations of hiring, even after controlling for tenure. However, for job separations, tenure plays a more significant role in determining how separation rates respond to business cycles. While previous research has used tenure and experience as predictors of hiring and separation probabilities, this chapter contributes by directly comparing their relative roles in explaining the cyclicity of these transitions.

The third chapter examines the evolution of the gender unemployment gap in Canada. Defined as women's unemployment rates minus men's rates, this gap was positive before 1990 but has remained negative since. [Albanesi and Şahin \(2018\)](#) suggest that the convergence of women's labour attachment explains the closure of the gender unemployment gap in the U.S. Canada has experienced a similar convergence in women's labour force participation toward men's. However, unlike in the U.S., women's unemployment rate in Canada has not only declined to match men's but has remained consistently lower since 1990. This chapter explores why this shift occurred.

To analyze the drivers of this change in the gender unemployment gap, I use a transparent and tractable three-state stock-flow decomposition to quantify the contributions of gender differences in labour market flows to the gender unemployment gap over time. Applying the decomposition method of [Forsythe and Wu \(2021\)](#), I approximate the unemployment rate by its steady-state value—a function of six transition flows—and decompose the gender unemployment gap into the sum of contributions from gender differences in these flows.

The results show that women-men differences in transitions between employment and non-participation have contributed positively to the gender unemployment gap. In contrast, women-men differences in employment-to-unemployment flows have been a significant negative contributor. Over time, the influence of employment-to-non-participation flows has diminished. As employment-to-unemployment flows continued to push the gap downward, the declining impact of employment-to-non-participation transitions explains the shift from a positive to a negative gender unemployment gap. Since 1990, men's higher employment-to-unemployment separation rates have been the primary driver of the negative gap. Further analysis reveals that gender differences in industry and occupation composition, part-time status, and tenure significantly explain the women-men difference in employment-to-unemployment transition rates.

While [Azmat, Güell and Manning \(2006\)](#) finds that the flows between employment and unemployment account for the emergence of gender unemployment gaps in OECD countries with high

gender disparities, this chapter shows that in Canada, it is the sustained negative contribution of employment-to-unemployment transitions, combined with the convergence in women's labour participation, that explains the reversal of the gender unemployment gap. Future work will involve applying this decomposition method to additional countries to compare the evolution of gender unemployment gaps across different labour markets.

Contents

Declaration of Authorship	ii
Acknowledgements	iii
Abstract	iv
General Introduction	v
List of Figures	x
List of Tables	xii
1 Older Individuals' Labour Force Participation	1
1.1 Introduction	1
1.2 Data	5
1.3 An Overview of the State of the Labour Market of Older Individuals	6
1.4 Stocks, Flows and Counterfactuals	8
1.4.1 Two Labour Market States: In and Out of Labour Force	8
1.4.1.1 Method	8
1.4.1.2 Empirical Difficulties	9
1.4.1.3 Results	10
1.4.2 Three Labour Market States: Employment, Unemployment and Out of Labour Force	10
1.4.2.1 Method	11
1.4.2.2 Results	11
1.4.3 Three Labour Market States: In Labour Force, Marginal Attachment and Non-Attachment	12
1.4.3.1 Method	12
1.4.3.2 Labour Attachment of Older Individuals during the Pandemic	13
1.5 Robustness Checks	13
1.6 Differences in Flows and Labour Participation between Older Men and Women	15
1.7 Comparison of COVID-19 with the Great Recession	16
1.8 Retirement Transitions during COVID-19	18
1.9 Conclusion	20
1.10 Figures and Tables	22

2	The Cyclicity of Hiring and Separation by Experience and Tenure	38
2.1	Introduction	38
2.2	Data	40
2.3	Replication of Forsythe (2022)	41
2.3.1	Methodology	41
2.3.2	Results	42
2.3.2.1	New Hires	42
2.3.2.2	Job Separations	45
2.3.2.3	Human Capital and Experience	46
2.4	Tenure	47
2.4.1	Potential Experience and Job Tenure on Job-to-Job Transitions	47
2.4.2	Potential Experience and Job Tenure on Job Separations	49
2.5	Robustness Checks	50
2.6	Conclusion	52
2.7	Figures and Tables	54
3	The Gender Unemployment Gap in Canada	63
3.1	Introduction	63
3.2	Trend and Cyclical Components of Unemployment Gap	65
3.3	Labour Force Compositions and Gender Unemployment Rate Gap	66
3.3.1	Age Composition	66
3.3.2	Education Composition	67
3.4	Decomposing the Unemployment Rate Gap into Difference in Flows	68
3.5	Decomposition Results	70
3.5.1	Contribution of Each Individual Flow	70
3.5.2	Evolution of the Gender Unemployment Gap	71
3.6	Explaining the Gender Difference in EU Flow	72
3.7	Conclusion	73
3.8	Figures and Tables	74
	Bibliography	85
	Appendices	90
A		91
A.1	Panel Creation and Sample Restrictions	91
A.2	Time Path and Steady States of Participation Rate	92
B		93
B.1	Appendix Figures and Tables	93
C		103
C.1	Marginal Contribution of Each Flow on Steady-State Unemployment Rate	103
C.2	Appendix Figures and Tables	106

List of Figures

1.1	Monthly participation rates of individuals aged 55 to 64 of US and Canada, 2017m1-2021m12.	22
1.2	Annual participation rates of age groups 15-54 and 55-64, 1976-2021.	22
1.3	Monthly participation rates of individuals 55-64 during the Great Recession and COVID-19, seasonally adjusted.	23
1.4	Actual vs counterfactual participation rates in I-O two-state.	23
1.5	Actual vs counterfactual participation rates in E-U-O three-state, holding flows between E and O, and U and O.	24
1.6	Actual vs counterfactual participation rates in E-U-O three-state, indirect channels.	25
1.7	Actual vs counterfactual participation rates in I-M-N three-state.	26
1.8	Percentage change in participation rate since February 2020, men and women, flows between I and O.	27
1.9	Percentage change in participation rate since February 2020, men and women, E-U-O three-state, inflows and outflows.	28
1.10	Percentage change in participation rate since February 2020, men and women, I-M-N three-state, inflows and outflows.	29
1.11	Actual vs counterfactual participation rates in I-O two-state, Oct 2008 - Oct 2009.	30
1.12	Actual vs counterfactual participation rates in E-U-O three-state, Oct 2008 - Oct 2009.	31
1.13	Actual vs counterfactual participation rates in I-M-N three-state, Oct 2008 - Oct 2009.	32
1.14	Monthly employment-to-non-attachment and employment-to-retirement rates, 2017m1-2022m6.	33
1.15	The effect of receiving employment benefit on retirement transition probability.	34
1.16	The effect of COVID-19 on retirement transition probability by industry, men and women.	35
2.1	Potential experience profile of hiring rates, by EIER unemployment rate quartiles.	54
2.2	Hiring over the business cycle.	55
2.3	Coefficients of unemployment rate on hiring probability for one-year potential experience bins.	56
2.4	Coefficients of unemployment rate on hiring probability for one-year potential experience bins, with and without college degree.	56
3.1	Gender unemployment rate gap, 1976-2019.	74

3.2	Labour force participation and employment of men and women, 1976-2019. . . .	75
3.3	Trend and cyclical component of unemployment rate of men and women, 1976-2019.	76
3.4	Age-adjusted gender unemployment rate gap, 1976-2019.	77
3.5	Education-adjusted gender unemployment rate gap, 1976-2019.	78
3.6	Monthly labour market transition rates (12-month centered moving average) by gender, 1976m1-2019m12.	79
3.7	Actual and steady-state unemployment rate by gender, 1976q1-2019q4.	80
3.8	Actual and approximated gender unemployment rate gap.	81
3.9	Gender Unemployment Gap Decomposed by Labour Market Flows	81
3.10	Actual and counterfactual gender unemployment rate gap, holding flows between employment and non-participation.	82
3.11	Actual and counterfactual gender unemployment rate gap, holding flows between unemployment and non-participation.	82
3.12	Actual and counterfactual gender unemployment rate gap, holding flow from employment to unemployment.	83
B.1	Detrended hiring rate over the business cycle, comparison between HP filter and Hamilton filter.	93
C.1	Actual gender unemployment rate gap and the proportion explained by individual flow, flows between employment and non-participation.	106
C.2	Actual gender unemployment rate gap and the proportion explained by individual flow, flows between unemployment and non-participation.	107
C.3	Actual gender unemployment rate gap and the proportion explained by individual flow, flows between employment and unemployment.	108
C.4	Actual gender unemployment rate gap and the proportion explained by each individual flow, with the partial derivative of steady-state unemployment rate with respect to each flow evaluated at women's flow rates.	109
C.5	Actual gender unemployment rate gap and the proportion explained by each individual flow, with the partial derivative of steady-state unemployment rate with respect to each flow evaluated at the mid-point flow rates.	110

List of Tables

1.1	Transition Rates before and during COVID-19	36
1.2	Probability of Returning to Employment, April-September Panel	37
2.1	Hiring over the Business Cycle: With and Without Controls	57
2.2	Hiring over the Business Cycle: Young and Experienced	58
2.3	Exits and Other Flows	58
2.4	Involuntary and Voluntary Separations to Unemployment	59
2.5	Hiring over the Business Cycle: Age Versus Education	59
2.6	Job-to-Job Transition over the Business Cycle: Experience and Tenure	60
2.7	Exits and Other Flows: Tenure Groups	61
2.8	Exits and Other Flows: Experience & Tenure Groups	62
2.9	Summary of Experience vs. Tenure Contributions to Flow Cyclicity	62
3.1	Gender Difference in Employment-to-Unemployment Transition Rate	84
B.1	Joint Distribution of Experience and Tenure Groups: Employed Individuals	94
B.2	Joint Distribution of Experience and Tenure Groups: Non-Employed but Worked in the Past Twelve Months	94
B.3	Hiring over the Business Cycle: With and Without Controls (Provincial)	95
B.4	Hiring over the Business Cycle: Young and Experienced (Provincial)	96
B.5	Exits and Other Flows (Provincial)	97
B.6	Involuntary and Voluntary Separations to Unemployment (Provincial)	98
B.7	Hiring over the Business Cycle: Age Versus Education (Provincial)	99
B.8	Job-to-Job Transition over the Business Cycle: Experience and Tenure (Provin- cial)	100
B.9	Exits and Other Flows: Tenure Groups (Provincial)	101
B.10	Exits and Other Flows: Experience & Tenure Groups (Provincial)	102
C.1	Gender Difference in Employment-to-Unemployment Transition Rate - Without Occupation	111
C.2	Proportion of Gender Unemployment Gap Explained by Each Individual Flow . .	111

*To my parents and grandparents, whose unconditional love and support
have made this journey possible.*

Chapter 1

Older Individuals' Labour Force Participation

1.1 Introduction

COVID-19 was an unprecedented shock that has brought enormous economic damage worldwide and dramatically changed peoples' lives. A sizable amount of literature has examined the impacts of COVID-19 from various perspectives. For instance, it has altered the production in different industries and affected labour productivity, as well as domestic consumption patterns (e.g., [Slade \(2022\)](#); [Blit, Skuterud and Veall \(2020\)](#); [Noda and Teramoto \(2024\)](#)). The pandemic's specific hygiene requirements, such as ventilation, and increased working from home, have also influenced energy consumption and greenhouse gas emissions (e.g., [Villeneuve, Abdeen, Papineau, Simon, Cruickshank and O'Brien \(2021\)](#); [Morissette, Deng and Messacar \(2021\)](#)). Additionally, COVID-19 has impacted community engagement and long-term care management (e.g., [Cattapan, Acker-Verney, Dobrowolsky, Findlay and Mandrona \(2020\)](#); [Pue, Westlake and Jansen \(2021\)](#)). In terms of labour market outcomes, studies have found a dramatic decline in aggregate working hours and the probability of being employed during the pandemic (e.g., [Lemieux, Milligan, Schirle and Skuterud \(2020\)](#); [Béland, Brodeur, Mikola and Wright \(2022\)](#); [Brodeur, Gray, Islam and Bhuiyan \(2021\)](#)).

Between February and April 2020, there was a 5.3 percentage point increase in the transition rate from the labour force to non-participation for older individuals in Canada. Although the outflow rate decreased after April 2020, it remained higher than pre-pandemic levels as of June 2022. Similar patterns were observed in the US, where COVID-19 led to a decrease in labour participation and an increase in retirements among older individuals (e.g., [Davis \(2021\)](#); [Quinby, Rutledge and Wettstein \(2021\)](#); [Montes, Smith and Dajon \(2022\)](#)).

However, the labour participation patterns of individuals near retirement age in the US and Canada differed during the pandemic. As shown in [Figure 1.1](#), the participation rate for individuals aged 55 to 64 in the US dropped at the onset of the pandemic and remained below pre-pandemic levels through the end of 2021. [Montes et al. \(2022\)](#) find that excess retirement

accounted for a large part of this decline, indicating that the pandemic led many older individuals to retire earlier than they might have otherwise.¹ In contrast, Canada experienced a different trend. After a sharp drop in the participation rate of older individuals in early 2020, the rate recovered to pre-pandemic levels within six months.

The substantial inflows from non-participation to the labour force since May 2020 explain the recovery in labour participation among older individuals in Canada. In the US, the flow from non-participation to unemployment saw a brief increase in May and June 2020, but then gradually decreased; the flow from non-participation to employment did not increase after May 2020.² COVID-19 has induced large and persistent flows in and out of the labour force in Canada. Understanding the dynamics of these flows could inform discussions on public policy aimed at mitigating the impact of similar shocks on older workers in the future. For instance, during the pandemic, public policy initiatives such as the Canada Emergency Wage Subsidy (CEWS) and Canada Emergency Response Benefit (CERB) appear to have supported older workers by mitigating their outflows from employment and unemployment to non-participation.³

In this paper, I map the relationship between labour market stocks and transition rates to illustrate the impulse responses of the labour force participation rate to changes in transition rates.⁴ Using the Labour Force Survey (LFS) data, I calculate the transition rates of older individuals among various labour states during COVID-19 and construct counterfactual participation rates based on pre-pandemic transition rates.⁵ These counterfactuals help gauge the relative importance of different flows in explaining the participation rate patterns observed for older individuals since the onset of COVID-19. As the effects of flow changes are cumulative, counterfactual participation rates effectively illustrate the impulse responses of older individuals' labour force participation to the COVID-19 shock.

Additionally, I take advantage of the rotating panel design of the LFS to construct six-month panels that follow the labour participation states of the same individuals during the pandemic. By examining self-reported retirement transitions, the study explores the retirement trajectories of older individuals, investigating whether COVID-19 affected the retirement probabilities of

¹Montes et al. (2022) find that in the US, while there was an increase in excess retirement among individuals aged 55 to 64, these excess retirements were concentrated among cohorts aged 65 and older at the start of the pandemic. The retired population is older in the US than in Canada, possibly because the full pension age and the early pension claiming age are higher in the US. In the US, for individuals born in 1960 or later, the full retirement age is 67. For those born before 1960, the full retirement age ranges from 65 to 66 and a few months, based on the specific year of birth. In Canada, the full eligible age for social security, such as Old Age Security (OAS), Guaranteed Income Supplement (GIS), and Canada and Quebec Pension Plans (CPP and QPP), has been 65 since 1967.

²These observations are based on the IPUMS Current Population Survey (CPS) data.

³In the US, the Paycheck Protection Program (PPP) played a similar role, boosting employment at eligible firms by between 2% and 5% as of May 2022 (Autor et al., 2022). A systematic comparison of policy differences—such as the design of emergency income support programs—falls beyond the scope of this study but could help explain differences in self-reported labour market status across countries.

⁴Changes in transitions among different labour states play a crucial role in explaining employment and labour participation fluctuations during recessions and subsequent recoveries (Gerbery and Miklošovič, 2020).

⁵COVID-19 has affected both labour supply and demand for older individuals. These shifts in supply and demand are reflected by the changes in observed transition rates, thus impacting older individuals' labour participation. Thus, I construct counterfactual participation rates using pre-pandemic level transition rates to mimic the labour market conditions without COVID-19 shock.

older men and women differently across various industries and whether receiving employment benefits during the pandemic influenced these retirement transitions.

Through counterfactual analysis, I identify the flows that drive labour force participation changes among older individuals during the COVID-19 pandemic. The results show that outflows from the labour force primarily explain the initial drop in participation rates, while significant inflows since May 2020 explain the subsequent recovery. Notably, the outflows in March and April 2020 were mainly from employment, with many individuals transitioning directly from employment to non-participation without going through unemployment. The increased inflows since May 2020 were mostly flows to employment. Additionally, COVID-19 led to substantial flows within the labour force, i.e. extensive flows between employment and unemployment from the spring to fall of 2020. However, these changes did not significantly affect the overall labour force participation rate.

COVID-19 affected older men and women differently. From February to April 2020, older women's participation rate dropped by 8.1 percent, taking almost a year to recover, while older men's participation rate dropped by 4.4 percent and recovered within five months. This difference is attributed to higher outflow rates (from February to October 2020) and lower inflow rates for older women compared to men. In terms of outflows from the labour force, higher flows from employment to non-participation for women explain the gender gap in participation rate changes. Regarding inflows to the labour force, both lower non-participation-to-employment and non-participation-to-unemployment rates for women contributed to the gender gap. However, flows in and out of marginal attachment did not explain the gender gap in change in participation rates.⁶

Besides finding that older women flowed out of the labour force more than men did, this paper also finds that COVID-19 increased the retirement transition probability for older women in industries such as mining, financial services, educational services, and arts and entertainment. Conversely, it decreased retirement transition probabilities for men in industries such as utilities, manufacturing, wholesale trade, transportation, finance and housing services, educational services, and hospitality. Receiving employment benefits during the pandemic reduced the retirement transition probability for individuals aged 55 to 59, with no significant difference between men and women. However, receiving employment benefits did not significantly affect the retirement probability of those aged 60 to 64. For men aged 65 to 69, receiving employment benefits reduced the retirement transition probability, while for women in the same age group, it increased the probability.

During the Great Recession, older individuals did not leave the labour force. Instead, many non-participants began actively looking for jobs during the recession. This aligns with findings in the US that the wealth loss during the Great Recession makes individuals postpone retirement

⁶In this paper, non-participation is divided into two categories: marginally attached to the labour market and non-attached to the labour market. While 'discouraged workers' typically refers to individuals not actively searching for work due to the perceived lack of opportunities, 'marginally attached workers' is a broader term. It includes individuals who want a job and may or may not be available to start soon, capturing a wider range of situations than just those who are discouraged. The identification of marginal and non-attachment in this paper conceptually follows Jones and Riddell (1999, 2006, 2019).

(e.g., [Gustman, Steinmeier and Tabatabai \(2010\)](#); [McFall \(2011\)](#)). In contrast, COVID-19 drove many individuals out of the labour force. Although the employment-to-non-attachment transition rate increased during the pandemic, the employment-to-retirement transition rate for near-pension-age individuals did not increase. Interestingly, many individuals transitioning to retirement or non-attachment during the early pandemic returned to employment later in 2020, indicating a possible reversal of their initial decision.

This paper contributes to the existing literature in three key ways. First, it adds to the body of work examining COVID-19's impact on the labour market. While much of the Canadian literature focuses on the effects of lockdown policies on employment, unemployment, and labour force stocks for different demographic groups (e.g., [Lemieux et al. \(2020\)](#); [Béland et al. \(2022\)](#)), few studies address COVID-induced labour flows and their implications (e.g., [Brochu, Créchet and Deng \(2020\)](#); [Jones, Lange, Riddell and Warman \(2021\)](#)). The net statistics of job losses and unemployment rates could sketch the adverse effect of COVID-19, but cannot reflect the exact motions of the working force and the shock that COVID-19 has brought to the natural capacity of the job market. This study complements the small body of literature on COVID-19 labour dynamics by investigating how older individuals' labour attachment changed in response to the pandemic.⁷ Additionally, I decompose the change in the participation rate by changes in stocks and flows, and build counterfactuals to show the accumulated effects of changes in flows on participation rate since the onset of the pandemic, i.e. how a change in transition rate in March affects the participation rate instantaneously and how this effect is transmitted to the following months.

Second, this study contributes to the literature on retirement. Prior to COVID-19, research on retirement decisions was well-documented, both theoretically and empirically. Theoretical life cycle models depict a representative individual solving a dynamic optimization problem to choose an optimal retirement age, i.e. the individual chooses consumption and leisure for each period to maximize her/his lifetime utility (e.g., [Gustman and Steinmeier \(1986\)](#); [French \(2005\)](#)). Empirical models estimate the effects of factors such as private savings, public pensions, total assets, health conditions, and life expectancy on retirement decisions (e.g., [Baker and Benjamin \(1999\)](#); [Gustman, Steinmeier and Tabatabai \(2010\)](#); [Finnie and Gray \(2018\)](#); [Milligan and Schirle \(2018b\)](#); [Munnell, Sanzenbacher and Rutledge \(2018\)](#)). In this paper, I add a new perspective by examining aggregate flows between working and non-working states for the older age group since the onset of the pandemic. Furthermore, this paper complements the literature on reverse retirement (e.g., [Smeaton, Di Rosa, Principi and Butler \(2018\)](#)) with evidence in a short-term framework, capturing transitions from retirement back to employment within six months using the LFS six-month panel.⁸

Third, this study relates to the literature examining how labour market transitions, such as job finding and job separation rates, impact labour market stocks (e.g., [Elsby, Michaels and Solon \(2009\)](#); [Campolieti \(2011\)](#); [Shimer \(2012\)](#); [Elsby, Hobijn and Şahin \(2015\)](#); [Elsby, Hobijn,](#)

⁷[Bédard and Michaud \(2021\)](#) find that this group of individuals is financially vulnerable, thus, might be less tolerant of adverse economic shocks.

⁸A quarter is considered sufficient time to observe labour state transitions for older individuals ([Blau, 1998](#)), making six-month panels effective in capturing retirement and re-employment transitions.

Karahan, Koşar and Şahin (2019)). Typically, this literature uses steady-state decomposition to analyze these effects. However, transitions at the participation margin are relatively low, and convergence to a steady-state can take a long time, making this technique less suitable for studying participation rate variations.⁹ Understanding how the type of shock and context will matter for how people near retirement will flow through the labour market. I provide a new technique that uses counterfactual analysis to disentangle the participation rate variation.

The rest of the paper is organized as follows. Section 1.2 discusses the data, while section 1.3 provides an overview of participation and transition rate changes during the pandemic. Section 1.4 presents the empirical strategy and results, with robustness checks in section 1.5. Section 1.6 compares the different impacts of COVID-19 on older men and women. Section 1.7 examines how flow changes affected older individuals' participation during the Great Recession and compares it to COVID-19. Section 1.8 discusses retirement transitions during the pandemic. Finally, section 1.9 concludes.

1.2 Data

I utilize the master files of the Canadian Labour Force Survey (LFS) to study how older individuals' labour force participation responds to the COVID-19 shock. The LFS is a large monthly survey (interviewing approximately 54,000 households per month) that gathers extensive labour force information of Canadian residents.¹⁰

One of the key characteristics of the LFS is its rotating panel design. Households are surveyed for six consecutive months, and one-sixth of the sample is replaced every month. With this design, I construct two-month panels to calculate the monthly transition rates and six-month panels to follow an older individual's retirement trajectory.¹¹

Moreover, the timeliness of the LFS and the high frequency of its data collection are important for my study. The LFS has a tight production timeline; it takes slightly less than four weeks between the data collection and its final release. The timeliness makes the LFS a critical tool for understanding the labour market in these uncertain pandemic times. The high frequency of the LFS data allows for the identification of the source of the dramatic changes in the participation rate that are observed from one month to the next.¹² Also, the rich labour market information and the large sample size are two additional advantages of the LFS. For example, given the detailed reasons for job leaving, voluntary retirement transitions can be identified. The large

⁹For more details on the convergence of the participation rate to its steady-state, please see Appendix A.2.

¹⁰Non-response has substantially increased in COVID-19 times, which has resulted in smaller monthly samples. Having said this, the samples remain quite large. Although this rise in non-response has raised some concerns about the quality of the data, Brochu and Cr chet (2022) find that the LFS is nevertheless a reliable data source for analyzing the economy during COVID-19.

¹¹I construct the panels with the unique household identifier and within-household person identifier. See more details on panel data construction in Appendix A.1

¹²For most data typically used in the retirement literature, the frequency of data collection is not very high; thus, the immediate response of older individuals to the COVID-19 shock and the development of the responses cannot be tracked. For example, the data collection frequency of the Retirement History Longitudinal Survey (RHLS) and the Health and Retirement Study (HRS) is every two years, whereas for the Census it is every five years.

sample sizes of the LFS are necessary, as my analysis focuses on a narrow age group (i.e. 55 to 64 year olds).

The sample is restricted to civilian Canadian residents who are 55 to 64 years of age. The age restrictions reflect the fact that the study focuses on the labour force participation of the individuals who are close to full public pension eligible age.¹³ Although my focus is on the pandemic period, I require the use of LFS data from January 2017 onward, as my empirical strategy necessitates the use of some pre-COVID years. Further, I require balanced two-month panels for the estimation of monthly transitions. As such, individuals who are not present in both months are dropped. Approximately one-sixth of respondents are part of the outgoing rotation as of the first month of the two-month panels and, as such, are not observed in the second month.¹⁴

1.3 An Overview of the State of the Labour Market of Older Individuals

Figure 1.2 illustrates the annual participation rate for individuals aged 55 to 64 and those aged 15 to 54 from 1976 to 2021. The participation rate is defined as the share of the population that is either employed or actively seeking work. The participation rate of individuals aged 55 to 64 shows clear distinctions in both levels and secular trends from the younger age group. The participation rate for individuals aged 55 to 64 has consistently been lower than that of those aged 15-54. However, since the mid-1990s, there has been a notable upward trend in the participation rate of older individuals. During the recessions of the 1980s and 1990s, the participation rate for those aged 55 to 64 declined. The magnitude of the drop was relatively small and does not appear substantial upon visual inspection. However, during the Great Recession, the participation rate for this age group did not decrease, possibly masked by the ongoing upward trend since the mid-1990s. The most recent recession, COVID-19, led to a drop in participation rates for individuals aged 55 to 64. Nevertheless, the annual rates may not reflect exactly the change in participation rate during this short-lived recession; examining monthly rates is necessary to better understand the impact of COVID-19.

Figure 1.3 presents the seasonally adjusted monthly participation rates for individuals aged 55 to 64, focusing on the most recent two recessions.¹⁵ The figure shows that COVID-19 brought a significant shock to the participation rate for individuals aged 55 to 64 in a short period, unlike during the Great Recession, when the participation rate remained relatively stable. At the onset of COVID-19, the participation rate was 67.1 percent, but it dropped to 62.7 percent within two months—a 4.4 percentage point decline. Notably, the participation rate rebounded quickly,

¹³In the LFS, labour force information for individuals 70 and up that is gathered in the birth interview is carried forward to the subsequent months in the survey. As a robustness check, I expand the age range of my sample to cover individuals aged from 55 to 69; the main conclusions still hold.

¹⁴The lack of two consecutive months of data can also occur due to changes in the eligible population, non-response, or individuals joining/leaving a surveyed dwelling partway through the six-month window. A change in the eligible population can, for example, result from immigration—a person migrated to Canada and joined the surveyed dwelling in the second month of the panel.

¹⁵Given that the two recessions do not start (nor end) in the same months, seasonality would be a confounding factor when comparing recessions. As such, I abstract from this concern by presenting seasonally adjusted rates.

reaching 66.2 percent six months after the pandemic began and returning to pre-COVID-19 levels after another two months, within the same year the pandemic started.

Table 1.1 presents periodical transition rates prior to and during the pandemic. Panel A shows transition rates among the three traditional labour market states: employment (E), unemployment (U), and out of the labour force (O).¹⁶ Panel B offers an alternative perspective by subdividing the non-participation group into two categories: marginally attached (M) and non-attached (N) individuals. I define marginally attached individuals as those who are not currently in the labour force but express a desire to work. In contrast, non-attached individuals are those who are not in the labour force and do not want a job. Column (1) shows the transition rates from February 2017 to January 2020, whereas columns (2) and (3) are for February 2020 to January 2021 and February 2021 to January 2022, respectively.¹⁷ Column (4) shows the F statistics that transition rates during Feb 2020 - Jan 2021 are statistically significantly different from the ones during Feb 2017 - Jan 2020. Column (5) shows the F statistics that transition rates during Feb 2021 - Jan 2022 are statistically significantly different from the ones during Feb 2017 - Jan 2020.

Several observations stand out from Panel A, column (1): first, transition rates from employment and non-participation tend to be much smaller than those originating from unemployment; second, conditioning on transiting from employment, non-participation is a more prevalent destination than unemployment for the age group 55 to 64; and third, those transitioning from non-participation are more likely to transition to employment directly instead of going through unemployment. From Panel B of column (1), one sees that marginally attached individuals are significantly more likely to transition into the labour force than non-attached individuals.

Column (2) consists of the first year of COVID-19. The transition rates in 2020 are much higher than those of the previous three years, which is evidence of intense labour market churning induced by COVID-19 (e.g., the substantial transition between employment and unemployment). Considering the size of the employed group, the effect of a 1.4 percentage point increase in the employment-to-unemployment transition rate on the employment rate is not negligible. Besides, there are more outflows from non-participation to both employment and unemployment in 2020, as compared to the three pre-COVID years.

Column (3) shows that, in 2021, most transition rates have returned to their pre-pandemic levels, with some exceptions. The transition from the labour force to marginal attachment decreases slightly but remains higher than its pre-COVID level. Additionally, the transitions from non-participation to both employment and unemployment stay high, as do the transitions from non-attachment to the labour force and marginal attachment. This indicates that older individuals tended to return to work after the pandemic.

¹⁶These states are defined according to the standard Labour Force Survey (LFS) classifications: individuals are considered employed if they performed any paid work or had a job from which they were temporarily absent; unemployed if they were without work but actively searching and available to start; and not in the labour force (non-participants) if they were neither employed nor unemployed.

¹⁷The changes in transition rates of the two subgroups (55-59 and 60-64) show a similar pattern before and during the pandemic, though some changes have larger magnitudes for the age group 60-64 (e.g., the transition from the labour force to non-attachment).

1.4 Stocks, Flows and Counterfactuals

In this section, I explore the source of change in the participation rate of individuals aged 55 to 64 during the pandemic. I take advantage of the fact that the participation rate can be expressed as a function of inflow and outflow rates and past labour market stocks. As such, I can construct counterfactual participation rates where I hold different flow rates constant at pre-COVID levels. This will allow me to gauge the relative importance of inflows and outflows in explaining the participation rate patterns observed for older individuals since the onset of COVID-19. I start my analysis by assuming only two labour market states (in and out of the labour force).

1.4.1 Two Labour Market States: In and Out of Labour Force

1.4.1.1 Method

Assume a discrete-time framework, in which an individual in period t can be in one of the two labour market states: in the labour force (I) or out of the labour force (O). Furthermore, assume no change over time in the eligible population. Let λ_t^{JK} represent the transition rate of going from state J in period $t-1$ to state K in period t . Assume λ_t^{JK} follows a Markov process, such that the transition probability depends only on the state of the past period. As such, the total number of individuals in the labour force as of period t can be expressed as:

$$I_t = I_{t-1} + O_{t-1} \cdot \lambda_t^{OI} - I_{t-1} \cdot \lambda_t^{IO}, \quad (1.1)$$

where I_t and O_t are the total number of individuals in and out of the labour force at time t , respectively. The second right-hand side term of equation (1.1) represents inflows into the labour force, whereas the third represents outflows. Dividing both sides of equation (1.1) by the population, one can write the participation rate, pr_t , as a first-order difference equation:

$$pr_t = pr_{t-1} + (1 - pr_{t-1}) \cdot \lambda_t^{OI} - pr_{t-1} \cdot \lambda_t^{IO}. \quad (1.2)$$

Although the size of the transition rates matter, equation (1.2) shows that their impact depends on the past period's participation rate, i.e. the size of the stock.¹⁸ Recursively solving equation (1.2), one can express the participation rate as:

$$\begin{aligned} pr_t &= \lambda_t^{OI} + \sum_{i=1}^{t-1} [\lambda_{t-i}^{OI} \cdot \prod_{i=1}^{t-1} (1 - \lambda_{t-i+1}^{IO} - \lambda_{t-i+1}^{OI})] + \prod_{i=1}^t (1 - \lambda_{t-i+1}^{IO} - \lambda_{t-i+1}^{OI}) \cdot pr_0 \\ &= f(pr_0; \boldsymbol{\lambda}^{OI}; \boldsymbol{\lambda}^{IO}), \end{aligned} \quad (1.3)$$

where $\boldsymbol{\lambda}^{OI} = (\lambda_1^{OI}, \lambda_2^{OI}, \dots, \lambda_t^{OI})$ and $\boldsymbol{\lambda}^{IO} = (\lambda_1^{IO}, \lambda_2^{IO}, \dots, \lambda_t^{IO})$. The fact that it can be written as a function of its initial stock and two sequences of transition rates is informative, as this

¹⁸For example, in the three years prior to the onset of COVID-19, twice as many individuals aged 55 to 64 were in the labour force than were out. Then, a small change in λ^{IO} will have a larger impact on the participation rate than a small change in λ^{OI} .

implies that a shock to the transition rates in one period can impact the participation rate multiple periods down the road.

One can determine the impact of changes in inflow and outflow rates by constructing counterfactuals where one keeps a set of transition rates at pre-COVID-19 levels, i.e.

$$pr_t^{\overline{IO}} = f(pr_0; \lambda^{\overline{IO}}; \lambda^{OI}) \quad (1.4)$$

$$pr_t^{\overline{OI}} = f(pr_0; \lambda^{IO}; \lambda^{\overline{OI}}). \quad (1.5)$$

Equation (1.4), for example, will help determine the importance of outflows in explaining the participation rate patterns observed during the pandemic.

1.4.1.2 Empirical Difficulties

Empirically, a key challenge lies in measuring transitions accurately due to the mismatch between flow data and cross-sectional estimates, a discrepancy commonly referred to as the margin error problem (e.g., [Abowd and Zellner \(1985\)](#); [Elsby et al. \(2015\)](#)). In this two-state world, it means that if one were to add up, using two-month panel data, the number of individuals that transition from out of the labour force (in period $t - 1$) to in the labour force (in period t) to those that remain in the labour force in both periods, it would not match the size of the labour force observed in the period t cross-section. Such a situation arises because one does not observe the same individuals over consecutive months in cross-sectional data. The survey design of the LFS implies that one can only follow (at best) 5/6th of the sample from one month to the next. As such, one cannot exactly recover the estimated participation rate using equation (1.2). This issue becomes more pronounced if the estimated transition rates are biased upward or downward due to measurement error, as such biases can accumulate over time (as presented in equation (1.3)). To address the margin error problem, I rely on the method proposed by [Elsby et al. \(2015\)](#). It adjusts the estimated transitions as to ensure that the flow data matches the cross-sectional data.¹⁹

Another difficulty in operationalizing the counterfactual analysis lies in determining what the transition rates would have looked like if the pandemic had not occurred, i.e. how to find a reference transition rate to use in the hypothetical world. To construct my counterfactuals, I take advantage of a few stylized facts: the monthly transition rates in the years prior to the pandemic (2017-2019) show no secular trend. In fact, apart from seasonal variations, the transition rates are quite stable over this time period. Therefore, one can use the transition rates of the same month, but of the year prior to the pandemic, to construct the counterfactual rates.²⁰

¹⁹More precisely, it is a minimum distance estimator which minimizes the weighted distance between adjusted and estimated rates, subject to the constraint that the flow data matches the cross-sectional data. In the robustness checks section, I carry out counterfactual analyses without making this margin error adjustment. The main conclusions remain unchanged.

²⁰As a robustness check I expand the number of pre-COVID years used to construct the counterfactual transitions. My main findings remain unchanged.

1.4.1.3 Results

Figure 1.4 compares the monthly actual participation rate of 55 to 64 year olds to the counterfactual ones that hold the outflow/inflow rates constant at pre-COVID-19 levels, for February 2020 to December 2020 period. The actual participation rate is estimated using cross-sectional data. The counterfactual is calculated using equations (1.4) and (1.5). The monthly transition rates are estimated using two-month panels and then adjusted to account for the margin error problem. The starting point of the counterfactual rate, i.e. February 2020 participation rate, is estimated using February 2020 cross-sectional data.

As shown in Figure 1.4, the large drop in the participation rate during the first two months of the pandemic is entirely driven by increased outflows. A significant gap emerges between the actual and counterfactual participation rates starting in March 2020, reaching 7.3 percentage points by June 2020. In contrast, the counterfactual participation rate that holds inflow rates constant at pre-pandemic levels closely tracks the actual rate before April 2020, indicating that inflows played no role in the initial participation rate decline.

A significant gap does emerge between the actual participation rate and counterfactual rate that holds λ^{OI} constant starting in May 2020, implying that inflows become an important factor as of early summer. This explains why there has been a dramatic rebound in the participation rate since this time. Notably, outflow rates as of May and June 2020 are still higher than those observed in the same months of the previous year, but the difference is now much more muted, and as such, the dramatic increase in inflows dominates.

It is important to recognize that it may take time for the distance between the actual participation rate and its counterfactual to converge, even if the flow shock was short-lived. For example, the gap between the actual and counterfactual rate that holds λ^{IO} constant decreases in the second half of 2020, but at a slow rate, despite outflow rates going back to pre-COVID levels. This holds true because both inflow and outflow rates are typically low, which means that it takes time for the participation rate to converge to its steady-state.²¹

1.4.2 Three Labour Market States: Employment, Unemployment and Out of Labour Force

A key finding of the two-state analysis is the presence of large outflows early in the pandemic. The critical question is whether these outflows originated primarily from employment or unemployment. Older individuals may have exited the labour force directly from employment, or through an intermediate step of unemployment. Similarly, the substantial inflows observed

²¹To better illustrate this issue, I carried out the following experiment. I construct another two counterfactual participation rates, holding both the inflow and outflow rates constant. It ensures the existence of a time-invariant steady-state and eliminates the effect of a moving steady-state on the convergence process. Specifically, both counterfactuals have .05 as inflow rates and .03 as outflow rates, so the two counterfactuals have the same steady-state—.625. The only difference between the two counterfactuals is the starting points. One starts at .675, while the other starts at .602. The initial difference between the two is 7.3 percentage points. After six months, the gap between the two is still 4.4 percentage points. After 24 months, the gap reduces to below 1 percentage point. After 72 months, both of the two rates get very close to the steady-state .625. This illustrates the rate convergence when the transition rates are low, isolating all other confounders.

from early summer 2020 could have followed different routes. One possibility is that non-participating individuals entered the labour force through unemployment before transitioning to employment. Another possibility is that they went directly to employment. This section explores these possibilities.²²

1.4.2.1 Method

More generally, assume again a discrete-time framework, in which an individual at time t is in a world of three states: employed (E), unemployed (U) or out of the labour force (O). Assume no population change over a short period. One can write the participation rate as a function of the initial stocks and six sequences of transition rates:

$$pr_t = g(pr_0, er_0, upr_0; \lambda^{EU}; \lambda^{EO}; \lambda^{UE}; \lambda^{UO}; \lambda^{OE}; \lambda^{OU}), \quad (1.6)$$

where $\lambda^{JK} = (\lambda_1^{JK}, \lambda_2^{JK}, \dots, \lambda_t^{JK})$ for $JK \in \{EU, EO, UE, UO, OE, OU\}$. It should be noted that not only do the flows between employment and non-participation, and between unemployment and non-participation, play a role in determining the participation rate, but also the flows within the labour force (i.e. flows between E and U states).

Using the participation rate, employment rate, and unemployment-to-population rate from February 2020 as initial points, I construct counterfactual participation rates to study the effect of different flows on the participation rate. As with the two-state case, I use transition rates adjusted for margin error. I also use 2019 transition rates to reflect pre-pandemic levels, as these six transition rates remain fairly stable in the years prior to the pandemic.

1.4.2.2 Results

The origin of outflows and destination of inflows to the labour force

Figure 1.5a illustrates the comparison between the actual and counterfactual participation rate holding flows from employment to non-participation and from non-participation to employment constant at pre-COVID levels, respectively. As illustrated by Figure 1.5a, many individuals who were employed in February flowed out of the labour force in March and April 2020. Initially, the inflows to employment did not explain the change in the participation rate but did explain the rebound of the participation rate in the later months. The figure suggests that most flows in and out of the labour force in 2020 were between employment and non-participation. By June 2020, 72.8% of the gap between the actual and counterfactual participation rates in the I to O counterfactual can be explained by the gap in the E to O counterfactual. Inversely, by November 2020, 65.1% of the gap in the O to I counterfactual can be explained by the gap in the O to E counterfactual.²³

²²However, during the pandemic, unemployment spells could be shorter than usual, and interviews representing the LFS reference week might miss some transitions in and out of unemployment. Thus, some employment-to-non-participation transitions through unemployment might not be captured here.

²³The employment-to-non-participation flows went back to the pre-pandemic level in August and September 2020. By November, the flows from non-participation to employment have reached a new stable level.

Since the flows between employment and non-participation explain most of the participation rate change, one can expect that the flows in and out of unemployment cannot explain much of the change. This is reflected in Figure 1.5b. The outflows from the labour force earlier in 2020 were not from unemployment, and the flows back later in 2020 were more to employment.

However, despite directly transitioning from employment to non-participation, older individuals may also gradually exit the labour market by moving from employment to unemployment and then from unemployment to non-participation. To determine which type of labour force separation drove the significant decline in the participation rate in March and April 2020, one can compare the two counterfactuals in Figure 1.6a. As illustrated, older individuals are more likely to transition directly from employment to non-participation, compared to taking the indirect route through unemployment. Nonetheless, a sizeable group of older individuals does move through unemployment. Additionally, the comparison of the two counterfactuals in Figure 1.6b suggests that when entering the labour force from non-participation, individuals are more likely to transition directly to employment rather than through unemployment.

The churning within the labour force: flows between employment and unemployment

Focusing on the three-state framework where dividing the labour force into employment and unemployment allows one to capture more variation within the labour market, i.e. flows between employment and unemployment. However, the flows within the labour force are not the main driver of the dive in the participation rate in early 2020 and the recovery later. Though the churning inside the labour force has increased unemployment, the leakage from unemployment to non-participation is not as significant as employment to non-participation. That is why the flows within the labour force have not significantly affected the size of the labour force.

1.4.3 Three Labour Market States: In Labour Force, Marginal Attachment and Non-Attachment

1.4.3.1 Method

Based on where the initial outflows from the labour force originated, this section focuses on the destination of those outflows of the labour market. Specifically, it examines whether people still want a job or have a job to start in the future after leaving the labour force, or if they are no longer attached to the labour force. This section provides insight into decomposing older individuals' labour participation by exploring more about non-participation, i.e., the marginally attached and non-attached groups.

As with the E-U-O three-state, one can write the participation rate of I-M-N three-state framework as

$$pr_t = h(pr_0, mr_0, nr_0; \lambda^{MI}; \lambda^{NI}; \lambda^{IM}; \lambda^{IN}; \lambda^{MN}; \lambda^{NM}), \quad (1.7)$$

where $\lambda^{JK} = (\lambda_1^{JK}, \lambda_2^{JK}, \dots, \lambda_t^{JK})$ for $JK \in \{IM, IN, MI, MN, NI, NM\}$. With the participation rate, marginal attachment rate and non-attachment rate of February 2020 as initial

points, I use the margin-error-adjusted transition rates of the same months of the year 2019 to reflect the pre-pandemic level of transitions and build counterfactuals.

1.4.3.2 Labour Attachment of Older Individuals during the Pandemic

As is shown in Figure 1.7a, both the flows from the labour force into marginal attachment and non-attachment significantly explain the slump in participation rate at the onset of COVID-19. In the three years before COVID-19, older individuals tended to arrive at non-attachment other than marginal attachment when they separated from the labour force.²⁴ However, during COVID-19, the flows to marginal attachment play an essential role. From March to April, the flow to marginal attachment is even more significant than to non-attachment.²⁵ With the previous findings that most of the outflows from the labour force are from employment, Figure 1.7a shows that many older individuals who separated from employment to non-participation during the early pandemic still wanted a job and did not seem to retire.

Although the percentage increase in transitions from the labour force to non-attachment during the pandemic was smaller than the corresponding percentage increase in flows to marginal attachment, the overall impact of non-attachment on the participation rate was larger. This is because the stock of non-attached individuals was much larger—nearly twenty times greater than that of the marginally attached group prior to COVID-19. Figure 1.7b shows that while marginal attachment experienced a greater proportional change, the increase in flows from non-attachment back into the labour force accounts for a larger share of the post-COVID recovery in the participation rate. In other words, a key driver of the participation rate rebound was individuals who had previously been out of the labour force and did not want a job returning to the labour force.

1.5 Robustness Checks

This section presents five robustness checks designed to address potential concerns about sample construction, data reliability, methodological assumptions, and sensitivity to parameter choices. Each test helps ensure that the findings regarding older individuals' labour force transitions during COVID-19 are not artifacts of arbitrary choices or data limitations.

First, I expand the sample to include individuals aged 55 to 69, rather than restricting it to those aged 55 to 64. While the full pension age in Canada is 65, many individuals aged 65 to 69 remain active in the labour force, and their participation has grown steadily over recent decades (Milligan and Schirle, 2018a). Including this older group helps assess whether the pandemic's impact on transitions and labour force dynamics differs at higher ages. The results show similar patterns: although participation levels are lower and labour market transitions are less frequent

²⁴The average transition rate from participation to non-attachment (λ^{IN}) from 2017 to 2019 is .021, while the average transition rate from participation to marginal attachment (λ^{IM}) of the same period is .008.

²⁵Researchers may argue that the classification of marginal attachment and unemployment might be obscure during the early pandemic. According to the decision rule of the LFS, besides temporal layoff and future start, an individual has to be both searching for a job and available to work to be classified as unemployed. Thus, re-labeling is not an issue at hand in the context.

for those aged 65 to 69, the qualitative findings remain the same. This suggests that the results are not sensitive to the age cutoff used in the baseline specification.

Second, I address concerns about the reliability of imputed labour market information in the Labour Force Survey (LFS). The LFS imputes current labour market status for respondents who previously participated but fail to respond in a given month. If these imputations are inaccurate, they may generate spurious transitions and bias the estimated transition rates. To address this concern, I exclude individuals with fully imputed labour market data and re-estimate the transition rates and counterfactuals. The findings remain consistent, indicating that the main results are not driven by imputed observations but reflect real labour market behaviour.

Third, I test whether the results are sensitive to the specific choice of pre-COVID reference year used to construct counterfactuals. The main analysis uses transition rates from 2019, the year immediately preceding the pandemic. To examine whether this choice unduly influences the results, I construct an alternative counterfactual using the average of 2017, 2018, and 2019 transition rates. The differences between the 2019 rates and the three-year averages are very small, typically in the third decimal place, and the counterfactual results are nearly identical. This demonstrates that the conclusions are robust to the choice of pre-COVID reference year.

Fourth, I conduct a placebo test to assess whether the iteration method used to build counterfactual participation rates introduces artificial divergence. Specifically, I use 2018 transition rates to construct counterfactual participation rates for 2019—a year with no significant labour market shock. If the counterfactual and actual participation rates for 2019 are closely aligned, this would suggest that the iteration method itself does not generate spurious discrepancies. The results confirm this: the counterfactual and actual participation rates for 2019 match very closely, validating the reliability of the method and confirming that the gaps observed in 2020 are not artifacts of the iteration process.

Lastly, I examine the robustness of results to the adjustment method used to correct for the margin error problem. In the main analysis, transition rates are adjusted using a minimum distance estimator that incorporates a weighting matrix based on the variance-covariance matrix of the raw transition rates. To verify that the findings are not sensitive to this method, I repeat the analysis using (i) a simplified identity matrix as the weighting matrix and (ii) unadjusted raw transition rates. Both robustness checks yield similar qualitative results. Although unadjusted rates can generate small cumulative discrepancies between iterated and actual participation rates, the overall trends and conclusions remain intact. This suggests that the key findings are not driven by the adjustment procedure.

1.6 Differences in Flows and Labour Participation between Older Men and Women

Having illustrated that the initial decline in labour participation rates among older individuals during the pandemic was primarily driven by substantial outflows, while the subsequent recovery has been attributed to increased inflows since May 2020, this section examines how the transitions among different labour states differ between older men and women, and how these differences affect their participation rates.

On average, the labour participation rate for older men was 11.6 percentage points higher than that for women in the three years preceding the pandemic. During the early phases of the pandemic, both older men and women experienced a significant drop in participation rates. However, men recovered more quickly than women. Figures 1.8 to 1.10 illustrate the percentage change in participation rates for both men and women since the onset of the pandemic, along with counterfactual scenarios assuming women had the same flow rates as men.

Figure 1.8 shows that older men's participation rate dropped by 4.4 percent in the first two months following the onset of the pandemic and returned to February 2020 levels five months later. In contrast, older women's participation rate dropped by 8.1 percent in the first two months and recovered more slowly, taking ten months to return to February 2020 levels. From February to October 2020, women's outflow rate from the labour force was higher than men's, while their inflow rate was lower during the year following the onset of the pandemic. This difference explains the gap in the recovery of participation rates between men and women. If women had experienced the same lower initial outflow rates as men, their percentage change in participation rate would have mirrored men's in the first four months after the onset of the pandemic. Furthermore, if women had consistently lower outflow rates like men, their participation rate would have been 3.9 percent higher than February 2020 levels twelve months after the onset of the pandemic. Conversely, if women had higher inflow rates like men, their participation rate would have been 12.4 percent higher than in February 2020, twelve months after the onset of the pandemic.

Figure 1.9a shows that, in terms of outflows from the labour force, flows from employment to non-participation explain more about the gender gap in percentage change in participation rate than flows from unemployment to non-participation. If the older women had experienced lower employment-to-non-participation rates similar to men, their percentage change in participation rate would have been very similar to men's in the first five months after the onset of the pandemic and would have been 3.9 percent higher than February 2020 levels twelve months later. In terms of inflows to the labour force, women have had both lower non-participation-to-employment and non-participation-to-unemployment rates than men since the onset of the pandemic. Figure 1.9b shows that both flows from non-participation to employment and unemployment explain the gender gap in percentage change in the participation rate.

The pandemic caused a greater number of older women to drop out of the labour force (especially from employment) compared to men. Additionally, more male non-participants entered

the labour force (both to employment and unemployment) as public health restrictions eased compared to women.

From the perspective of the destination of outflows from the labour force, both flows to marginal attachment and non-attachment played an important role in the initial drop in participation rates. However, older men and women do not differ in flows to marginal attachment but do differ in flows to non-attachment. The flow from the labour force to non-attachment is higher for women than for men. Compared to men, more women flowed out of the labour force and did not want a job during the pandemic. Figure 1.10a shows that outflows to non-attachment explain more about the gender difference in percentage change in participation rate than flows to marginal attachment. Figure 1.10b shows that if older women had higher inflows from non-attachment to the labour force so that they were similar to men, women's participation rates could have returned to February 2020 levels five months after the onset of the pandemic. If women had the same inflow rates from non-attachment to the labour force as men during the entire year, their participation rate would have been 7.5 percent higher than February 2020 levels by February 2021. However, flows from marginal attachment to the labour force do not seem to explain the gender gap in percentage change in participation rates. Although flows in and out of marginal attachment were significantly different during the pandemic compared to periods before, they do not explain the gender difference in participation rate changes.

1.7 Comparison of COVID-19 with the Great Recession

Similar to the COVID-19 recession, the Great Recession was also relatively short-lived, lasting from October 2008 to May 2009. However, the declines in GDP and employment during the Great Recession were not as severe and intense as those experienced during the COVID-19 pandemic. Additionally, the labour force participation of older individuals during the Great Recession differed from that observed during the COVID-19 recession. This section examines the changes in flows and their impact on the participation rate of older individuals during the Great Recession, and compares it to the case of the COVID-19 recession.

During the Great Recession, the participation rate of older individuals experienced only a slight drop of 0.7 percentage points in the first two months, followed by a rebound to higher levels. The outflow rate during the recession remained similar to levels seen three years prior, and even fell below pre-recession levels afterward. Figure 1.11 shows that outflows from the labour force did not significantly explain the variation in participation rates during the Great Recession. If the outflow rate had remained at pre-recession levels after the recession (and not have become lower), the participation rate could have been around 0.5 percentage points lower than it actually was between May 2009 and September 2009. The Great Recession did not appear to prompt older individuals to exit the labour force; on the contrary, older individuals had a higher probability to remain in the labour force during the recession.

In contrast, inflows into the labour force were higher during the Great Recession compared to three years before. Figure 1.11 indicates that these inflows played an important role in explaining the pattern of labour force participation of older individuals during and right after

the Great Recession. Had the inflow rate not exceeded pre-recession levels, the participation rate could have been 1.1 percentage points lower than the actual rate by May 2009.

Figure 1.12 illustrates the impact of different outflows from and inflows into the labour force. During the Great Recession, the flow from unemployment to non-participation decreased. There was a slight increase in flows from employment to non-participation in the first two months, but they declined to below pre-recession levels in the subsequent months. Figure 1.12a indicates that decreased unemployment-to-non-participation flows had a more substantial impact on older individuals' participation rates during the recession, while both decreased outflows from employment and unemployment contributed significantly to the increase in participation rates in the six months following the recession.

Regarding inflows to the labour force, flows from non-participation to unemployment increased by an average of 0.36 percentage points compared to pre-recession levels, while flows from non-participation to employment saw a slight increase during the recession and a decrease afterward. Figure 1.12b shows that in the first four months, increased inflows to employment accounted for much of the change in participation rates among older individuals, while later, the effect of increased inflows to unemployment became more significant. If transitions from non-participation to unemployment had remained at pre-recession levels, the participation rate could have been 1.29 percentage points lower by December 2009.

Examining the destinations of outflows from the labour force, outflows to marginal attachment increased, while outflows to non-attachment decreased starting in October 2008. Figure 1.13a demonstrates that reduced flows from the labour force to non-attachment played a more crucial role in explaining the changes in older individuals' participation rates. Conversely, when looking at the origins of inflows, flows from marginal attachment to the labour force did not vary significantly before and after the recession, while flows from non-attachment to the labour force increased during the recession. Figure 1.13b indicates that if transitions from non-attachment to the labour force had not exceeded pre-recession levels, the participation rate could have been 0.71 percentage points lower by May 2009.

Overall, Figures 1.12 and 1.13 show that during the Great Recession, unemployed older individuals did not appear discouraged and did not drop out of the labour force. When individuals did leave, they tended to remain marginally attached to the labour market rather than out of the labour force and did not want a job. This contrasts with the COVID-19 recession, where many individuals initially exited the labour force and did not want a job.

The Great Recession significantly differed from the COVID-19 recession in that older individuals generally remained in the labour force rather than leaving it. Additionally, during the Great Recession, increased inflows from non-participation to unemployment significantly contributed to the rise in participation rates post-recession, whereas during COVID-19, inflows from non-participation to employment were the primary drivers of recovery in participation rates.

1.8 Retirement Transitions during COVID-19

Though previous findings show that substantial flows from employment to non-participation explain the initial drop in the participation rate of older individuals during the pandemic, the transition to non-participation may not necessarily indicate an immediate or permanent decision to leave the labour force. This section investigates whether older individuals separated from their jobs to enter retirement since the onset of the pandemic by examining self-identified retirement transitions.²⁶

The decision of older individuals to retire depends on whether the value of retirement exceeds the value of continuing to work or actively seeking a job if they are getting laid off. At the onset of the pandemic, negative psychological shocks, such as difficulties adjusting to changes in the working environment and concerns about COVID-19 infection, may have pushed older individuals toward retirement. However, federal financial aid programs, such as the Canada Emergency Response Benefit (CERB)/Canada Recovery Benefit (CRB), provided them with an incentive to remain in the labour market. These benefits, which were being greater than typical pension income, required recipients to not voluntarily leave their jobs. As a result, some older individuals might have found it more profitable to stay attached to the labour market and claim CERB/CRB rather than retire. Moreover, as the negative psychological shock diminished over time, those who initially left their jobs might find the value of working outweighs the value of retirement again and return to the labour force.

I test the above conjecture using LFS data. Unlike the CPS in the US, the LFS does not include a direct question identifying current retirees.²⁷ However, the LFS includes a question for non-employed individuals who had a job in the previous twelve months, asking why they left or lost their jobs. One of the possible responses is retirement. This question allows for the identification of retirement transitions, i.e., transitions from employment to retirement.²⁸ Additionally, one can construct a six-month panel to follow the same individuals to determine when they transition into retirement, and whether they return to work later. It is helpful to explain the recovery of the labour participation rate in later 2020, i.e. whether the same individuals who dropped out of the labour force came back or were new older workers participating in the labour market.

The rate at which individuals move from being employed to retirement shows that transitions from employment to retirement accounted for only a small portion of the initial decrease in labour participation (with 14.4% of the flow from employment to non-participation in March

²⁶Retirement transitions in this study are identified based on respondents' reported reason for leaving their last job, with one of the options being "retired." This measure captures the individual's own characterization of their labour market exit, rather than formal retirement status tied to the receipt of public or private pension benefits. Some individuals may report retirement as their reason for leaving, even if they are not yet collecting pension income; while others may receive pensions but not report retirement as the reason for leaving.

²⁷A new variable about non-participants' main activities (NLFMACT) can identify if a respondent is currently retired, but it has been included in the survey only since March 2020.

²⁸The variable is of good quality, and Statistics Canada has used it to report retirement statistics. The series (Table: 14-10-0060-01) is widely cited (e.g., Gower (1997)). However, this variable cannot identify retirement transitions that do not follow the employment-to-retirement path. This is not a significant issue for this research, as during the pandemic, substantial outflows from the labour force were from employment.

and 4.3% in April). Figure 1.14 shows that the monthly retirement transition rate did not increase during the pandemic.

During the pandemic, the receipt of employment benefits is associated with a lower probability of transitioning to retirement.²⁹ Figure 1.15 illustrates that receiving at least one type of employment benefits reduced the probability of transitioning to retirement for both men and women aged 55 to 59.³⁰ However, receiving employment benefits did not significantly affect retirement transitions for men and women aged 60 to 64. For men aged 65 to 69, receiving employment benefits reduced the retirement transition probability, while for women in the same age group, it increased the probability.³¹

Retirement behaviour may differ substantially across industries due to variation in job characteristics such as physical demands, workplace safety, and access to benefits like employer-sponsored retirement plans. These differences became particularly relevant during the COVID-19 pandemic, when industry-specific risks and disruptions may have shaped older workers' retirement decisions. Figure 1.16 presents retirement transition probabilities by industry for men and women aged 55 to 64. The results suggest that the pandemic is associated with a reduced probability of retirement transitions for men in industries such as utilities, manufacturing, wholesale trade, transportation, finance and housing services, educational services, and hospitality. In contrast, for women in the same age group, the probability of retirement transitions increased in industries such as mining, financial services, educational services, and arts and entertainment. However, it is important to note that most of these industry-specific estimates are not statistically significant and should be interpreted with caution.

Using a six-month panel to track individuals' labour market trajectories, Table 1.2 compares the probability of transitioning back to employment between 2019 and 2020 for those who lost employment and dropped out of the labour force (or became non-attached to the labour force) in March or April of both years. Column (1) shows that the unconditional probability of returning to employment in May for those who were out of the labour force in March or April 2020 is 36.9%, compared to 27.3% in 2019. By September, the probability of returning to employment for those who were out of the labour force in March or April 2020 is 69.8%, which is twice as high as during the same period in 2019. Unlike in 2019, when individuals who became non-attached to the labour market in March or April had a low probability of returning to employment, column (2) shows that by September 2020, 58.1% of individuals who became non-attached in March or April returned to employment. A significant portion of older individuals who dropped out of the labour force and did not want a job in the early pandemic returned later.

²⁹These patterns should be interpreted cautiously, as they may reflect underlying differences in labour market attachment or job characteristics. For instance, individuals who receive employment benefits may be more attached to their jobs, which could in turn lower their likelihood of retirement.

³⁰The controls of the regressions are demographic characteristics, working industry, and income levels.

³¹Among the various employment benefits, the reception of a workplace pension plan increased the probability of a retirement transition for individuals aged 55 to 59, 60 to 64, and 65 to 69. Other benefits, such as paid vacation and sick leave, reduced the probability of transitioning to retirement.

1.9 Conclusion

In this paper, I find that during the COVID-19 pandemic, there was a significant shift in labour flow dynamics for older individuals in Canada. Compared to pre-pandemic patterns, there were substantial outflows from the labour force to the state of non-participation. Many older individuals transitioned directly to non-participation without gradual detachment from the labour force. However, since May 2020, there were significant and persistent flows back into the labour force, with many older individuals entering employment from both the state of marginal attachment and non-attachment.

Moreover, I find that older women experienced a slower recovery in labour force participation compared to older men. From February to October 2020, more women exited the labour force, and fewer women entered the labour force during the subsequent year. The differences in outflows from employment and inflows to both employment and unemployment explain the gender gap in participation rate changes. Additionally, the pandemic affected retirement transition probabilities differently for older men and women across various industries. Employment benefits impacted retirement transitions similarly for men and women aged 55 to 64, but differently for those aged 65 to 69.

During the Great Recession, older individuals did not tend to leave the labour force, and many non-participants began to actively look for jobs. The loss in total wealth prompted older individuals to stay in the labour force and postpone retirement. In contrast, COVID-19 led to a substantial number of older individuals exiting the labour force, particularly during the first lockdown, due to a combination of factors, including health concerns and changes in work environments. These outflows persisted even after the economy reopened, but significant inflows back into the labour force were also observed. Consequently, labour participation among older individuals recovered eight months after the pandemic began.

Similar to the findings by [Quinby, Rutledge and Wettstein \(2021\)](#) in the US, retirement transitions explained only a small portion of the drop in participation rates in Canada. Tracking the same individuals over six months revealed instances of reverse retirement, where individuals who transitioned from employment to retirement or to non-attachment returned to work later. The pandemic did not entirely disrupt work incentives for older individuals. Many individuals became marginally attached to the labour force and eventually resumed employment. Employment benefits and federal emergency financial aid helped near-retirement-age individuals remain in the labour force.

However, this study is subject to several limitations. Firstly, there is no direct identification in the LFS of whether flows into employment are new entrants or re-entrants.³² As a result, though knowing that half of the older workers who separated from their jobs earlier returned to employment in the later months of 2020, it is still difficult to determine if the participation rate recovery includes new labour force entrants. Secondly, the variable in the LFS that identifies retirement transition is so specific that few transitions can be captured. If one were to follow the

³²The variable FLOWUNEM can identify different types of flows into unemployment.

labour trajectory of the same individual for six months, only a handful of retirement transitions and transition back could be tracked.

1.10 Figures and Tables




FIGURE 1.1: Monthly labour force participation rates of individuals aged 55 to 64 of US and Canada, 2017m1-2021m12. Data source: the LFS public-use files and CPS IPUMS files. The solid line shows the participation rates of the age group 55-64 for Canada, and the dashed line shows the rates of the same age group for the US.




FIGURE 1.2: Annual labour force participation rates of age groups 15-54 and 55-64, 1976-2021. Data source: the LFS public-use files. The solid line shows the participation rates for the age group 55-64, and the dashed line shows the rates for the age group 15-54. The vertical red dashed lines indicate the starting and end dates of the recessions.




FIGURE 1.3: Monthly labour force participation rates of individuals 55-64 during the Great Recession and COVID-19, seasonally adjusted. Data source: Statistics Canada, Table 14-10-0287-01. The horizontal axis is the number of months before/after the start of each recession. The vertical red dashed line indicates the start of the two recessions.






FIGURE 1.4: Actual vs counterfactual participation rates in In-Out two-state. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. The dashed line is the counterfactual participation rate that holds λ^{IO} constant at 2019 level. The dash-dot line is the counterfactual participation rate that holds λ^{OI} constant at 2019 level. 90% bootstrapping confidence interval is shown in the shaded area.




(A) flows between Employment and Non-Participation




(B) flows between Unemployment and Non-Participation

FIGURE 1.5: Actual vs counterfactual participation rates in Employment-Unemployment-Non-Participation three-state. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. In panel (a), the dashed and dash-dot lines are the counterfactual participation rates that hold λ^{EO} and λ^{OE} constant at 2019 level, respectively. In panel (b), the dashed and dash-dot lines are the counterfactual participation rates that hold λ^{UO} and λ^{OU} constant at 2019 level, respectively. 90% bootstrapping confidence interval is shown in the shaded area.




(A) direct and indirect transiting out of the labour force




(B) direct and indirect transiting into the labour force

FIGURE 1.6: Actual vs counterfactual participation rates in Employment-Unemployment-Non-Participation three-state. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. In panel (a), the dashed line is the counterfactual participation rate that simultaneously holds λ^{EU} and λ^{UO} constant at 2019 level, and the dash-dot line is the counterfactual participation rate that holds λ^{EO} constant at 2019 level. In panel (b), the dashed line is the counterfactual participation rate that simultaneously holds λ^{OU} and λ^{UE} constant at 2019 level, and the dash-dot line is the counterfactual participation rate that holds λ^{OE} constant at 2019 level. 90% bootstrapping confidence interval is shown in the shaded area.



(A) outflows from the labour force



(B) infows into the labour force

FIGURE 1.7: Actual vs counterfactual participation rates in Participation-Marginal-Non-Attachment three-state. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. In panel (a), the dashed and dash-dot lines are the counterfactual participation rates that hold λ^{IN} and λ^{IM} constant at 2019 level, respectively. In panel (b), the dashed and dash-dot lines are the counterfactual participation rates that hold λ^{NI} and λ^{MI} constant at 2019 level, respectively. 90% bootstrapping confidence interval is shown in the shaded area.

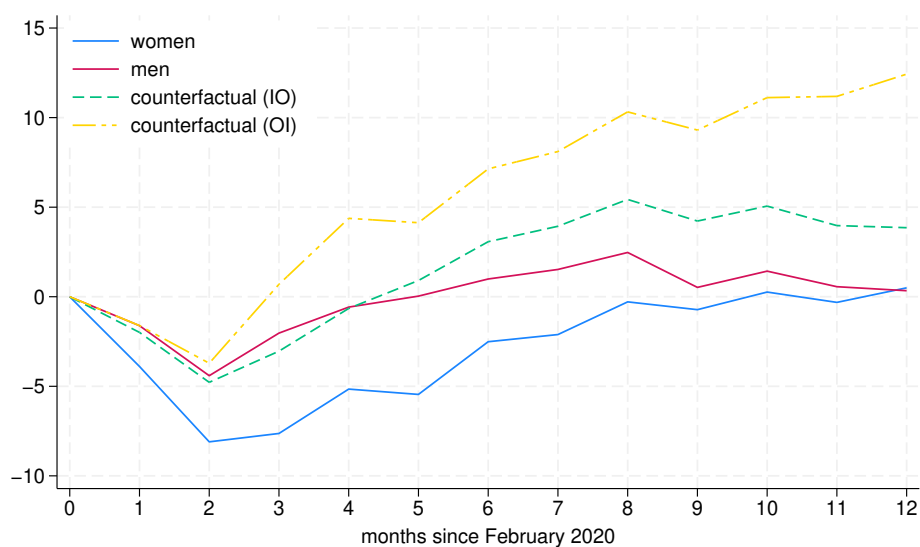




FIGURE 1.8: Percentage change in participation rate since February 2020, men and women. Data source: the LFS master files, individuals aged 55 to 64. The red solid line is the percentage change in participation rate of men, and the blue solid line is the percentage change in participation rate of women. The dashed line represents the counterfactual percentage change in women's participation rate if they had the same λ^{IO} as men. The dash-dot line represents the counterfactual percentage change in women's participation rate if they had the same λ^{OI} as men.

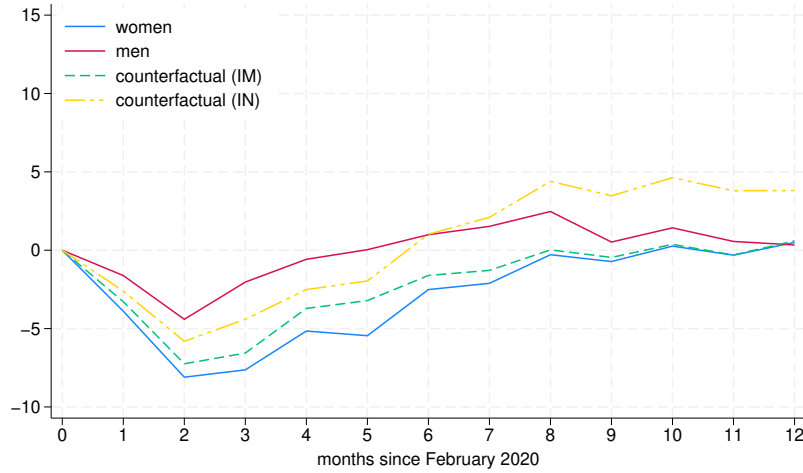


(A) outflows from the labour force

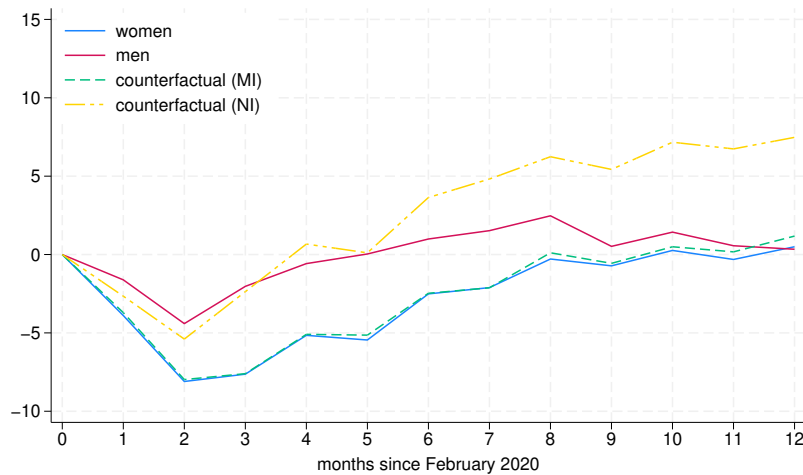


(B) inflows into the labour force

FIGURE 1.9: Percentage change in participation rate since February 2020, men and women. Data source: the LFS master files, individuals aged 55 to 64. The red solid line is the percentage change in participation rate of men, and the blue solid line is the percentage change in participation rate of women. In panel (a), the dashed line represents the counterfactual percentage change in women's participation rate if they had the same λ^{EO} as men. The dash-dot line represents the counterfactual percentage change in women's participation rate if they had the same λ^{UO} as men. In panel (b), the dashed line represents the counterfactual percentage change in women's participation rate if they had the same λ^{OE} as men. The dash-dot line represents the counterfactual percentage change in women's participation rate if they had the same λ^{OU} as men.



(A) outflows from the labour force



(B) inflows into the labour force

FIGURE 1.10: Percentage change in participation rate since February 2020, men and women. Data source: the LFS master files, individuals aged 55 to 64. The red solid line is the percentage change in participation rate of men, and the blue solid line is the percentage change in participation rate of women. In panel (a), the dashed line represents the counterfactual percentage change in women's participation rate if they had the same λ^{IM} as men. The dash-dot line represents the counterfactual percentage change in women's participation rate if they had the same λ^{IN} as men. In panel (b), the dashed line represents the counterfactual percentage change in women's participation rate if they had the same λ^{MI} as men. The dash-dot line represents the counterfactual percentage change in women's participation rate if they had the same λ^{NI} as men.


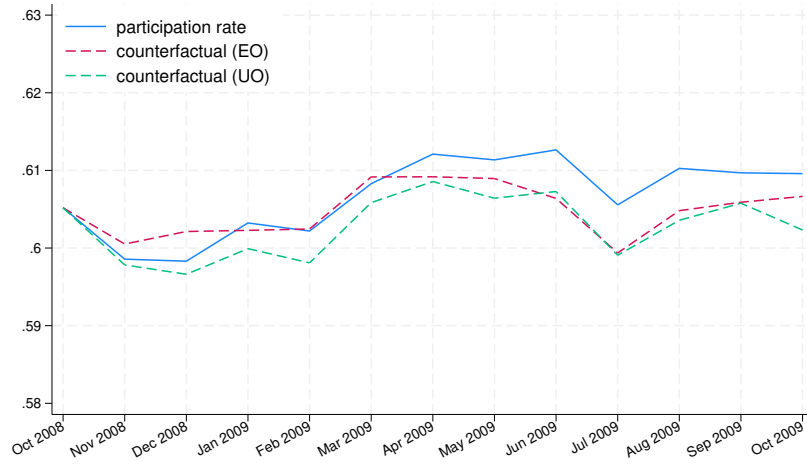



FIGURE 1.11: Actual vs counterfactual participation rates in In-Out two-state, Oct 2008 - Oct 2009. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. The red dashed line is the counterfactual participation rate that holds λ^{IO} constant at the pre-recession level. The green dashed line is the counterfactual participation rate that holds λ^{OI} constant at the pre-recession level. The pre-recession level of transitions is calculated with the average of each month of the three years before the Great Recession.

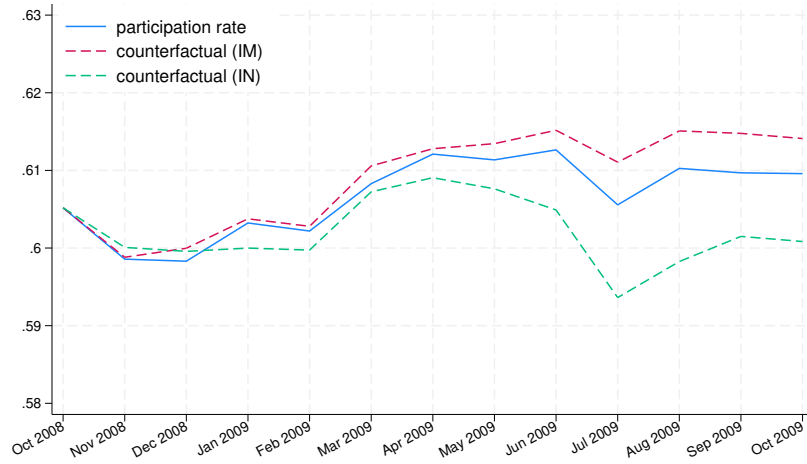


(A) outflows from the labour force




(B) inflows into the labour force

FIGURE 1.12: Actual vs counterfactual participation rates in Employment-Unemployment-Non-Participation three-state, Oct 2008 - Oct 2009. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. In panel (a), the red dashed line is the counterfactual participation rate that holds λ^{EO} constant at the pre-recession level. The green dashed line is the counterfactual participation rate that holds λ^{UO} constant at the pre-recession level. In panel (b), the red dashed line is the counterfactual participation rate that holds λ^{OE} constant at the pre-recession level. The green dashed line is the counterfactual participation rate that holds λ^{OU} constant at the pre-recession level. The pre-recession level of transitions is calculated with the average of each month of the three years before the Great Recession.



(A) outflows from the labour force



(B) inflows into the labour force

FIGURE 1.13: Actual vs counterfactual participation rates in Participation-Marginal-Non-Attachment three-state, Oct 2008 - Oct 2009. Data source: the LFS master files, individuals aged 55 to 64. The solid line is the actual participation rate. In panel (a), the red dashed line is the counterfactual participation rate that holds λ^{IM} constant at the pre-recession level. The green dashed line is the counterfactual participation rate that holds λ^{IN} constant at the pre-recession level. In panel (b), the red dashed line is the counterfactual participation rate that holds λ^{MI} constant at the pre-recession level. The green dashed line is the counterfactual participation rate that holds λ^{NI} constant at the pre-recession level. The pre-recession level of transitions is calculated with the average of each month of the three years before the Great Recession.




FIGURE 1.14: Monthly Employment-to-Non-Attachment and Employment-to-Retirement Rates, 2017m1-2022m6. Data source: the LFS master files, individuals aged 55 to 64. The solid line denotes the transition rate from employment to non-attachment, while the dashed line denotes the transition rate from employment to retirement.




FIGURE 1.15: The Effect of Receiving Employment Benefit on Retirement Transition Probability. Sample period: 2020m3-2022m7. Data source: the LFS master files, individuals aged 55 to 69. The horizontal axis denotes the effect of receiving at least one type of employment benefit on retirement transition probability in percentage points. The regressions control for education, marital status, earning level, industry, province and month fixed effects, and individuals' tenure before the retirement transition.

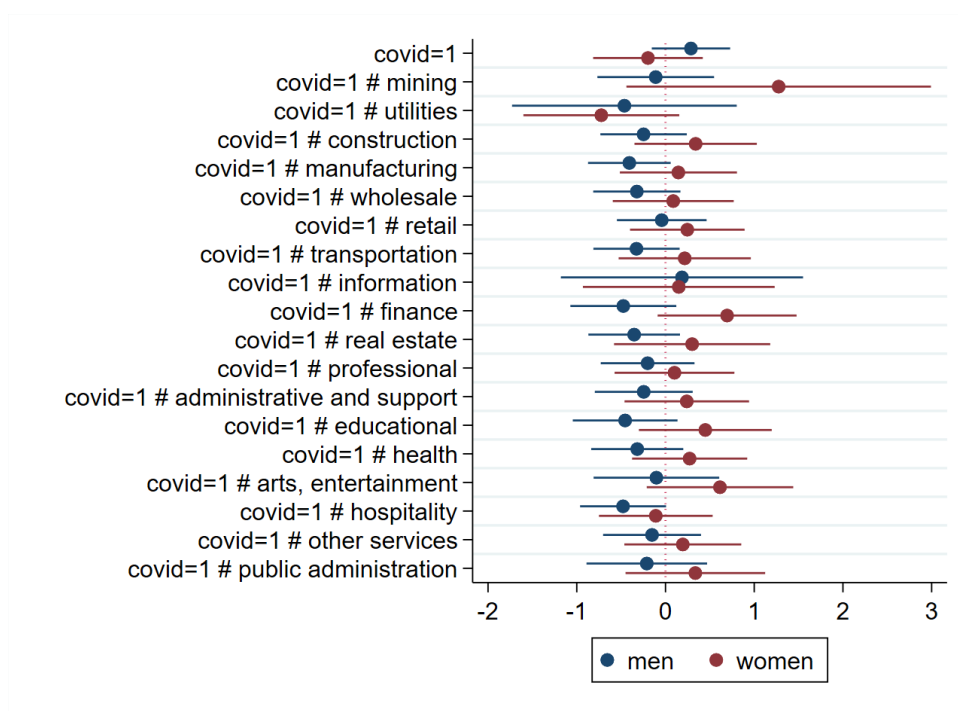


FIGURE 1.16: The Effect of COVID-19 on Retirement Transition Probability by Industry, Men and Women. Sample period: 2017m1-2022m7. Data source: the LFS master files, individuals aged 55 to 64. The horizontal axis denotes the effect of COVID-19 on retirement transition probability in percentage points, while the vertical axis represents the interaction of industries with the COVID dummy variable. The regression controls for education, marital status, earning level, province and month fixed effects, and individuals' tenure before the retirement transition.

TABLE 1.1: Transition Rates before and during COVID-19

	Time Period			F Tests	
	2017m2 - 2020m1 (1)	2020m2 - 2021m1 (2)	2021m2 - 2022m1 (3)	2020m2 - 2021m1 (4)	2021m2 - 2022m1 (5)
<i>Panel A: Employment, Unemployment and Non-participation</i>					
From employment					
to unemployment(λ^{EU})	.010 (.0002)	.024 (.0007)	.012 (.0005)	360.822 (.0000)	14.279 (.0002)
to OLF(λ^{EO})	.021 (.0003)	.031 (.0008)	.025 (.0007)	145.192 (.0000)	20.756 (.0000)
From unemployment					
to employment(λ^{UE})	.195 (.0037)	.247 (.0067)	.198 (.0065)	45.229 (.0000)	.112 (.7379)
to OLF(λ^{UO})	.155 (.0035)	.176 (.0060)	.163 (.0061)	8.775 (.0031)	1.099 (.2945)
From non-participation					
to employment(λ^{OE})	.030 (.0005)	.043 (.0013)	.038 (.0012)	98.839 (.0000)	34.767 (.0000)
to unemployment(λ^{OU})	.022 (.0005)	.032 (.0011)	.029 (.0011)	72.386 (.0000)	40.924 (.0000)
<i>Panel B: Labour Force Attachment</i>					
From labour force					
to marginal attachment(λ^{IM})	.008 (.0002)	.018 (.0006)	.012 (.0005)	276.679 (.0000)	56.980 (.0000)
to non-attachment(λ^{IN})	.021 (.0003)	.025 (.0007)	.023 (.0007)	30.789 (.0000)	10.971 (.0009)
From marginal attachment					
to INLF(λ^{MI})	.312 (.0065)	.384 (.0103)	.343 (.0113)	35.190 (.0000)	5.929 (.0149)
to non-attachment(λ^{MN})	.395 (.0069)	.338 (.0096)	.391 (.0113)	22.698 (.0000)	.072 (.7879)
From non-attachment					
to INLF(λ^{NI})	.039 (.0006)	.046 (.0013)	.046 (.0014)	27.866 (.0000)	21.674 (.0000)
to marginal-attachment(λ^{NM})	.018 (.0004)	.029 (.0011)	.026 (.0010)	95.739 (.0000)	51.693 (.0000)

Notes. Transition rates for individuals aged 55 to 64, 2017m2-2022m2. Standard errors are in parentheses. Column (4) shows the F statistics that transition rates 2020m2-2021m1 are statistically significantly different from the ones 2017m2-2020m1. Column (5) shows the F statistics that transition rates 2021m2-2022m1 are statistically significantly different from the ones 2017m2-2020m1. P-values for the F tests are shown in parentheses. Panel A shows the transitions between Employment (E), Unemployment (U) and Non-participation (O). Panel B shows the transitions between In labour force (I), Marginal attachment (M), and Non-attachment (N). λ^{JK} denotes the transition rate between labour states J and K from time $t - 1$ to t . Data source: LFS master files.

TABLE 1.2: Probability of Returning to Employment, April-September Panel

	From non-participation (1)		From non-attachment (2)	
	2020	2019	2020	2019
Back to employment in				
May	.369	.273	.414	.238
June	.515	.483	.575	.414
July	.600	.301	.534	.197
August	.582	.337	.539	.291
September	.698	.347	.581	.292
Weighted sample size	29,230	12,691	12,297	10,008

Notes. The proportion of individuals aged 50 to 64 transiting from non-participation and non-attachment back to employment, respectively. The sample of column (1) consists of individuals who are in non-participation and have been non-employed for no greater than two months as of April, and the sample of column (2) consists of individuals who are in non-attachment and have been non-employed for no greater than two months as of April. The proportions are calculated using weights.

Chapter 2

The Cyclicalities of Hiring and Separation by Experience and Tenure

2.1 Introduction

Labour market experience and job tenure affect workers' vulnerability to economic downturns in different ways. This paper shows that experience matters most for job mobility—particularly job-to-job transitions—while job tenure is more important in determining job security. During recessions, younger workers with short tenure face steeper declines in hiring and higher separation rates. In contrast, experienced or long-tenured workers are more insulated from labour market disruptions.¹

Existing research on the cyclical patterns of labour market flows has documented that younger workers are more sensitive to economic fluctuations than their older counterparts, particularly in terms of hiring and separation rates (e.g., [Xu and Couch, 2017](#); [Jaimovich and Siu, 2009](#); [Clark and Summers, 1981](#)). However, some of this observed age effect may stem from the fact that younger individuals are more likely to hold short-tenure jobs. As highlighted by [Jung and Kuhn \(2019\)](#), tenure is strongly correlated with age and plays an important role in shaping labour market transitions.

There is a substantial body of research indicating the importance of tenure in explaining labour market flows and cyclicalities in wages. For instance, [Topel and Ward \(1992\)](#) finds that young workers' job mobility depends on their tenure, and young workers tend to have short-tenured jobs in the first ten years of their careers. [Krolikowski \(2017\)](#) notes that the job-to-job transition probability decreases with increased tenure, with this trend becoming flatter after two years of job tenure. Further, [Albagli, Contreras, Tapia and Wlasiuk \(2022\)](#) observes that earnings are less cyclical for workers with longer tenure. Since job mobility of young workers is linked with wages ([Light and McGarry, 1998](#); [Topel and Ward, 1992](#)), the cyclicalities in wages could, in turn, influence the cyclicalities in their labour market transitions.

¹The precise definitions of experience and tenure used in this chapter are provided in the following paragraphs. Briefly, experience refers to potential labour market experience (based on one's age and years of schooling), while tenure measures the length of time with the current employer.

These findings suggest that some of the age or experience effects identified in earlier studies may have actually been driven by job tenure. The overlap between age and tenure makes it challenging to disentangle their respective contributions to the cyclicity of job mobility. Many prior studies rely on the U.S. Current Population Survey (CPS), which does not contain monthly job tenure information. This data limitation makes it difficult to separately identify the roles of experience and tenure. In contrast, the Canadian Labour Force Survey (LFS) provides monthly tenure data, enabling a comprehensive analysis of job transitions across both dimensions.

This paper uses Canadian LFS microdata to study the cyclicity of labour market transitions across experience and tenure groups. Experience is defined as general labour market experience, calculated as age minus years of schooling and the standard school entry age of six. Tenure is defined as the length of a worker's employment with the same employer. Building on Forsythe (2022), who document that hiring cyclicity in the U.S. is largely driven by differences in potential experience, this study asks whether similar patterns hold in Canada, and whether job tenure also plays an independent role in shaping cyclicity. I estimate individual-level monthly transition probabilities—such as job-to-job moves, separations to unemployment or non-participation, and entries from non-employment—using linear probability models with region, time, and demographic controls.

The findings show that hiring probabilities are procyclical, particularly for younger workers (those with less than 10 years of potential experience). A one percentage point (p.p.) increase in the regional unemployment rate is associated with a 0.13 p.p. decline in their probability of being newly hired, compared to only 0.02 p.p. for more experienced workers.² This difference holds across different types of hires—job-to-job, unemployment-to-employment, and non-participation-to-employment transitions. Moreover, when unemployment rises, young workers are more likely to experience job separations to both unemployment and non-participation, and these separations are more often involuntary.

Going beyond Forsythe (2022), this study compares the relative importance of potential experience and job tenure in driving these cyclical patterns. The results show that the dominant factor varies by transition type. For job-to-job transitions, experience accounts for 99% of the total cyclical variation explained by experience and tenure combined. For separations to unemployment and to non-participation, tenure accounts for 77% and 87%, respectively. In other words, experience matters most for hiring and upward job mobility, while tenure better predicts job retention and stability.

These findings are consistent with institutional layoff practices such as last-in, first-out (LIFO), where workers with shorter tenure—especially those still in probationary periods—are typically the first to be dismissed during downturns. Specifically, among short-tenured workers (less than 3 years), those with less experience face a much sharper decline in hiring probability—0.085 p.p.

²New hires are defined as individuals who transition into employment in a given month after not being employed or being employed in a different job in the previous month.

versus 0.026 p.p. —in response to a one p.p. rise in the unemployment rate.³ However, among less experienced workers, differences in hiring cyclicalities across tenure groups are negligible. Conversely, for separations, workers with long tenure are significantly less likely to separate when the unemployment rate increases, but this protective effect does not vary much by experience within the short-tenure group.⁴

This paper contributes to the literature on labour market flows and business cycles. While much of the existing work has examined the aggregate cyclicalities of hiring, separations, and job-to-job transitions (e.g., [Shimer, 2005](#); [Fujita and Ramey, 2006, 2009](#); [Nakamura, Nakamura, Phong and Steinsson, 2019](#); [Hairault, Le Barbanchon and Sopraseuth, 2015](#)), fewer studies have looked at how these patterns vary by experience and tenure jointly. This paper fills that gap using monthly microdata from Canada’s LFS and highlights the importance of considering both general and match-specific human capital in understanding labour market dynamics.

In addition, this work complements recent efforts to estimate the effects of experience and tenure on wages or hiring probabilities (e.g., [Schmieder, von Wachter and Heining, 2023](#); [Williams, 2009](#)), by focusing specifically on the cyclicalities of these outcomes. The results emphasize that experience and tenure are not interchangeable: they shape different dimensions of labour market vulnerability and resilience across the business cycle.

The remainder of the paper is structured as follows: Section 2.2 presents the data and Section 2.3 analyzes the cyclicalities of new hires and separations by experience. Section 2.4 breaks down the cyclicalities of hires and separations into components attributable to experience and tenure. Section 2.5 presents robustness checks, and Section 2.6 offers concluding remarks.

2.2 Data

This study relies on the master files of the Canadian Labour Force Survey (LFS). The LFS is a large household survey conducted by Statistics Canada that contains rich labour market and demographic information for Canadian residents. On a monthly basis, the dataset encompasses approximately 100,000 individuals.

I take advantage of the rotating panel design of the LFS to construct two-month mini-panels.⁵ To do so, I adopt the method provided by Statistics Canada for matching individuals across time, relying on a unique combination of household IDs and individual-specific IDs within those households. This information is only made available in the master files.

³I chose the 3-year cutoff based on the empirical distribution of transition rates by tenure. This cutoff also aligns with prior literature that observes non-linearities in mobility and wage growth within the first few years of tenure — for example, [Topel and Ward \(1992\)](#) and [Krolikowski \(2017\)](#).

⁴Splitting the sample by tenure introduces stronger selection concerns than using experience. Tenure reflects not only time with the employer but also underlying job match quality, firm stability, and possibly unobserved worker characteristics. So I don’t interpret the tenure coefficient as a causal effect of tenure per se. Rather, I treat tenure as a proxy for a bundle of factors associated with job stability — and my results should be understood as descriptive patterns in how workers with different observed job attachment respond to economic fluctuations.

⁵Households are followed for six consecutive months, with one-sixth of the sample being replaced every month. See [Brochu \(2021\)](#) for more details on the structure of the LFS.

For the part of the analysis that explores the role of employer-employee match duration, I take advantage of the rich tenure data available in the LFS. As part of the regular monthly questionnaire, the LFS asks respondents when they started working for the present employer, and the in-progress tenure duration is recorded in months.

For my main analysis, I try as much as possible to impose restrictions similar to those in Forsythe (2022). This includes restricting my attention to two-month mini-panels that span two consecutive months and focusing on the 1995 to 2019 period. Similar to Forsythe (2022), I impose the same lower age restriction of 16. I must, however, impose an upper age restriction of 69 years of age, which she does not. This is done because the LFS, in order to ease the interviewee's burden, only asks for labour market information of those 70 years and up in the initial interview, and the information is carried forward in subsequent months, as long as the households remain in the survey. As such, one cannot capture labour market transitions for this older age group. This is not a concern since most Canadians retire well before the age of 70.⁶ Additional restrictions that are typically linked to the creation of mini-panels are discussed in Appendix A.1.

2.3 Replication of Forsythe (2022)

2.3.1 Methodology

Following Forsythe (2022), I estimate how the hiring probability varies with business cycles for experienced and inexperienced young individuals by the following specification:

$$D_{ikst}^{\text{hired}} = \alpha_s + \delta_t + Z'_{ikst}\theta + \sum_{k=1}^K (\beta_k D_k^{\text{PE}} + \gamma_k D_k^{\text{PE}} \times \text{Unemp. Rate}_{st}) + \epsilon_{ikst}, \quad (2.1)$$

where D_{ikst}^{hired} is an indicator for being hired (rescaled to 100) for individual i in potential experience group k , residing in Employment Insurance Economic Region (EIER) s at time t . I define new employment matches—referred to as new hires—as transitions either from non-employment to employment or from one job to another.⁷

D_k^{PE} is an indicator for potential experience group k . I calculate potential labour market experience as: potential experience = age – years of education – 6. When this formula yields a negative value, I impute it as zero—this occurs in about 1% of the sample. Based on this measure, individuals are classified into two groups: those with fewer than ten years of potential experience are categorized as ‘young’, and those with more than ten years are considered ‘experienced’.

Unemp. Rate_{st} denotes the unemployment rate in EIER s at time t . EIERS are geographic units defined by the Government of Canada to reflect relatively homogeneous local labour markets. There are 55 EIERS in Canada. While this number is similar to the number of U.S. states,

⁶Individuals aged above 70 years old represent only a small portion of the workforce, i.e. approximately 1% over my sample period.

⁷In accordance with Forsythe (2022), I do not include transitions into self-employment as new job matches.

each EIER represents a smaller population and geographic area due to Canada's lower overall population.⁸

Additional controls include regional fixed effects (α_s), month-year fixed effects (δ_t), and a vector of demographic controls (Z_{ikst}), which are interacted with a corresponding coefficient vector θ . The demographic controls consist of dummy variables for a respondent's gender and highest level of educational attainment.⁹

In the given econometric model, the parameter γ_k holds particular significance as it captures the interaction effect between the potential experience group indicator, D_k^{PE} , and the unemployment rate, Unemp. Rate_{st} , on the likelihood of an individual being hired. Interpretation of γ_k centers on understanding how the probability of being hired for individuals within experience group k is influenced by changes in the regional unemployment rate. A positive value of γ_k suggests that as the unemployment rate increases, the probability of individuals with potential experience level k being hired is higher, relative to the base probability captured by β_k . Conversely, a negative value of γ_k would indicate that a higher unemployment rate reduces the hiring probability for individuals in that experience group. Essentially, γ_k reflects the sensitivity of hiring prospects for different experience levels to the economic conditions of a region, as described by the unemployment rate.

I will also be re-estimating equation (2.1) for other transitions, such as job-to-job transitions, to better understand the hiring effect.¹⁰ I also estimate equation (2.1) for job separations and flows between unemployment and non-participation. I break down job separations into separations to unemployment and separations to non-participation. Moreover, I distinguish job separations as involuntary separations and voluntary separations.

2.3.2 Results

In this section, I present the findings on the cyclicity of the new hires, emphasizing differences between young and experienced workers.¹¹ Additionally, I delve into the cyclicity of job separations and other labour flows and their variations between the different experience groups.

2.3.2.1 New Hires

Figure 2.1 illustrates that individuals with fewer than ten years of potential work experience are generally more likely to be hired compared to those with over ten years of experience. Notably,

⁸One may note that certain labour market policies, such as minimum wage legislation, are set at the provincial level and could influence hiring probabilities. As a robustness check, I replicate the main analysis at the provincial level, and the results are reported in Section 2.5.

⁹There are 14 demographic groups in my sample, defined by the intersection of gender and education level: male/female and one of seven education categories (no high school diploma, high school diploma, some post-secondary, trades certificate or diploma, college/CEGEP/University certificate or diploma below Bachelor's, Bachelor's degree, and university credential above the Bachelor's level).

¹⁰I considered a job-to-job transition valid when an individual was employed for two consecutive months but reported a job tenure of only one or two months for the second month. I also do a robustness check on a stricter definition of job-to-job transition as part of the new hires and present the results in section 2.5.

¹¹Actually, all hires are new hires. Here I follow the definition used in Forsythe (2022) and use the word 'new hires' to emphasize that hires are very recent, i.e. happened in one month.

this dynamic shifts in the context of elevated unemployment rates, wherein the likelihood of being hired for individuals with 10 to 45 years of experience experiences a notable increase (around 0.5 p.p.). It suggests that firms are more likely to hire experienced workers when the economic condition is bad.

Recessions tend to have a greater effect on young workers compared to experienced workers, as depicted in Figure 2.2, which is in line with the findings of Forsythe (2022) in terms of the U.S. labour market. During the 2001 downturn, hiring rates decreased for both groups.¹² However, during the financial crisis, young workers faced a decrease in hiring rates—approximately 0.5 p.p.—while experienced workers’ hiring rates remained relatively stable during the same period. Panel (b) of Figure 2.2 applies the Hodrick-Prescott (HP) filter to separate the cyclical components of these rates.¹³ The resulting de-trended plot of hiring rates for young versus experienced workers shows a clear pattern: young workers’ hiring rates are more cyclical and react more noticeably to economic changes, whereas experienced workers’ hiring rates exhibit a degree of stability to such fluctuations.

Panel A of Table 2.1 demonstrates how the hiring probability fluctuates with changes in the unemployment rate within EIERs for the overall working-age population.¹⁴ All the estimates in this panel and subsequent tables have standard errors clustered at the EIER level. Across columns (1) to (4), the analysis progressively incorporates controls for EIER, demographic characteristics, and month-year fixed effects. The estimates turn negative and become statistically significant as soon as one controls for region fixed effects (column (2)), and remain this way when demographic controls and month-year fixed effects are also included.

In panel B of Table 2.1, the focus shifts to the cyclical effects on hiring probabilities for young versus experienced workers. More precisely, it shows the estimates of equation (2.1) with sequential addition of controls. Initially, column (1) reveals that a one p.p. increase in the unemployment rate corresponds to a 0.055 p.p. reduction in the likelihood of hiring young workers, while it slightly raises the probability by 0.06 p.p. for experienced workers. Subsequent insights from columns (2) and (3), which introduce controls for EIER and demographic characteristics, demonstrate that both groups are adversely affected by increasing unemployment rates in EIERs. However, this impact is more pronounced for young workers. These trends remain consistent even after considering month-year fixed effects. Specifically, column (4) shows that a one p.p. rise in unemployment leads to a 0.13 p.p. decline in hiring chances for young workers and a 0.02 p.p. decrease for experienced workers. Notably, the latter effect is statistically significant only at the 10% level. This highlights the disproportionate impact of rising unemployment

¹²According to the NBER Business Cycle Dating Committee, there was a recession from March 2001 to November 2001 in the U.S. However, in Canada, Cross and Bergevin (2012) refers to 2001 as a “period of considerable weakness in economic activity”.

¹³Researchers (e.g., Hamilton (2018)) may have concern that HP filter cannot capture the true cyclical behaviour of a time series, but instead capture a mechanical behaviour by its construction. Alternatively, I use the Hamilton filter to separate the long-term trend from the variation at business cycle frequencies of hiring rates. The result is shown in Figure B.1, which also suggests that hiring rates are more cyclical for young workers than experienced workers.

¹⁴The labour income dynamics for the workers who are close to retirement ages are very different from those of the prime working-age workers. I also did the same analysis on individuals aged 16 to 49 and got very similar results.

rates on young workers, compared to experienced workers with comparable demographics within specific EIERs and time frames.

In comparing my findings with those of Forsythe (2022) regarding the U.S. labour market, one observes significant similarities with the Canadian context. Specifically, the effect of business cycles is really felt in the hiring probabilities of young workers, which tend to move in the same direction as business cycle fluctuations. In both the U.S. and Canadian scenarios, the impact of rising unemployment rates on the hiring probabilities of experienced workers is considerably smaller in absolute terms and less statistically significant compared to that of younger workers.

Table 2.2 delves into the nuances of hiring-probability cyclicity between young and experienced workers in relation to changes in the EIER unemployment rate. This table dissects the origins of these hires, categorizing them as either emerging from employment, unemployment, or non-participation. In all cases, the full set of controls is included.

For ease of presentation, column (1) of Table 2.2 replicates the patterns found in column (4) of Table 2.1. Thereafter, columns (2) through (4) shed light on outcomes for hires sourced from employment, unemployment, and non-participation, respectively. Across these columns, a consistent narrative unfolds: young workers are particularly vulnerable during economic downturns compared to their more experienced counterparts. In detail, column (2) indicates that the adverse influence of the EIER unemployment rate on the likelihood of young workers moving from one job to another is fourfold higher compared to experienced workers. Column (3) highlights that the detrimental impact of the EIER unemployment rate on the transition from unemployment to employment is 0.68 p.p. more pronounced for young workers. Similarly, column (4) reveals that young workers experience a tenfold greater negative effect from unemployment when it comes to hiring from non-participation. The aggregated effects of a one p.p. increase in the EIER unemployment rate on young and experienced workers are summarized in column (1). Utilizing the Wald test, I confirm that the impact of the EIER unemployment rate on young and experienced workers differs in a statistically significant manner across all columns presented.

The results presented in Table 2.2 are in line with the observations made by Forsythe (2022) regarding the U.S. labour market. The findings of the two countries illustrate a common pattern: young workers suffer more during economic contractions, irrespective of their prior employment status, whether employed, unemployed or not participating in the labour force. While Forsythe (2022) reports positive coefficients for experienced workers in columns (1), (2), and (4), these figures are not only small in magnitude but also exhibit low statistical significance. This suggests that experienced workers do not face the same level of vulnerability as younger workers.

To provide a more nuanced and detailed understanding, Figure 2.3 shows how unemployment rates affect hiring probabilities at different years of potential experience, segmented into one-year increments.¹⁵ Panel (a) of Figure 2.3 offers a comprehensive view of all hires, showing a

¹⁵Figure 2.3 provides a detailed representation of estimates of equation (2.1). Unlike the broader categorization into young and experienced workers, this figure refines the analysis by estimating the impact of unemployment rates on hiring probabilities at more granular, one-year intervals.

prominent hump-shaped pattern. Notably, an uptick in unemployment rates seems particularly harmful for those with under 15 years or over 40 years of experience. This observation deviates from the U.S. findings documented in Forsythe (2022). The lower hiring probability among older workers may be linked to their higher reservation wages, posing a hurdle to employment during economic slumps. Additionally, the over-40-year segment might be over-represented by less educated individuals if those with higher education have opted for retirement, thus amplifying the observed decline in hiring probabilities facing rising unemployment rates.

Delving into job-to-job transitions for each experience year, panel (b) presents a concave trend, with most coefficients falling below zero. This pattern suggests that, during economic hardships, job-to-job transitions become less frequent for both young and more experienced workers. However, within a given EIER and month-year, there seems to be a greater willingness to hire individuals with over 15 years of experience from other jobs.

The narrative in panel (c) shows a universally negative impact of higher unemployment on the recruitment of individuals from unemployment, regardless of experience level. While this effect diminishes for those with more years of potential experience, it remains in line with Forsythe (2022)'s U.S. findings.

Contrasting with U.S. patterns, panel (d) of Figure 2.3, which examines hires from non-participation in relation to potential experience, reveals a distinct trend in the Canadian context. Here, the adverse impact of unemployment rates on the hiring likelihood decreases as potential experience increases, but it never shifts to a positive effect. This negative influence persists even beyond 15 years of potential experience. In the United States, in contrast, the impact of unemployment rates on the likelihood of hiring from non-participation turns positive after 30 years of experience. This difference underscores a unique characteristic of the Canadian labour market: a consistent negative effect of rising EIER-level unemployment rates on hiring opportunities across all experience levels, unlike the U.S., where extensive experience can eventually mitigate this negative influence.

2.3.2.2 Job Separations

Table 2.3 shifts the focus from the cyclicity of new hirings discussed in Tables 2.1 and 2.2 to the cyclicity in various forms of job separations and transitions between unemployment and non-participation. Taking a broad perspective, column (1) reveals that an uptick in the EIER unemployment rate by one p.p. makes it more common for young workers to experience separations from their jobs, compared to experienced workers.¹⁶ Moving to column (2), there is a clear indication that, for both young and experienced workers, a one p.p. rise in the EIER unemployment rate correlates with a higher probability of separating from their jobs to enter unemployment. Although the Wald test confirms a statistically significant difference in the effects on the two potential experience groups, the actual magnitude of these effects is similar.

¹⁶The difference is not statistically significant, though.

In contrast, column (3) illustrates that in a contracting economy, young workers exhibit a decreased tendency to switch to another job.¹⁷ Statistically, this adverse impact is significantly more pronounced for young workers compared to experienced workers. Column (4) reveals that when the EIER unemployment rate is one p.p. higher, both young and experienced workers are more likely to separate from a job to non-participation. The effect is stronger for young workers than for experienced workers.¹⁸

In an interesting turn, column (5) demonstrates that economic downturns see both young and experienced workers moving from non-participation to active job-seeking. This movement suggests a heightened sense of competition or slack in the labour market, prompting individuals to re-enter the job hunt. Conversely, column (6) indicates a reluctance to transition from unemployment to non-participation during these times. Notably, the Wald tests show that these cyclical flows between unemployment and non-participation do not significantly vary between the young and the experienced.

To understand the role of workers' willingness to quit a job in explaining the cyclicalities of job separations, I disaggregate separations from employment into voluntary and involuntary separations in Table 2.4. Column (1) highlights that an increase of one p.p. in the EIER unemployment rate correlates with a rise in the probability of involuntary separations for both younger and experienced workers. Notably, this effect is significantly more pronounced for the former, indicating that economic downturns disproportionately impact younger individuals in terms of job security. Interestingly, the results do not suggest a statistically significant disparity in the cyclicalities of voluntary separations to unemployment between the younger and experienced workers when faced with adverse economic conditions.

Synthesizing the insights from Tables 2.3 and 2.4, it appears that economic downturns—marked by rising unemployment rates—primarily lead to increased separations of younger individuals from employment. Notably, under adverse economic conditions, younger workers do not show a greater likelihood of transitioning out of employment via voluntary separations compared to their more experienced counterparts. This suggests that the elevated job separation rates among young workers during economic downturns are more likely driven by involuntary separations rather than voluntary quits.

2.3.2.3 Human Capital and Experience

In addition to labour market experience, human capital accumulation may also affect the probability of a worker finding a new job. Are workers with a college degree more prone to be hired by prospective employers? This section examines the role of university education in modulating the cyclical nature of hiring probabilities.

¹⁷Note that the coefficients in column (3) of Table 2.3 are different from those in column (2) of Table 2.2, though they are both for job-to-job transition. The reason is that, when defining hires, I exclude those who transit to self-employed; when defining separations, I didn't exclude this group of people.

¹⁸Nevertheless, when the analysis is replicated at the provincial level (detailed results are in section 2.5), the rise in provincial unemployment rates does not distinctly affect the separation probabilities for young versus experienced workers, with the singular exception being the likelihood of transitioning to a new job.

In Figure 2.4, the impact of one p.p. escalation in the EIER-level unemployment rate on the employment prospects across various one-year potential experience bins, segregated by educational attainment (college degree holders versus non-holders), is graphically represented. The illustration reveals a pronounced disparity: during adverse economic conditions, as evidenced by a heightened unemployment rate, individuals lacking a college degree within their initial fifteen years in the labour market are disproportionately affected. In contrast, beyond fifteen years of labour market experience, the possession of a college degree does not confer enhanced resilience during economic downturns. In fact, post-fifteen years of experience, those who do not hold a college degree exhibit a heightened probability of employment amidst rising unemployment rates.¹⁹

Table 2.5 shows the aggregate impact of unemployment rate increases on the hiring prospects of both educational cohorts. The analysis incorporates sequentially controlling for the EIER, demographic, and month-year fixed effects, which progressively attenuate the coefficient magnitudes. Column (4) of Table 2.5 indicates that an uptick in the unemployment rate adversely affects both educational groups, with a more pronounced impact on college degree holders. Nevertheless, the Wald test ascertains no statistically significant difference in the effects experienced by these groups. Summarizing these insights with the findings from Figure 2.4 suggests that while a college degree initially shields workers from the adverse effects of economic downturns, its protective advantage diminishes after accruing fifteen years of experience.

2.4 Tenure

2.4.1 Potential Experience and Job Tenure on Job-to-Job Transitions

In this part of the study, I take advantage of the tenure information available in the LFS to explore whether the varying impact of EIER unemployment rates on the likelihood of job-to-job transitions for young and experienced workers is actually influenced by the length of job tenure. Given that young workers typically have shorter tenures, it is conceivable that job tenure could be the underlying variable responsible for the observed differences in how unemployment rates affect young versus experienced workers. This section aims to determine if, when holding tenure length constant, young and experienced workers are affected differently by EIER unemployment rates.

¹⁹Because potential experience is calculated using age and years of schooling, comparisons across education groups at fixed experience levels may confound age effects. For example, a college graduate with 5 years of potential experience is typically 4–5 years older than a non-college counterpart. Consequently, observed differences between education groups could partially reflect age rather than education alone.

To do so, I further augment equation (2.1) to account for the cyclicity of job-to-job transition probability by different job tenure groups:²⁰

$$D_{ijkst}^{\text{hired}} = \alpha_s + \delta_t + Z'_{ijkst}\theta + \sum_{k=1}^K (\beta_k D_k^{\text{PE}} + \gamma_k \times D_k^{\text{PE}} \times \text{Unemp. Rate}_{st}) + \sum_{j=1}^{J-1} (\eta_j D_j^{\text{TEN}} + \rho_j \times D_j^{\text{TEN}} \times \text{Unemp. Rate}_{st}) + \epsilon_{ijkst}, \quad (2.2)$$

where D_{ikst}^{hired} is an indicator for the event of a job-to-job transition (rescaled to 100) for individual i in tenure group j and experience group k from EIER s at time t . Equation (2.2) is equation (2.1) augmented by the tenure level effect and the heterogeneous effect of the unemployment rate by different tenure groups.²¹ I denote workers with fewer than three years of job tenure as short-tenured workers, and workers with more than three years of tenure as long-tenured workers.²²

The parameters γ and ρ in the equation are critical for understanding how the heterogeneous effect of the unemployment rate on the hiring probability is influenced by both varying levels of tenure and potential experience. The coefficient γ represents how the unemployment rate affects those with different levels of potential experience when holding tenure length constant. ρ captures the influence of unemployment on individuals with varying tenure lengths when holding potential experience constant. In the regression, the omitted category is long-tenured workers.²³

Table 2.6 assesses the relative influence of potential experience versus job tenure on job-to-job transitions. Column (1) of Table 2.6 duplicates the results from column (2) in Table 2.2. As previously discussed, when the EIER-level unemployment rate goes up, young workers are hit harder; they are 0.089 p.p. less likely to jump from one job to another. For their more experienced counterparts, this change means a decrease in transition likelihood by 0.024 p.p.. Instead of controlling only for potential experience, column (2) shows how tenure itself influences job switches. It reveals that workers with under three years of tenure are more likely to make

²⁰The current specification does not appear to be symmetric between experience and tenure. Alternatively, I can include Unemp. Rate_{st} as one regressor and then have interactions between Unemp. Rate_{st} and one category of tenure and one category of experience. In that case, the coefficients before the interaction terms would show the marginal effects of changing a potential experience group or a tenure group. However, for ease of comparing the results with the last section, I use the specification currently presented in the paper.

²¹The current specification separately interacts tenure and experience bins with the unemployment rate, assuming additive effects. A fully saturated specification—interacting experience bins, tenure bins, and the unemployment rate—would allow for more flexible estimation of heterogeneous cyclical sensitivity. For example, it could reveal whether job tenure plays a more important role for workers with less potential experience. While this extension would be valuable, I did not implement it due to data access limitations. I plan to explore this richer interaction structure in future work.

²²For a summary of the joint distribution of experience and tenure groups, please refer to Appendix B.1 Tables B.1 and B.2.

²³For example, denoting γ_1 the effect of the unemployment rate on the dependent variable for young workers and γ_2 the effect for experienced workers, the difference between γ_1 and γ_2 reflect that, conditional on the same tenure length, how the EIER unemployment rate affects young and experienced workers differently. Denoting ρ_1 the effect of the unemployment rate on the dependent variable for short-tenured workers, it reflects that, conditional on the same level of the potential experience, how the EIER unemployment rate affects short-tenured workers differently from the long-tenured workers.

job-to-job transitions than those who are settled longer, and when the EIER unemployment rate increases, those short-tenure workers (under three years) face a steeper decline (0.057 p.p. decrease) in job-to-job transitions compared to their counterparts (0.023 p.p. decrease) who have held onto their jobs for over three years.

Column (2) might lead one to believe that tenure is the driving factor behind the negative impact of unemployment on job-to-job transitions. Yet, the findings in column (3) challenge this interpretation, allowing both potential experience and tenure effects to be present. Holding tenure level constant (e.g., tenure fewer than three years), one sees that when the EIER unemployment rate rises by one p.p., workers with more than ten years of potential experience face a smaller reduction of 0.026 p.p. (-0.018 p.p. - 0.008 p.p.) in transition probability, compared to their less experienced counterparts, who see a 0.085 p.p. (-0.077 p.p. - 0.008 p.p.) reduction. However, holding potential experience constant (e.g., potential experience of fewer than ten years), when the EIER unemployment rate increases by one p.p., one does not see a big difference in the decrease in job-to-job transition probability between short-tenured (0.085 p.p.) and long-tenured (0.077 p.p.) workers. This evidence points to potential experience playing a more pivotal role in explaining job-to-job transitions than job tenure. It could suggest that general skills are more critical than employer-specific skills in navigating job changes, or it may reflect suboptimal job matching for those with shorter tenure.

2.4.2 Potential Experience and Job Tenure on Job Separations

Table 2.7 extends the analysis from Table 2.3, exploring the relationship between job stability and tenure. The findings, spanning columns (1) to (4), reveal a disparity in how workers of different tenures respond to economic shifts in terms of job separations. Those with fewer than three years of job tenure exhibit a heightened sensitivity to economic fluctuations. Similar to the observations in Table 2.3, a rise in the EIER unemployment rate suggests an increased likelihood of these individuals exiting their jobs to join unemployment for both tenure groups. However, the increase in the EIER's unemployment rate decreases the likelihood of moving to another job.²⁴ In terms of shifting from employment to non-participation, a one p.p. increase in the unemployment rate tends to push short-tenure workers towards non-participation, whereas long-tenure workers show a reduced tendency to move into non-participation under similar conditions.

Columns (5) and (6) of Table 2.7 reveal that long-tenured workers have a higher propensity to transition from non-participation to unemployment with an increase in the EIER unemployment rate, compared to their short-tenured counterparts. However, while this difference is statistically significant, its economic significance is less pronounced. Conversely, under conditions of higher unemployment rates, short-tenured workers show a decreased likelihood of moving from unemployment to non-participation. In contrast, long-tenured workers are 0.02 p.p. more likely to make this transition. These findings indicate that workers, irrespective of their tenure, tend

²⁴Note that the coefficients of column (3), Table 2.7 are different from those in column (2), Table 2.6. The reason is that I excluded transitions to self-employment from new hires, while I didn't exclude those transitions from job separations.

to respond to labour market slackness during economic downturns, often re-entering the labour market from non-participation. Notably, short-tenured unemployed individuals are less prone to leave the labour market due to discouragement in such periods. It is important to note that for these analyses (columns (5) and (6)), the sample is composed of individuals who are non-employed but have worked in the past twelve months.

Table 2.8 delves into whether potential experience or job tenure is more predictive of job separations during economic downturns. In all columns of Table 2.8, I include both interaction terms of unemployment rate with potential experience and with job tenure. Column (1) shows that, within the low experience group (potential experience fewer than ten years), a worker with long tenure is 0.156 p.p. less likely to leave their job, given a one p.p. increase in the EIER unemployment rate. Conversely, a worker with short tenure is 0.322 p.p. (-0.156 p.p. + 0.478 p.p.) more likely to experience a job separation under the same conditions. On the other hand, within the short-tenured workers, the less experienced are 0.322 p.p. more likely to separate from their jobs when the EIER unemployment rate is one p.p. higher, while the more experienced ones are 0.489 p.p. more likely to experience the job separation. Conditional on the same tenure level, one does not see a large difference across experience groups, but conditional on the same experience group, one sees a big variation across tenure groups. This suggests that job tenure is a more significant factor in determining job separations during tough economic times than potential experience. Columns (2) and (4) support this, showing that separations to both unemployment and non-participation are more strongly influenced by job tenure. The exception is transitions to another job, where potential experience plays a more significant role, as seen in column (3) of Table 2.6 (job-to-job transition probability).

Columns (5) and (6) of Table 2.8 further indicate that potential experience more significantly impacts transitions from non-participation to unemployment, whereas job tenure has a greater influence on transitions from unemployment to non-participation than experience.

Summarizing the information from Tables 2.3, 2.7, and 2.8, it becomes clear that job tenure is a more critical factor in job separations during economic downturns than potential labour market experience.

2.5 Robustness Checks

This section presents a set of robustness checks designed to validate the core findings of the chapter. Each check addresses a specific potential concern—such as sample period sensitivity, aggregation level choices, measurement definitions, and data quality—to rule out alternative explanations and confirm that the findings are not dependent on specific sample selection, variable definitions or data imputations.

The first robustness check addresses the concern that the study's main conclusions may be specific to the post-2000 sample period. To test this, I extend the analysis back to 1990, thereby including the early 1990s recession and broadening the range of macroeconomic conditions under

study.²⁵ This expanded window allows for a more comprehensive test of the cyclicalities of hiring and separation probabilities across experience groups. The results remain consistent with the main analysis: workers with fewer than ten years of potential labour market experience exhibit more cyclical hiring probabilities and are more vulnerable to job loss during downturns. This consistency over a longer historical period suggests that the original findings are not merely an artifact of the selected sample window.

The second robustness check considers whether the findings depend on the geographic level at which unemployment rates are measured. While the main analysis uses the EIER unemployment rate, this check instead uses unemployment rates at the provincial level, which aligns with policy jurisdictions in Canada and mirrors the state-level analysis in the U.S. by Forsythe (2022). This test evaluates whether the main findings persist when labour market variation is aggregated differently. The results confirm that the cyclicalities in hiring probabilities among less experienced workers is still present and, in some cases, slightly more pronounced. Nevertheless, despite these variations in magnitude, the overarching conclusion remains unchanged, affirming the consistency of the study's primary findings across different levels of regional analysis.²⁶

The third check explores the possibility that the definition of a new hire—particularly job-to-job transitions—may influence the estimated cyclicalities of hiring. To address this concern, I adopt a stricter definition of job-to-job transitions in this robustness check. While the main analysis classifies individuals with one or two months of tenure in their second month of continuous employment as new hires from another job, the robustness specification restricts this to only those with exactly one month of tenure. This narrower definition helps exclude potentially ambiguous or less clearly defined transitions. Although this stricter definition reduces the measured magnitude of cyclicalities—coefficient sizes are about half of those in the baseline—the key qualitative result remains: less experienced workers are more likely to experience declines in hiring during economic downturns. This suggests that while the magnitude of the estimated effects may depend on how new hires are defined, the direction and significance of the result are robust.

The fourth robustness check assesses whether the results are driven by fully imputed labour force information. Since the LFS uses Whole Record Imputation (WRI) for individuals with missing labour market information, there is a risk that transitions constructed by imputed data may not reflect actual labour dynamics. To test this, I re-estimate all specifications after excluding individuals flagged as WRI. The results remain highly consistent with the baseline analysis, indicating that the use of imputed records does not materially affect the estimated relationship between experience and cyclical labour market outcomes. This reinforces the reliability of the findings and confirms that they are not driven by potential imputation biases in the data.

²⁵A notable change in the educational attainment variable in the LFS occurred in 1990. To maintain a consistent measure of potential labour market experience, the analysis does not extend to pre-1990 periods.

²⁶In terms of job separations, the heterogeneous effect of the unemployment rate on job separation is not statistically significant for separations to unemployment and non-participation at the provincial level. The reason might be that the provincial unemployment rate cannot encompass the same level of variation as the EIER unemployment rate, thus resulting in a smaller gap in RSS between the restricted model and unrestricted model, and thus smaller F statistics.

2.6 Conclusion

This study examines how economic downturns affect the probabilities of being hired and experiencing job separations across different worker groups, with a particular focus on differences by potential labour market experience and job tenure. It addresses the question: to what extent do economic downturns affect young and experienced workers differently, and are these effects primarily driven by short- or long-tenure employment relationships? To answer this, I estimate a linear probability model using data from the master files of the Canadian Labour Force Survey, incorporating a rich set of fixed effects and tracking individual-level transitions across employment states over the business cycle.

The findings show that the hiring probability is more cyclical for young workers. Moreover, the cyclicalities of the job-to-job transition probability is primarily influenced by potential labour market experience rather than the length of job tenure. This suggests that, during economic downturns, employers may prefer hiring individuals with broader labour market experience, who likely possess more general human capital, over those who have simply stayed longer with a previous employer. On the separation side, job tenure emerges as a stronger predictor of job stability during downturns—consistent with the idea that workers with longer tenure have accumulated firm-specific human capital, or are protected by implicit contracts, probation rules, or last-in-first-out layoff practices. In other words, experience drives upward mobility, while tenure buffers against downward transitions.

I summarize these findings in Table 2.9, which compares the relative contributions of potential experience and job tenure to the cyclicalities of various labour market flows. The results show that the cyclicalities of job-to-job transitions is overwhelmingly driven by experience-related effects, whereas the cyclicalities of separations to unemployment or non-participation is more strongly associated with tenure. For instance, nearly 99% of the variation in the cyclicalities of job-to-job flows is attributed to potential experience, while 77% and 87% of the variation in EU and EN transitions, respectively, are driven by tenure. This reinforces the interpretation that general human capital plays a larger role in helping workers find new jobs in downturns, while match-specific human capital plays a larger role in protecting jobs.

From a policy perspective, these findings highlight the importance of labour market initiatives that aim to bolster employment stability and security for younger workers, particularly during economic downturns. Potential measures could include the expansion of vocational training and apprenticeship programs, as well as the introduction of subsidies for early career development or wage subsidy schemes designed to encourage the hiring of less experienced workers during downturns. Additionally, enhancing severance packages for employees with fewer than three years of tenure could provide crucial financial support during involuntary downward transitions.

That said, several limitations merit discussion. First, while the empirical strategy relies on plausibly exogenous local unemployment rate shocks, the analysis remains descriptive rather than fully causal. I estimate OLS regressions with region, time, and demographic fixed effects, but unobserved heterogeneity in job characteristics, match quality, or firm behaviour

could still confound the observed relationships. Second, although the LFS is a high-quality, nationally representative survey, it primarily captures labour supply behaviour and does not include information from employers. This limits the ability to observe firm-side decisions, such as changes in vacancy posting, hiring standards, or wage offers, and constrains the analysis of demand-side responses. Additionally, the key variables—such as tenure and experience—are either self-reported or imputed based on age and education, which may introduce modest measurement error. While occasional misreporting is possible (e.g., rounding tenure to the nearest year), these errors are unlikely to systematically bias the main results. Still, they may affect the precision or magnitude of some estimated relationships.

Future research can build on this work by using linked employer-employee data to study how firm characteristics interact with worker experience and tenure in shaping cyclical labour market outcomes. In particular, incorporating information on job ladders, promotions, wage offers, and probationary policies could provide deeper insight into how general and match-specific capital influence employment dynamics. Furthermore, theoretical models of cyclical hiring behaviour—such as those based on search frictions, wage posting, or learning about match quality—could help clarify the mechanisms behind employers' decisions during downturns.

In summary, this study sheds light on the vulnerability of younger and less tenured workers during recessions, and it also highlights the complexity of labour market adjustments to business cycles, as well as the need for multifaceted policy responses. Understanding how experience and tenure shape workers' mobility and job stability over the business cycle remains crucial for designing effective labour market policies.

2.7 Figures and Tables

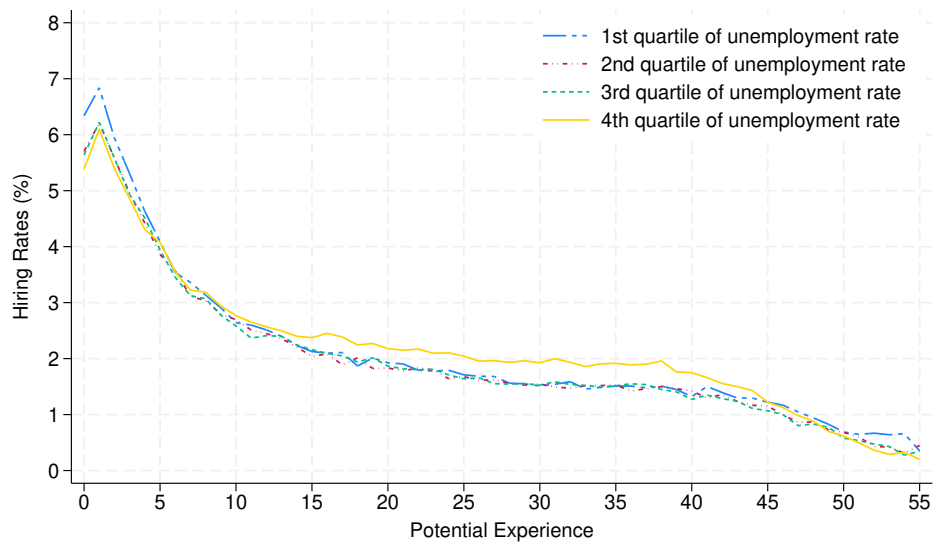


FIGURE 2.1: Potential Experience Profile of Hiring Rates, by EIER unemployment rate quartiles.

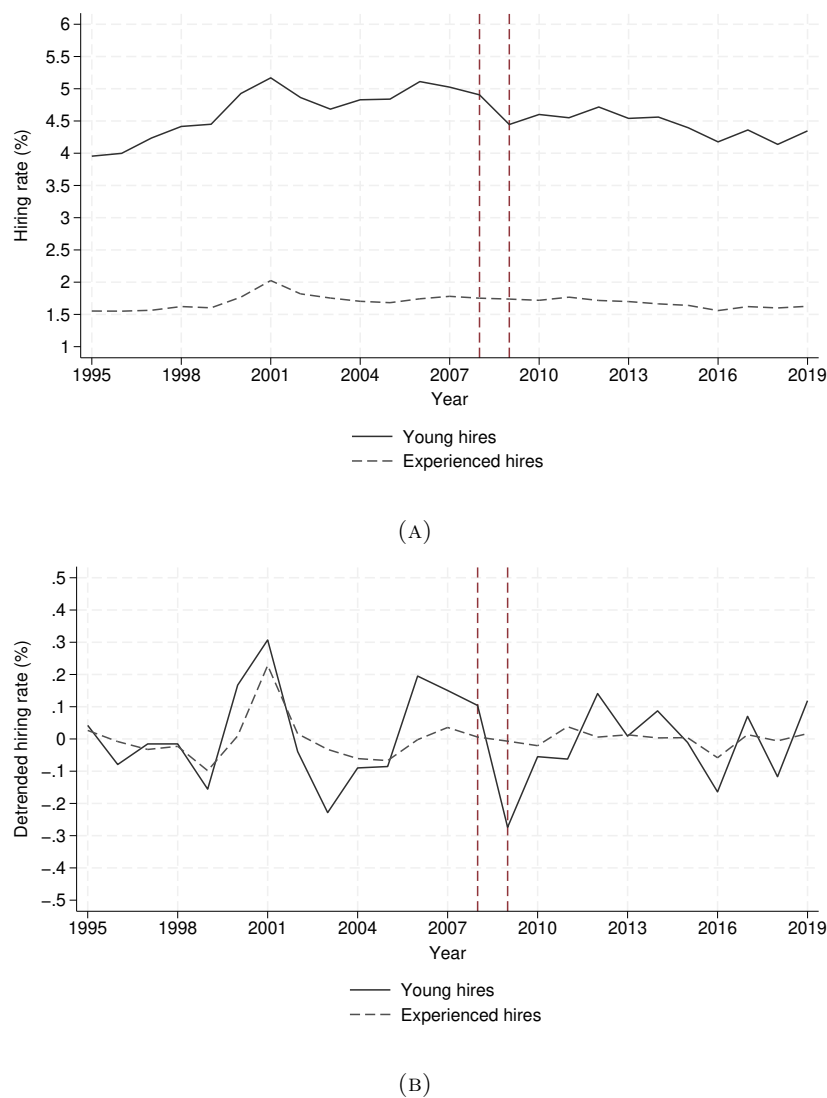


FIGURE 2.2: Hiring Over the Business Cycle. In panel (a), each line reflects the raw share of workers in the potential experience category hired each month in the LFS, weighted using LFS sampling weights and smoothed by averaging across the year. In panel (b), each series is seasonally adjusted by removing month-fixed effects, detrended by HP filter with a monthly smoothing factor (14,400), and then averaged across the year. Red dashed vertical lines indicate recent recession dates in Canada. Cross & Bergevin (2012) classify the date of the recent Canadian recessions as January 1980-June 1980, June 1981-October 1982, March 1990-April 1992, and October 2008-May 2009.




FIGURE 2.3: Coefficients of Unemployment Rate on Hiring Probability for One-Year Potential Experience Bins, Partialling out Main Effects and Employment Insurance Economic Region (EIER), Demographic, and Month-Year Fixed Effects and Weighted using LFS Sampling Weights. Areas between two dashed lines represent 95% confidence intervals, based on s.e. clustered at the EIER level.




FIGURE 2.4: Coefficients of Unemployment Rate on Hiring Probability for One-Year Potential Experience Bins, Partialling out Main Effects and Employment Insurance Economic Region (EIER), Demographic, and Month-Year Fixed Effects and Weighted using LFS Sampling Weights. The black line represents individuals without a college degree, and the grey line represents individuals with a college degree. Shaded areas represent 95% confidence intervals, based on s.e. clustered at the EIER level.

TABLE 2.1: Hiring over the Business Cycle: With and Without Controls

Outcome: Pr(Hired) \times 100	(1)	(2)	(3)	(4)
<i>Panel A: Aggregate effect</i>				
U. rate	0.0225 (0.0221)	-0.0832*** (0.0141)	-0.0977*** (0.0147)	-0.0507*** (0.0110)
R^2	0.000	0.000	0.002	0.004
<i>Panel B: Disaggregate by potential experience</i>				
PE \leq 10	4.661*** (0.173)	4.597*** (0.179)	4.435*** (0.180)	4.455*** (0.177)
PE \leq 10 \times U. rate	-0.0551* (0.0269)	-0.167*** (0.0231)	-0.173*** (0.0234)	-0.130*** (0.0163)
PE $>$ 10 \times U. rate	0.0660** (0.0218)	-0.0587*** (0.0130)	-0.0709*** (0.0139)	-0.0248* (0.0111)
R^2	0.009	0.010	0.011	0.012
EIER fixed effect	No	Yes	Yes	Yes
Demographic fixed effect	No	No	Yes	Yes
Month-year fixed effect	No	No	No	Yes
N	21,253,704	21,253,704	21,253,704	21,253,704

Notes: The dependent variable is a binary variable which refers to a worker beginning a job at a new firm. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE 2.2: Hiring over the Business Cycle: Young and Experienced

	(1)	(2)	(3)	(4)
Outcome: Pr(Hired) \times 100	All	EE	UE	NE
PE \leq 10	4.455*** (0.177)	2.264*** (0.129)	11.82*** (0.589)	9.220*** (0.599)
PE \leq 10 \times U. rate	-0.130*** (0.0163)	-0.0885*** (0.0123)	-2.286*** (0.120)	-0.374*** (0.0343)
PE $>$ 10 \times U. rate	-0.0248* (0.0111)	-0.0237*** (0.00472)	-1.601*** (0.0808)	-0.0386* (0.0188)
Sample	All	Employed	Unemployed	NILF
R^2	0.012	0.007	0.034	0.027
Wald test (p-value)	6.85e-10	8.83e-6	1.03e-15	1.05e-10
N	21,253,704	14,499,841	1,165,829	5,588,034

Notes: The dependent variable is a binary variable which refers to beginning a job at a new firm. ‘All’ refers to all new hires. ‘EE’ refers to new hires from another job. ‘UE’ refers to new hires from unemployment. ‘NE’ refers to new hires from non-participation. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE 2.3: Exits and Other Flows

	(1)	(2)	(3)	(4)	(5)	(6)
	E exits	EU	EE	EN	NU	UN
PE \leq 10	4.680*** (0.173)	0.785*** (0.0587)	2.319*** (0.0881)	1.575*** (0.130)	6.405*** (0.591)	3.944*** (0.364)
PE \leq 10 \times U. rate	0.154*** (0.0209)	0.176*** (0.0156)	-0.0840*** (0.00973)	0.0622*** (0.0153)	0.314*** (0.0626)	-0.150* (0.0693)
PE $>$ 10 \times U. rate	0.129*** (0.0197)	0.144*** (0.00827)	-0.0279*** (0.00500)	0.0124 (0.0103)	0.409*** (0.0140)	-0.208* (0.0893)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	0.293	8.35e-4	7.93e-6	6.94e-3	0.167	0.115
R^2	0.020	0.007	0.007	0.011	0.019	0.024
N	14,499,841	14,499,841	14,499,841	14,499,841	5,588,034	1,165,829

Notes: ‘E exists’ is a binary variable which refers to a job separation. ‘EU’ refers to a job separation to unemployment. ‘EE’ refers to a job separation ending in a new job. ‘EN’ refers to a job separation to non-participation. ‘NU’ and ‘UN’ refer to transitions between non-participation and unemployment. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE 2.4: Involuntary and Voluntary Separations to Unemployment

	(1)	(2)
	Pr(Involuntary) \times 100	Pr(Voluntary) \times 100
PE \leq 10	0.173* (0.0698)	0.423*** (0.0231)
PE \leq 10 \times U. rate	0.173*** (0.0127)	0.0153*** (0.00370)
PE $>$ 10 \times U. rate	0.128*** (0.00614)	0.0162*** (0.00187)
Wald test (p-values)	4.13e-5	0.802
R^2	0.005	0.002
N	14,499,841	14,499,841

Notes: ‘Involuntary’ is a binary variable which refers to involuntarily separate from a job to unemployment. ‘Voluntary’ refers to voluntarily separating from a job to unemployment. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Estimates include main effects and the EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE 2.5: Hiring over the Business Cycle: Age Versus Education

Outcome: Pr(Hired) \times 100	(1)	(2)	(3)	(4)
College	-0.726*** (0.143)	-0.934*** (0.125)	-1.711*** (0.164)	-1.583*** (0.181)
No College \times U. rate	0.0189 (0.0232)	-0.0892*** (0.0141)	-0.0973*** (0.0142)	-0.0478*** (0.0106)
College \times U. rate	-0.0173 (0.0102)	-0.0935*** (0.0221)	-0.0994*** (0.0232)	-0.0660** (0.0212)
EIER fixed effect	No	Yes	Yes	Yes
Demographic fixed effect	No	No	Yes	Yes
Month-year fixed effect	No	No	No	Yes
Wald test (p-value)	0.0684	0.797	0.906	0.310
R^2	0.001	0.001	0.002	0.004
N	21,253,704	21,253,704	21,253,704	21,253,704

Notes: ‘Hired’ is a binary variable which refers to beginning a job at a new firm. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘College’ refers to having a university degree. Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the Col. \times U. rate and No col. \times U. rate coefficients are statistically distinct.

TABLE 2.6: Job-to-Job Transition over the Business Cycle: Experience and Tenure

Outcome: $\text{Pr}(\text{EE}) \times 100$	(1)	(2)	(3)
$\text{PE} \leq 10$	2.264*** (0.129)		1.688*** (0.0953)
Tenure < 3		1.926*** (0.122)	1.284*** (0.0944)
$\text{PE} \leq 10 \times \text{U. rate}$	-0.0885*** (0.0123)		-0.0765*** (0.00908)
$\text{PE} > 10 \times \text{U. rate}$	-0.0237*** (0.00472)		-0.0177** (0.00528)
Tenure < 3 \times U. rate		-0.0570*** (0.00953)	-0.00815 (0.00910)
Tenure $\geq 3 \times$ U. rate		-0.0230*** (0.00433)	
R^2	0.007	0.008	0.010
Wald test (p-value)	8.83e-6	3.513e-3	2.35e-6
N	14,499,841	14,499,841	14,499,841

Notes: The dependent variable is a binary variable which denotes the happening of a job-to-job transition, re-scaled to 100. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). ‘Tenure’ refers to job tenure (in years). Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald tests for columns (1) (3) indicate whether the $\text{PE} \leq 10 \times \text{U. rate}$ and $\text{PE} > 10 \times \text{U. rate}$ coefficients are statistically distinct. The Wald test for column (2) indicates whether the Tenure < 3 \times U. rate and Tenure $\geq 3 \times$ U. rate coefficients are statistically distinct.

TABLE 2.7: Exits and Other Flows: Tenure Groups

	(1) E exits	(2) EU	(3) EE	(4) EN	(5) NU	(6) UN
Tenure < 3	2.750*** (0.193)	0.135 (0.0752)	2.011*** (0.0755)	0.604*** (0.140)	4.842*** (0.494)	3.376*** (0.203)
Tenure < 3 × U. rate	0.407*** (0.0138)	0.318*** (0.0104)	-0.0519*** (0.00735)	0.141*** (0.0123)	0.728*** (0.0685)	-0.120 (0.0813)
Tenure ≥ 3 × U. rate	-0.00272 (0.0232)	0.0631*** (0.0106)	-0.0281*** (0.00431)	-0.0377** (0.0134)	0.865*** (0.0502)	0.0204 (0.0877)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	3.82e-21	2.58e-31	0.00506	7.06e-12	0.0119	3.31e-10
R ²	0.028	0.013	0.008	0.012	0.021	0.016
N	14,499,841	14,499,841	14,499,841	14,499,841	1,208,085	820,675

Notes: ‘E exists’ is a binary variable which refers to a job separation. ‘EU’ refers to a job separation to unemployment. ‘EE’ refers to a job separation ending in a new job. ‘EN’ refers to a job separation to non-participation. ‘NU’ and ‘UN’ refer to transitions between non-participation and unemployment. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘Tenure’ refers to the tenure of the previous job. Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the Tenure < 3 × U. rate and Tenure ≥ 3 × U. rate coefficients are statistically distinct.

TABLE 2.8: Exits and Other Flows: Experience & Tenure Groups

	(1) E exits	(2) EU	(3) EE	(4) EN	(5) NU	(6) UN
PE \leq 10	4.009*** (0.176)	0.789*** (0.0597)	1.691*** (0.0730)	1.529*** (0.115)	1.537* (0.623)	5.371*** (0.348)
Tenure < 3	1.268*** (0.186)	-0.126 (0.0867)	1.376*** (0.0627)	0.0185 (0.131)	4.530*** (0.711)	1.314*** (0.213)
PE \leq 10 \times U. rate	-0.156*** (0.0301)	-0.0113 (0.0183)	-0.0778*** (0.00734)	-0.0671*** (0.0173)	0.436*** (0.0783)	-0.0516 (0.0739)
PE > 10 \times U. rate	0.0111 (0.0249)	0.0727*** (0.00983)	-0.0246*** (0.00527)	-0.0370* (0.0144)	0.971*** (0.0571)	0.000764 (0.0868)
Tenure < 3 \times U. rate	0.478*** (0.0266)	0.282*** (0.0128)	0.000586 (0.00638)	0.196*** (0.0195)	0.0450 (0.0881)	-0.0708*** (0.0201)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	1.74e-08	3.83e-11	0.00000202	0.0532	4.90e-08	0.123
R^2	0.030	0.013	0.011	0.014	0.024	0.020
N	14,499,841	14,499,841	14,499,841	14,499,841	1,208,085	820,675

Notes: ‘E exists’ is a binary variable which refers to a job separation. ‘EU’ refers to a job separation to unemployment. ‘EE’ refers to a job separation ending in a new job. ‘EN’ refers to a job separation to non-participation. ‘NU’ and ‘UN’ refer to transitions between non-participation and unemployment. ‘U. rate’ refers to the Employment Insurance Economic Region (EIER) unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). ‘Tenure’ refers to the tenure of the previous job. Estimates include main effects and EIER, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the EIER level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE > 10 \times U. rate coefficients are statistically distinct.

TABLE 2.9: Summary of Experience vs. Tenure Contributions to Flow Cyclicity

Flow Type	Experience (p.p.)	Tenure (p.p.)	Dominant Factor (% of the cyclicity explained)
Job-to-job (EE)	0.0532	0.0006	Experience (99%)
Employment to Unemployment (EU)	0.084	0.282	Tenure (77%)
Employment to Non-participation (EN)	0.0301	0.196	Tenure (87%)

Notes: Each cell in the Experience and Tenure columns reports the marginal effect of having low experience and having short tenure on the cyclicity of separation probabilities (i.e. absolute value of the difference in interaction coefficients between subgroups, based on regressions that include both experience-by-unemployment and tenure-by-unemployment interaction terms, Table 2.8). Specifically, the Experience value is calculated as the absolute difference between the coefficients on PE \leq 10 \times Unemp. Rate and PE > 10 \times Unemp. Rate. The Tenure value is the absolute value of the coefficient on Tenure < 3 \times Unemp. Rate, with Tenure \geq 3 as the omitted group. The “Dominant Factor” column reports the share of total explained variation attributed to the larger of the two effects, calculated as the ratio of each factor’s effect to the sum of both (e.g., $0.0532 / (0.0532 + 0.0006) \approx 99\%$ for job-to-job transitions). This summary is intended to provide an intuitive comparison of relative cyclical sensitivity and does not reflect a formal decomposition of model variance.

Chapter 3

The Gender Unemployment Gap in Canada

3.1 Introduction

In Canada, the gender unemployment gap—defined as women’s unemployment rates minus men’s unemployment rates—has undergone a significant change over the decades.¹ Before 1990, most of the time, women experienced higher unemployment rates than men. Since then, the trend has reversed, with women maintaining lower unemployment rates than men (Figure 3.1).² During this time, women’s labour force participation and employment have steadily converged toward men’s levels, though they still remain lower (Figure 3.2). While [Albanesi and Şahin \(2018\)](#) find that the convergence in labour force participation can explain the closure of the gender unemployment gap in the U.S., it is still not clear why women’s unemployment rates have not only decreased to match men’s levels but now consistently remain lower than men’s in Canada.

The unemployment rate at any given time is determined by the flows into and out of unemployment. Thus, the gender unemployment gap can be explained by gender differences in these transition flows. Prior research has shown that differences in unemployment rates between groups stem from differences in inflows to and outflows from unemployment (e.g., [Baker, Corak and Heisz \(1998\)](#); [Azmat, Güell and Manning \(2006\)](#); [Forsythe and Wu \(2021\)](#)). [Elsby et al. \(2015\)](#) emphasize that flows involving non-participation, i.e. movements between unemployment and non-participation (UN and NU), account for roughly one-third of unemployment rate fluctuations. To examine the evolution of the gender unemployment gap in Canada, I decompose it

¹In the gender pay gap literature, the gender gap is usually defined as men’s earnings minus women’s earnings or women’s earnings as a fraction of men’s earnings. However, as women’s unemployment rates are historically higher than men’s unemployment rates in many OECD countries, here I define the gender unemployment gap as women’s unemployment rate minus men’s unemployment rate, following [Azmat et al. \(2006\)](#) and [Albanesi and Şahin \(2018\)](#).

²Historical data from [Usalcaas and Kinack \(2017\)](#) indicate that the gender unemployment rate gap was negative in Canada from 1946 to 1968. However, unemployment data in the Labour Force Survey (LFS) before and after 1976 are not directly comparable due to changes in survey design. Prior to 1976, the LFS classified individuals’ labour force status based on their main activity. Those who did not report job-seeking as their primary or secondary activity were not counted as unemployed. For instance, an unemployed woman who reported homemaking as her main activity would not have been classified as unemployed, whereas under today’s classification, she would be.

into the contributions of six transition flows across the three labour market states: employment (E), unemployment (U), and non-participation (N).

The decomposition reveals two main findings. First, transitions between employment and non-participation (EN and NE flows) have been positive contributors to the gender unemployment gap, but their impact has diminished over time. Before 1990, women had much higher EN transition rates and lower NE transition rates than men, contributing to a positive gender unemployment gap. However, as women's labour force participation has gradually converged toward men's (Figure 3.2), the gender differences in these flows have also narrowed. This pattern is similar to what [Albanesi and Şahin \(2018\)](#) document for the U.S. Second, employment-to-unemployment (EU) transitions have consistently acted as a negative contributor to the gender unemployment gap, with men exhibiting higher EU rates than women, particularly during recessions.

Taken together, the gender unemployment gap reflects the combined effects of these six transition flows. Before 1990, the positive contributions of EN and NE flows outweighed the negative effects of EU flows, keeping the gap positive. However, as the impact of EN and NE flows has decreased over time, the persistent negative contribution of EU transitions has become dominant, leading to the reversal of the gender unemployment gap from positive to negative. Since 1990, men's higher EU separation rates have been the primary driver of the negative gender unemployment gap.

To better understand the gender difference in EU transition rates, I investigate the extent to which this disparity can be explained by differences in job characteristics. After accounting for industry and occupation composition, the gender gap reverses sign—suggesting that women's lower average EU rates are largely driven by their concentration in more stable sectors. Within similar job categories, however, women tend to have shorter tenure and are more likely to work part-time, both of which are associated with higher job turnover. Together, industry and occupation composition, part-time status, and tenure account for 69.6% of the observed gender difference. Nonetheless, a residual gap remains that may reflect unobserved job characteristics, labour supply preferences, or potential discrimination.

This paper contributes to the literature on the gender gap by broadening the scope beyond the well-studied gender pay gap (e.g., [Goldin \(1990, 2006, 2014, 2024\)](#); [Goldin, Kerr, Olivetti and Barth \(2017\)](#); [Blau and Kahn \(1994, 2000, 2003, 2006, 2017\)](#); [Blau, Kahn, Boboshko and Comey \(2024\)](#); [Fortin, Bell and Böhm \(2017\)](#); [Fortin \(2019\)](#); [Olivetti, Pan and Petrongolo \(2024\)](#)). Previous studies using Oaxaca-Blinder decomposition have shown that factors such as human capital (including education and labour market experience), industry, and occupation play key roles in explaining the gender pay gap ([Blau and Kahn, 2000, 2017](#); [Goldin et al., 2017](#)). However, far less attention has been paid to the gender unemployment gap. Notable exceptions include [Azmat et al. \(2006\)](#), who examine cross-country differences in the gender unemployment gap in OECD countries; [Albanesi and Şahin \(2018\)](#), who study the convergence in the gender unemployment gap in the U.S.; and [Myatt and Murrell \(1990\)](#), who study the Canadian context with data ending in 1987—before men's unemployment rates began to exceed women's. This

paper addresses this gap by explaining the evolution of the gender unemployment gap in Canada in recent decades. Unlike [Azmat et al. \(2006\)](#), who examine cross-country differences in the gender unemployment gap, this paper focuses on its evolution over time in Canada, showing how gender differences in labour flows have led to the reversal of its sign.

In addition, this paper applies insights from the cyclical unemployment fluctuation literature (e.g., [Blanchard and Diamond \(1990\)](#); [Davis, Faberman and Haltiwanger \(2006\)](#); [Yashiv \(2007\)](#); [Elsby, Michaels and Solon \(2009\)](#); [Fujita and Ramey \(2009\)](#); [Shimer \(2012\)](#); [Elsby, Hobijn and Şahin \(2013, 2015\)](#); [Elsby, Hobijn, Karahan, Koşar and Şahin \(2019\)](#)) to study the gender unemployment gap. Similar applications, such as explaining group differences in unemployment rates using labour flows, include [Baker, Corak and Heisz \(1998\)](#), [Jones and Riddell \(1998\)](#), [Lauerova and Terrell \(2007\)](#), [Koutentakis \(2015\)](#), [Lydon and Simmons \(2020\)](#), and [Forsythe and Wu \(2021\)](#). This paper borrows from the decomposition method in [Forsythe and Wu \(2021\)](#) and applies it to study the gender unemployment gap in Canada. While [Forsythe and Wu \(2021\)](#) focus on how cyclical fluctuations in unemployment rates of different groups are explained by flows and the relative contributions of flows on level differences in the unemployment rate between two groups, this paper emphasizes the time path of each flow's contribution and how it explains the flip in the gender unemployment gap.

The remainder of the paper is organized as follows. Section [3.2](#) presents the trends and cyclical patterns of unemployment for men and women. Section [3.3](#) investigates how changes in age and education composition of the labour force have affected the gender unemployment rate gap. In Section [3.4](#), a flow decomposition analysis is used to determine the contribution of each labour market transition to the gender unemployment rate gap. Section [3.5](#) presents the decomposition results. Section [3.6](#) examines how employment composition in terms of industry, occupation, and part-time work status explains the women-men gap in employment-to-unemployment transitions. Finally, section [3.7](#) concludes.

3.2 Trend and Cyclical Components of Unemployment Gap

This study relies on the Canadian Labour Force Survey (LFS) covering the period from 1976 to 2019. The LFS is a large monthly household survey that collects comprehensive labour market data for Canadian residents. Around 55,000 households are interviewed each month. It serves as the official source of employment and unemployment statistics for Canada.

Using the unemployment information in the LFS, [Figure 3.3](#) illustrates the trend and cyclical components of the unemployment rates for men and women from 1976 to 2019 (filtered by HP filter). Panel (a) of [Figure 3.3](#) shows that the trend component of women's unemployment rate is higher than men's from January 1976 to May 1983 and from September 1984 to September 1989. There was a brief period from June 1983 to August 1984 when the trend of women's unemployment rate was slightly lower — nearly equal to men's. Since October 1989, however, the trend component of the unemployment rate for women has consistently been lower than that for men, continuing through to the end of 2019. This pattern closely mirrors the overall gender unemployment rate gap. Before 1990, the men's unemployment rate was, on average,

0.81 percentage points lower than the women's, with the trend explaining 94.75% of this gap. From 1990 to 2019, men's unemployment rate was, on average, 0.88 percentage points higher than women's, with 97.76% of this difference attributed to the trend component.

Panel (b) of Figure 3.3 shows that while the cyclical component of unemployment for men and women displays minimal differences in general between 1976 and 2019, men's cyclical unemployment rates rise more sharply than women's during recessions. In the 1981-82 recession, 61.6% of the gender unemployment gap was explained by differences in their cyclical unemployment rates. Similarly, during the 1990-92 recession, 55.6% of the gap was due to the cyclical component. During the 2008-09 recession, the cyclical component accounted for 26.5% of the gender unemployment gap, which was a smaller share compared to previous recessions.

In summary, Figure 3.3 suggests that the gender unemployment rate gap is largely driven by differences in the trend component of unemployment rates, though the cyclical component plays a significant role in explaining the gap during recessions.

3.3 Labour Force Compositions and Gender Unemployment Rate Gap

As shown in section 3.2, the gender unemployment rate gap is primarily driven by long-term trends in the unemployment rates of men and women. In this section, I explore whether these trends are influenced by demographic shifts over time, specifically changes in age and educational attainment within the labour force.

3.3.1 Age Composition

Given that younger workers typically face higher unemployment rates and that women in the labour force have historically been younger than their male counterparts, one might assume that age composition plays a significant role in driving the gender unemployment gap. Figure 3.4, panel (a), shows that women in the labour force have been, on average, younger than men. Since 1980, the average age of both male and female workers has steadily increased, with the age gap between them narrowing over time. Prior to 1990, female workers were, on average, 2.2 years younger than male workers, which could help explain why women had higher unemployment rates during that period. However, since 1990, the average age gap has stabilized at around 0.9 years, while women's unemployment rates have been lower than men's.

To assess whether changes in the age composition of the labour force contribute to the gender unemployment rate gap, I apply the methodology of Albanesi and Şahin (2018), adjusting the female unemployment rate by estimating what it would be if women had the same age composition as men. I divide the labour force into three age groups: youth (15-24), prime working-age (25-54), and older workers (55+). Let $A = \{15-24, 25-54, 55+\}$ represent the age groups. For each group, I calculate the unemployment rates for men and women at time t , denoted as $u_t^m(i)$ and $u_t^f(i)$, where i refers to the age group. The male unemployment rate is then computed as $u_t^m = \sum_{i \in A} s_t^m(i) u_t^m(i)$, where $s_t^m(i)$ is the proportion of men in age group i in the labour force. Similarly, the female unemployment rate is calculated as $u_t^f = \sum_{i \in A} s_t^f(i) u_t^f(i)$, where $s_t^f(i)$ represents the share of women in each age group. To compute the counterfactual

female unemployment rate, assuming women had the same age composition as men, I use $\tilde{u}_t^f = \sum_{i \in A} s_t^m(i) u_t^f(i)$.

Panel (b) of Figure 3.4 compares the actual and counterfactual gender unemployment gaps. It shows that if the female labour force were older, as in the case of men, women's unemployment rates would be slightly lower. As a result, the counterfactual gap is below the actual gender unemployment gap. However, on average, age composition can only explain about 11% of the gender unemployment gap. After 1990, the actual and counterfactual gaps are nearly identical. This indicates that changes in the age composition of men and women in the labour force do not significantly explain the evolution of the gender unemployment rate gap in recent decades.

3.3.2 Education Composition

Since individuals with lower education levels typically experience higher unemployment rates, one might assume that differences in educational attainment between men and women contribute to the gender unemployment rate gap and its evolution over time. To examine this conjecture, I present the average years of schooling for men and women in the labour force from 1976 to 2019 in panel (a) of Figure 3.5.³ The average years of schooling are calculated as a weighted average of the schooling years of each education group for both men and women. The weights are the shares of individuals in a certain education group.⁴

I categorize the labour force into four education groups: below high school, high school, post-secondary certificate or diploma, and university degree or higher. For the calculation, following Albanesi and Şahin (2018), I assign 10 years of schooling for the “below high school” group, 12 years for the “high school” group, 14 years for the “post-secondary certificate or diploma” group, and 18 years for the “university degree or higher” group. Using the shares of men and women in each education category as weights, I then compute the weighted average years of schooling for both genders. Figure 3.5, panel (a), reveals that women have, on average, more years of schooling than men, with this gap widening over time. Before 1992, the difference in average years of schooling between men and women was not significant. However, by 2019, women were, on average, half a year more educated than men.

To determine if changes in the education composition of the labour force can explain the gender unemployment gap, I calculate a counterfactual female unemployment rate by assuming women have the same educational distribution as men.⁵ Figure 3.5, panel (b), compares the actual and counterfactual gender unemployment gap. The results show that if women had lower education levels similar to men's, their unemployment rate would be higher, resulting in a counterfactual gap that lies above the actual gap. On average, education composition explains 12% of the

³Panel (a) of Figure 3.5 shows the average years of schooling for individuals aged 25 and above to reflect those who have completed their education. However, in panel (b) of Figure 3.5, the analysis includes the entire working-age population for consistency.

⁴In the LFS, the education variable before 1990 records the number of years of schooling completed. After 1990, it reflects the highest level of educational attainment. Individuals with 11 to 13 years of schooling before 1990 did not necessarily have a high school diploma; however, they are treated as high school graduates in this analysis.

⁵The method I use to calculate the counterfactual female unemployment rate here is similar to the approach used to calculate it under the assumption that women had the same age composition as men.

gender unemployment gap. However, this effect is not substantial enough to account for the persistent negative gender unemployment gap observed since 1990, and the flip of the sign of the gender unemployment gap.

3.4 Decomposing the Unemployment Rate Gap into Difference in Flows

Section 3.3 shows that changes in demographic composition do not drive the evolution of the gender unemployment rate gap. In this section, I examine how the dynamics of labour market flows impact this gap between men and women. The fact that men exhibit higher employment levels but also higher unemployment rates than women may be explained by higher inflows to and lower outflows from unemployment for men. To determine which flows are the primary contributors to the gender unemployment gap, I apply a stock-flow decomposition to quantify the contribution of each flow.

To calculate the six labour flows across the three labour market states—employed (E), unemployed (U), and not in the labour force (N)—I use the LFS master files to match the same individuals over two consecutive months, and I exclude cases with age and gender inconsistencies.⁶ I adjust the monthly transition rates for seasonality, and adjust the rates for margin error following [Elsby et al. \(2015\)](#). For the decomposition, I convert the discrete-time transition probabilities to flow hazard rates following the method in [Shimer \(2012\)](#).

Figure 3.6 presents the monthly transition rates for men and women from 1976 to 2019. Panels (a) and (c) of Figure 3.6 illustrate transitions between employment and unemployment, while panels (b), (d), (e) and (f) show transitions at the participation margin. Panel (b) shows that women's EN transition rates have been higher than men's, though the gap has declined over time, especially during the late 1970s and 1980s. The gender gap in EN transition rates decreased from 2 percentage points in 1976 to about 0.5 percentage points by 2019. Panel (e) indicates that men generally have higher NE rates than women, with the gap narrowing from 3 percentage points in 1976 to about 0.5 percentage points by 2019.

Similarly, panel (d) of Figure 3.6 shows that women's UN rates are higher than men's, although the gap has been decreasing over the past four decades. Panel (f) shows that both men's and women's NU rates follow a cyclical pattern—both are increasing during recessions. Men's NU rates are higher than women's, but the gender gap in NU rates has been converging over time. Finally, panels (a) and (c) of Figure 3.6 reveal cyclical patterns in EU and UE transitions for both genders, with men showing slightly higher UE transition rates, particularly before 1981, and persistently higher EU rates, with this gap widening during recessions and stabilizing around 0.36 percentage points after 1990.

Overall, Figure 3.6 shows how the gender gap has evolved across the six transition rates. The convergence in the four transitions at the participation margin (EN, NE, UN, and NU) aligns with the convergence in labour force participation rates between men and women. In the

⁶For the details for the construction of the flow data and the exclusion of age and gender inconsistency, please see Appendix A.1.

following analysis, I decompose the gender unemployment rate gap into contributions from each transition flow, illustrating how this gap shifted from positive to persistently negative after 1990.

To perform the flow decomposition, I approximate the unemployment rate using its steady-state form, which is a function of the six transition rates. Changes in the employment rate (e_t) and the unemployment-to-population rate (u_t) at time t can be expressed as follows:

$$\Delta e_t = -(\lambda_t^{EU} + \lambda_t^{EN})e_{t-1} + \lambda_t^{UE}u_{t-1} + \lambda_t^{NE}(1 - e_{t-1} - u_{t-1}) \quad (3.1)$$

$$\Delta u_t = -(\lambda_t^{UE} + \lambda_t^{UN})u_{t-1} + \lambda_t^{EU}e_{t-1} + \lambda_t^{NU}(1 - e_{t-1} - u_{t-1}), \quad (3.2)$$

where λ_t^{JK} denotes the flow rate from state J to K from time $t-1$ to t , with $J, K \in \{E, U, N\}$.⁷ In each equation, the first term on the right-hand side represents outflows, while the second and third terms represent inflows.

Equations (3.1) and (3.2) can be rewritten in matrix form as:

$$\Delta s_t = P_t \cdot s_{t-1} + q_t, \quad (3.3)$$

$$\text{where } s_t = \begin{bmatrix} e_t \\ u_t \end{bmatrix}, P_t = \begin{bmatrix} -\lambda_t^{EU} - \lambda_t^{EN} - \lambda_t^{NE} & \lambda_t^{UE} - \lambda_t^{NE} \\ \lambda_t^{EU} - \lambda_t^{NU} & -\lambda_t^{UE} - \lambda_t^{UN} - \lambda_t^{NU} \end{bmatrix}, \text{ and } q_t = \begin{bmatrix} \lambda_t^{NE} \\ \lambda_t^{NU} \end{bmatrix}.$$

As in the steady state, there is no change in e_t and u_t ($\Delta s_t = 0$); one has the steady-state levels of e_t and u_t as $\bar{s} = \begin{bmatrix} \bar{e}_t \\ \bar{u}_t \end{bmatrix} = (-P_t)^{-1} \cdot q_t$. The steady-state unemployment rate (u_t^{ss}) can thus be expressed as:

$$u_t^{ss} = \frac{\bar{u}_t}{\bar{e}_t + \bar{u}_t} = \frac{\lambda_t^{NE}\lambda_t^{EU} + \lambda_t^{EN}\lambda_t^{NU} + \lambda_t^{EU}\lambda_t^{NU}}{(\lambda_t^{UE}\lambda_t^{NE} + \lambda_t^{UN}\lambda_t^{NE} + \lambda_t^{NU}\lambda_t^{UE}) + (\lambda_t^{NE}\lambda_t^{EU} + \lambda_t^{EN}\lambda_t^{NU} + \lambda_t^{EU}\lambda_t^{NU})} \quad (3.4)$$

Equation (3.4) links the relationship between the unemployment rate and six flow rates. I use equation (3.4) to approximate the unemployment rates for both men and women, and then decompose the gender unemployment rate gap into the contribution of the gender differences in six flow rates.

Figure 3.7 evaluates the fit of the steady-state approximation to the actual unemployment rates for men and women. For men, panel (a) shows that the steady-state approximation aligns closely with the actual unemployment rate, with a correlation of 0.97 and an R^2 value of 0.94. Similarly, panel (b) demonstrates that the steady-state approximation captures women's unemployment rate well, with a correlation of 0.98 and an R^2 of 0.96. This good fit suggests that the steady-state approximation is reliable in approximating the unemployment rate, allowing me to further analyze the contributions of each flow to the gender unemployment rate gap.

Taking a first-order Taylor approximation of the natural log of women's steady-state unemployment rate near the point of men's steady-state unemployment rate and flow rates, one

⁷For example, λ^{EU} denotes the flow rate from employment to unemployment from time $t-1$ to time t .

has:

$$\begin{aligned} \ln u_t^{ss,f} &\approx \ln u_t^{ss,m} + \sum_{i \in JK} \left. \frac{\partial \ln u_t^{ss,f}}{\partial \lambda_t^i} \right|_{\Lambda_t = \Lambda_t^m} \cdot (\lambda_t^{i,f} - \lambda_t^{i,m}) \\ &\approx \ln u_t^{ss,m} + \sum_{i \in JK} \left. \frac{\partial \ln u_t^{ss,f}}{\partial \lambda_t^i} \right|_{\Lambda_t = \Lambda_t^m} \cdot \lambda_t^{i,m} \cdot (\ln \lambda_t^{i,f} - \ln \lambda_t^{i,m}), \end{aligned} \quad (3.5)$$

where $\lambda_t^{i,m}$ and $\lambda_t^{i,f}$ denote flow rate i for men and women at time t , respectively. Here, $JK = \{EU, EN, UE, UN, NE, NU\}$ represents the set of the six transition flows, and Λ_t^f and Λ_t^m are the vectors of these flows for women and men at time t , respectively.

Since the unemployment rate of men and women can be well approximated by their steady-state levels, the log unemployment rate gap between women and men can be expressed as:

$$\ln u_t^f - \ln u_t^m \approx \sum_{i \in JK} \left. \frac{\partial \ln u_t^f}{\partial \lambda_t^i} \right|_{\Lambda_t = \Lambda_t^m} \cdot \lambda_t^{i,m} \cdot (\ln \lambda_t^{i,f} - \ln \lambda_t^{i,m}). \quad (3.6)$$

Equation (3.6) indicates that the gender unemployment rate gap can be written as the sum of the contributions of gender differences in six flow rates. Each flow's contribution depends on its marginal effect on the steady-state unemployment rate and the gender difference in that flow.

That is, the gender unemployment gap can be decomposed as:

$$F_t^{gap} = F_t^{EU} + F_t^{EN} + F_t^{UE} + F_t^{UN} + F_t^{NE} + F_t^{NU} + \epsilon_t, \quad (3.7)$$

where F_t^{gap} represents the log difference between women's and men's unemployment rates. On the right-hand side of the equation, for instance, F_t^{EU} represents the contribution of the gender difference in EU flow rate to F_t^{gap} . The term ϵ_t captures the portion of the gender unemployment rate gap that is not explained by this flow decomposition.

Figure 3.8 compares the actual gender unemployment gap with the sum of the contributions of the six flows (i.e. the left-hand side versus the right-hand side of equation (3.6)). The explained gap aligns closely with the actual gap, with a correlation of 0.98 and an R^2 value of 0.96. It suggests that equation (3.6) effectively approximates the gender unemployment rate gap.

3.5 Decomposition Results

3.5.1 Contribution of Each Individual Flow

Figure 3.9 provides a summary decomposition of the gender unemployment rate gap over time, breaking it down into the contributions from each individual labour market flow.⁸ The EU flow—transitions from employment to unemployment—emerges as the dominant driver since the early 1990s, closely tracking the overall gap and mirroring its cyclical fluctuations. Flows

⁸Figure 3.9 provides an integrated view of these flow-level contributions. For detailed graphs of each flow's standalone contribution, see Appendix Figures C.1–C.3.

between employment and non-participation (EN and NE) contribute positively to the gap in earlier decades, but their role has diminished over time. Contributions from flows between unemployment and non-participation (UN and NU) are smaller and negative, while differences in job-finding rates (UE flow) have had a limited impact on the gap since the 1980s.

These patterns reflect the underlying differences in transition rates between men and women. The EN flow increases the unemployment rate by shrinking the labour force while holding unemployment constant. Since women historically had higher EN flow rates, this transition increased their unemployment rate relative to men, thereby widening the gap. However, as gender differences in EN flows have narrowed in recent decades, this effect has become less pronounced.

Similarly, men were more likely than women to move from non-participation into employment (higher NE flows), which increased the size of their labour force and reduced their unemployment rate. As a result, the lower NE flow also contributed to a higher unemployment rate for women, positively contributing to the gender gap. However, the difference in NE flow rates between men and women has gradually declined over time, and so has its contribution to the gap.

UN and NU flows also reduced the gender gap but with declining influence over time. In contrast, gender differences in the UE flow have largely converged since 1980, resulting in minimal contribution to the gap. The EU flow is now the primary driver: women's lower separation rates from employment consistently reduce their unemployment rate relative to men, especially during recessions when men's separations spike.

3.5.2 Evolution of the Gender Unemployment Gap

The contributions of flows at the participation margin—EN, NE, UN, and NU—have declined over time as gender differences in these transitions have narrowed. To better understand how these changes in flows have shaped the evolution of the overall gender unemployment gap, I conduct a series of counterfactual exercises that isolate the role of flows in driving the observed decline in the gender unemployment gap.

Figure 3.10 compares the actual gender unemployment gap with a counterfactual scenario where women's flows between employment and non-participation remain at pre-1980 levels. The results indicate that if women's EN rates had not decreased, and their NE rates had not increased, the gender unemployment gap would have remained positive (around one percentage point) until 2019. This suggests that the convergence in EN and NE transitions played a key role in the shift from a positive to a negative gender unemployment gap.

Similarly, Figure 3.11 examines a counterfactual scenario where women's flows between unemployment and non-participation remain at pre-1980 levels. The results show that if UN and NU flows had not converged, the gender unemployment gap would have become negative right after 1980, and the magnitude of the gap would have been larger than the actual one. Together with Figure 3.10, these results indicate the significant role of converging labour force participation in the evolution of the gap.

Finally, Figure 3.12 presents a counterfactual scenario to explain the persistent negative gender unemployment gap since 1990. Assuming that women's EU flow rates had not decreased and remained at pre-1990 levels, the gap would have been near zero, except during recessions. This suggests that women's lower EU rates compared to men are the main contributors to the persistence of the negative gap since 1990.

In summary, the convergence of flows at the labour force participation margin—EN, NE, UN, and NU—explains the long-term decline in the gender unemployment gap, particularly the shift from a positive gap in the 1980s to a negative one by the 1990s. In contrast, the persistent negative gap observed since 1990 is primarily driven by women's lower EU flow rates compared to men. These findings indicate that while participation flows explain the trend change, employment-to-unemployment separations account for the persistence of the negative gap in more recent decades.

3.6 Explaining the Gender Difference in EU Flow

Knowing that women's lower EU flow rates are the main driver of their lower unemployment rates compared to men since 1990, this section examines the factors contributing to this difference.

Table 3.1 examines how the gender gap in EU transitions changes when progressively adding different fixed effects. The sample period is from 1987 to 2019. Column (1) shows that without any controls, women's EU transition rates are 0.362 percentage points lower than men's. In column (2), controlling for age and education fixed effects reduces the gender gap slightly, but it remains at 0.358 percentage points, indicating that differences in age and education composition do not explain the gap.

In column (3), adding industry and occupation fixed effects results in a significant change.⁹ The coefficient for the female dummy becomes positive (0.062 percentage points), suggesting that women's concentration in stable industries and occupations accounts for most of their lower EU rates compared to men. However, when the part-time status and tenure are included in column (4), the coefficient for the female dummy becomes negative again (-0.110 percentage points). This indicates that women's higher likelihood of working part-time, which is associated with greater job turnover, explains a portion of their higher EU rates within the same industry and occupation. Comparison between column (4) and column (1) suggests that after taking into account industry and occupation composition, part-time status and job tenure, only 30.4% of the initial gender gap in EU transitions remains.

Since occupation information in the LFS has been available only since 1987, I do another regression analysis without controlling for occupation to have a complete picture from 1976 to 2019. The results are shown in Table C.1. Without controlling for age and education fixed effects, industry fixed effects, and part-time status, the gender gap in EU transitions is -0.343

⁹Industry classification is based on the 2-digit codes from the North American Industry Classification System (NAICS) 2017, while occupation classification follows the 1-digit codes from the National Occupational Classification (NOC) 2016.

percentage points in column (1), slightly smaller than in Table 3.1. Adding age and education fixed effects in column (2) has a minimal impact, with the gap changing to -0.355 percentage points. In column (3), controlling for industry fixed effects reduces the gender gap significantly by 80.4%. Finally, including part-time status and tenure in column (4) leads to a negative gap of -0.258 percentage points. That is, 24.8% of the difference in EU transition rates between men and women can be explained by industry fixed effects, part-time status and deviation from industry average tenure. Without controlling for occupation, there is a larger unexplained portion of the EU gender gap compared to Table 3.1.

3.7 Conclusion

This paper examines the evolution of the gender unemployment gap in Canada over the past four decades, focusing on its evolution from positive to negative and identifying the key contributors driving this change. Using data from the Canadian Labour Force Survey and a stock-flow decomposition, I analyze how six labour market flows contribute to the gender unemployment rate gap.

The findings reveal two distinct phases in the evolution of the gap. Before 1980, the positive gender unemployment gap was primarily driven by gender differences in employment-to-non-participation transitions, as women were more likely than men to exit the labour force. Since 1990, the persistent negative gap has been largely explained by gender differences in employment-to-unemployment flows, with women experiencing lower separation rates than men.

Further analysis shows that industry and occupation composition, part-time status, and tenure play a significant role in explaining the gender difference in employment-to-unemployment flows. Women's lower separation rates are largely attributed to their higher concentration in stable industries and occupations. However, within the same industry and occupation, women exhibit slightly higher separation rates than men, likely due to their greater likelihood of holding part-time jobs and having lower tenure. Collectively, these factors account for 69.6% of the gender difference in employment-to-unemployment flows.

As the decomposition results indicate that women's increasing labour force participation explains the shift from a positive to a negative gender unemployment gap, policies supporting women's labour market participation, such as maternity leave and childcare programs, may have played a role in improving women's employment outcomes over time. However, further research is needed to assess the direct impact of these policies and to examine other potential factors, such as job quality, that may contribute to gender disparities in the labour market. While [Azmat, Güell and Manning \(2006\)](#) finds that the flows between employment and unemployment explain the emergence of gender unemployment gaps in OECD countries with high gender disparities, this paper demonstrates that in Canada, it is the sustained negative contribution of employment-to-unemployment transitions, along with the convergence in women's labour participation, that explains the reversal of the gender unemployment gap. Future research will apply this decomposition framework to other countries to compare how gender unemployment gaps have evolved across different labour markets.

3.8 Figures and Tables

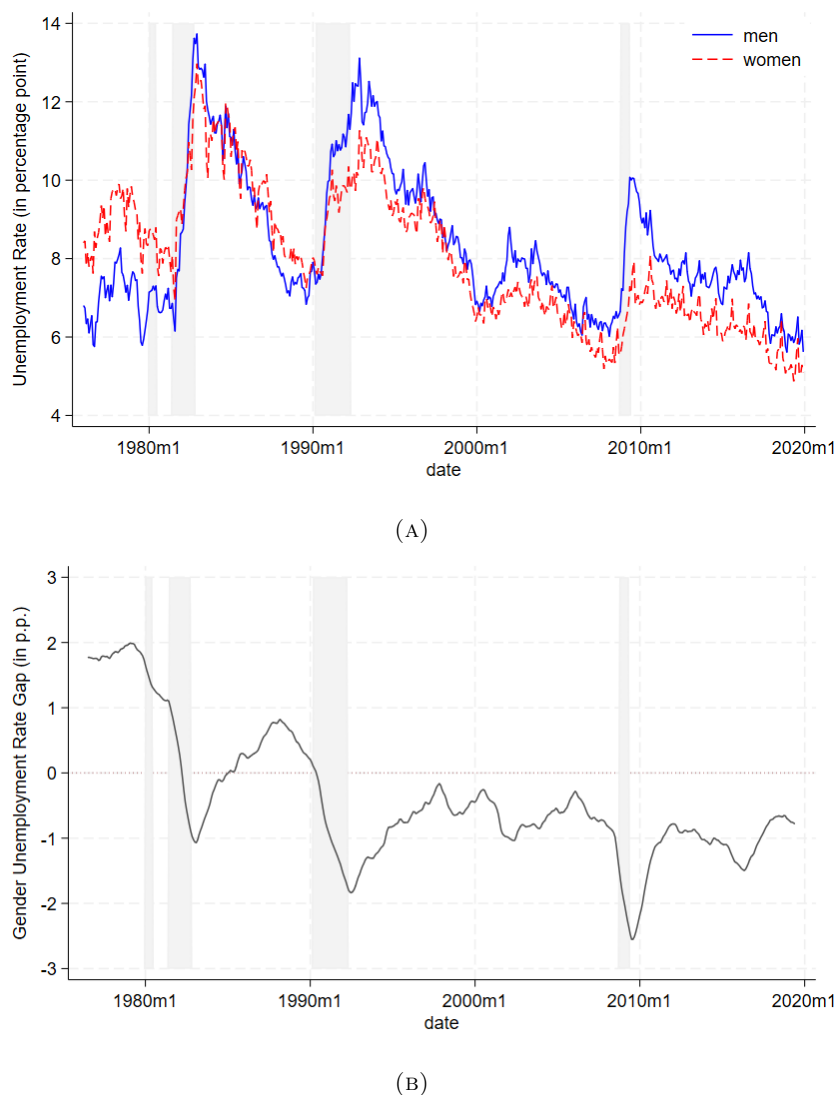

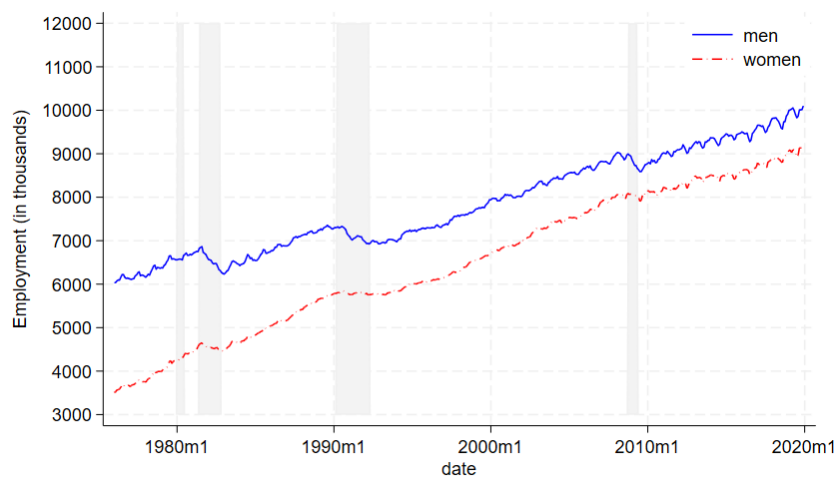


FIGURE 3.1: Gender Unemployment Rate Gap, 1976-2019. Panel (a) is the monthly unemployment rate of the working-age population by gender, and panel (b) is the centered 12-month moving average of the gender unemployment rate gap, defined as the difference between women's unemployment and men's unemployment rate. The rates are weighted using LFS sampling weights. Shaded areas indicate dates of recent recessions. Data Source: the LFS public-use files.

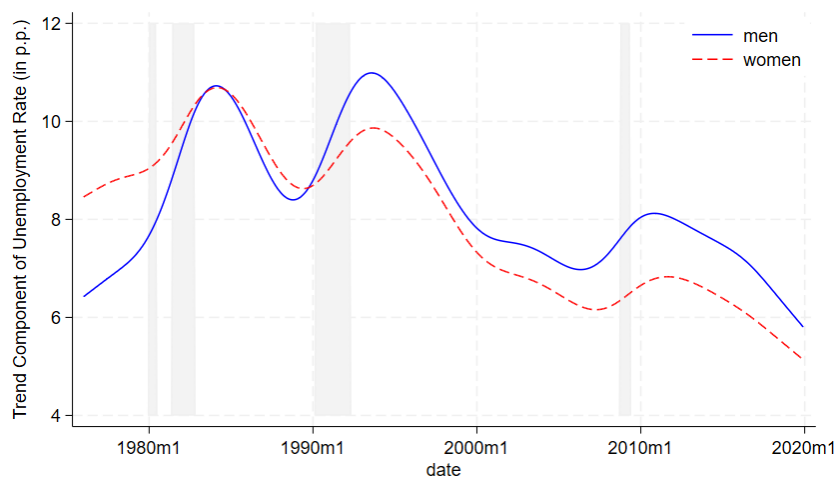


(A)




(B)

FIGURE 3.2: Labour Force Participation and Employment of Men and Women, 1976-2019. Panel (a) is the monthly participation rate of the working-age population by gender, and panel (b) is the monthly employment by gender. The participation rates and employment are weighted using LFS sampling weights. Shaded areas indicate dates of recent recessions. Data Source: the LFS public-use files.




(A)




(B)

FIGURE 3.3: Trend and Cyclical Component of Unemployment Rate of Men and Women, 1976-2019. Panel (a) is the trend component of the unemployment rate, while panel (b) is the cyclical component of the unemployment rate. The time trend and cyclical component of the unemployment rate are separated using the HP filter, with the smoothing parameter 10^5 . Shaded areas indicate dates of recent recessions. Data Source: the LFS public-use files.



(A)



(B)

FIGURE 3.4: Age-Adjusted Gender Unemployment Rate Gap, 1976-2019. Panel (a) is the average age of men and women in the labour force, while panel (b) is the actual and counterfactual gender unemployment gap. The gender unemployment gap is defined as the difference between female and male unemployment rates. The counterfactual unemployment rate is calculated by assuming women have the same age composition as men. The share of each age group for men and women and unemployment rate of each age group are weighted using the LFS sampling weight. Shaded areas indicate dates of recent recessions. Data Source: panel (a), the LFS master files; panel (b), the LFS public-use files.

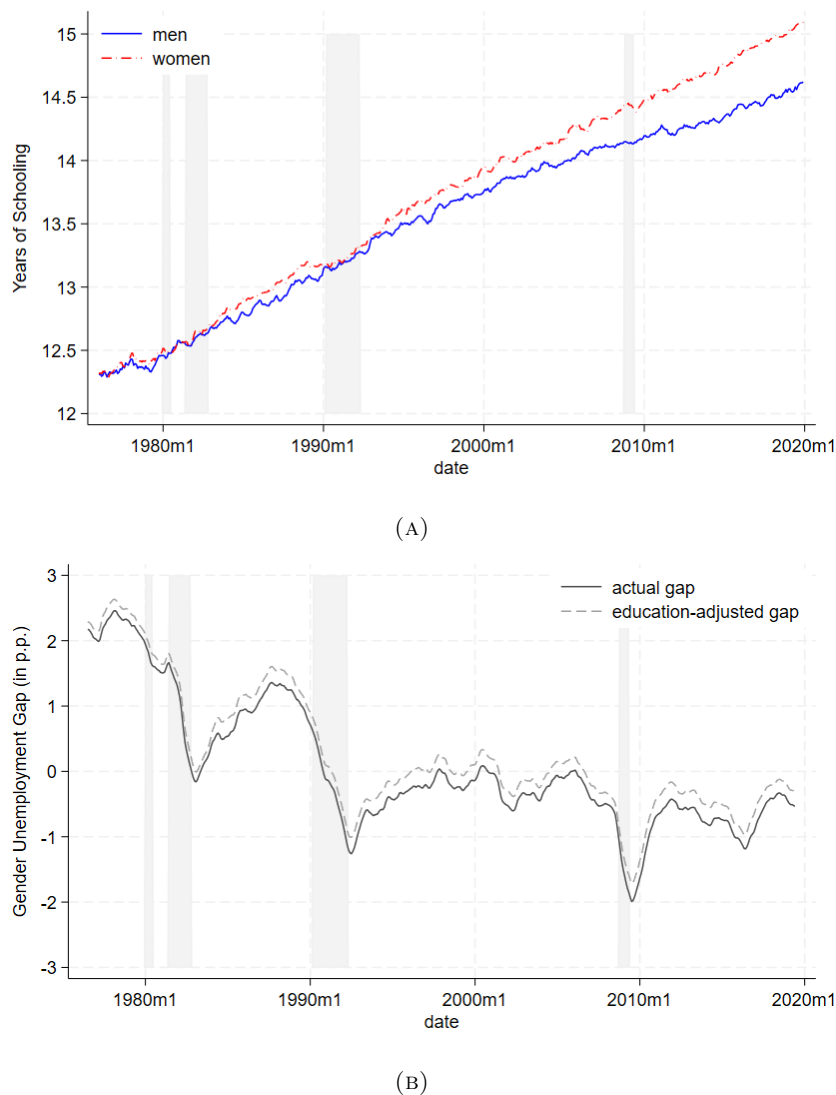


FIGURE 3.5: Education-Adjusted Gender Unemployment Rate Gap, 1976-2019. The sample consists of individuals aged 25 and above for complete education attainment. Panel (a) is the average years of schooling of men and women in the labour force, while panel (b) is the actual and counterfactual gender unemployment gap. The gender unemployment gap is defined as the difference between female and male unemployment rates. The counterfactual unemployment rate is calculated by assuming women have the same education attainment distribution as men. The share of each educational attainment group for men and women and unemployment rate of each educational attainment group are weighted using the LFS sampling weight. Shaded areas indicate dates of recent recessions. Data Source: the LFS public-use files.

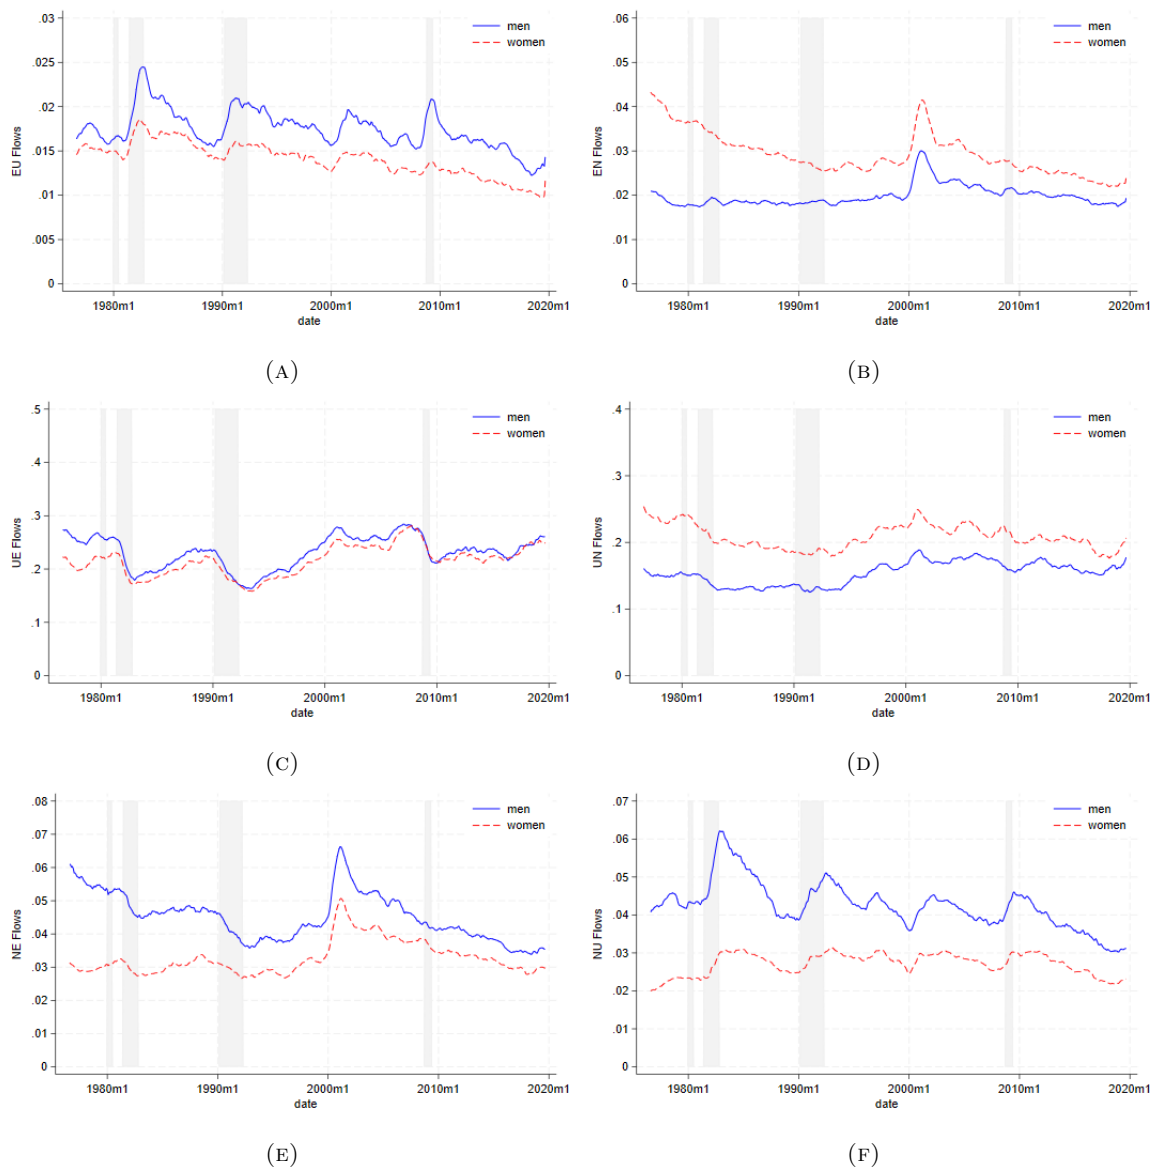
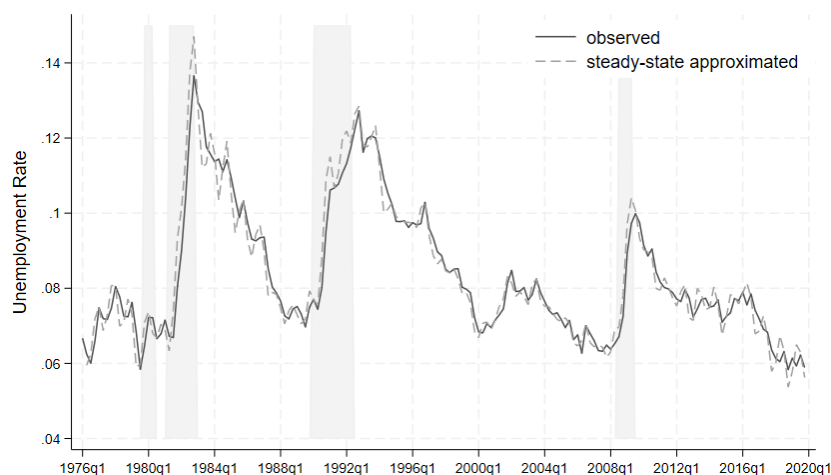


FIGURE 3.6: Monthly Labour Market Transition Rates (12-month centered moving average) by Gender, 1976m1-2019m12. Panel (a) is the transition rate from employment to unemployment, panel (b) is the transition rate from employment to non-participation, panel (c) is the transition rate from unemployment to employment, panel (d) is the transition rate from unemployment to non-participation, panel (e) is the transition rate from non-participation to employment and panel (f) is the transition rate from non-participation to unemployment. The rates are weighted using the LFS sampling weight. Shaded areas indicate dates of recent recessions. Data Source: the LFS master files.



(A) men



(B) women

FIGURE 3.7: Actual and Steady-State Unemployment Rate by Gender, 1976q1-2019q4. Panel (a) is the unemployment rate for men, while panel (b) is the unemployment rate for women. The solid line denotes the actual unemployment rate, and the dashed line denotes the steady-state unemployment rate, which is calculated by the six transition rates between employment, unemployment and non-participation. Shaded areas indicate dates of recent recessions. Data Source: the LFS master files.




FIGURE 3.8: Actual and Approximated Gender Unemployment Rate Gap. The solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the approximated first-order log-linearized unemployment gap. Shaded areas indicate dates of recent recessions. Data Source: the LFS master files.




FIGURE 3.9: Gender Unemployment Gap Decomposed by Six Labour Market Flows. Each coloured area represents the contribution of a labour market flow to the gender unemployment rate gap (female minus male unemployment rate). Contributions are calculated as the product of the marginal effect of each flow on the unemployment rate and the gender difference in that flow (see Appendix C.1). Positive values indicate that a flow increases the gap; negative values indicate a narrowing effect.




FIGURE 3.10: Actual and Counterfactual Gender Unemployment Rate Gap. The solid line represents the observed unemployment rate gap between women and men, and the dashed line represents the counterfactual gender unemployment rate gap, assuming women's λ^{EN} and λ^{NE} remain at the pre-1980s levels. Data Source: the LFS master files.

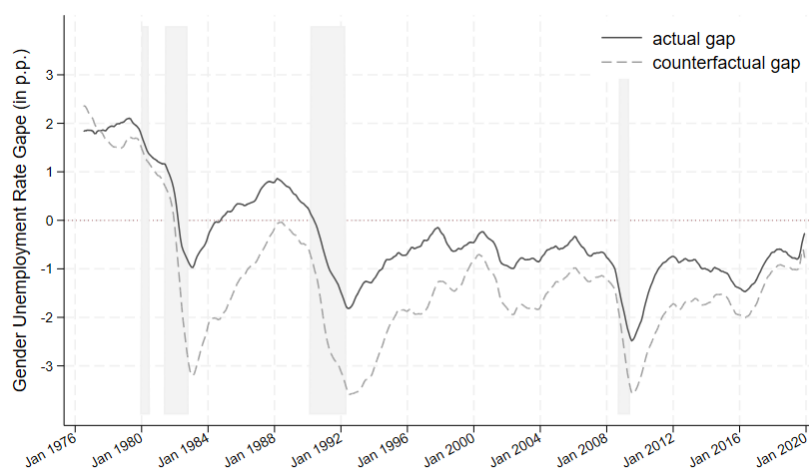


FIGURE 3.11: Actual and Counterfactual Gender Unemployment Rate Gap. The solid line represents the observed unemployment rate gap between women and men, and the dashed line represents the counterfactual gender unemployment rate gap, assuming women's λ^{UN} and λ^{NU} remain at the pre-1980s levels. Data Source: the LFS master files.




FIGURE 3.12: Actual and Counterfactual Gender Unemployment Rate Gap. The solid line represents the observed unemployment rate gap between women and men, and the dashed line represents the counterfactual gender unemployment rate gap, assuming women's λ^{EU} remain at the pre-1990s levels. Data Source: the LFS master files.

TABLE 3.1: Gender Difference in Employment-to-Unemployment Transition Rate

	(1)	(2)	(3)	(4)
Female	-0.362*** (0.0494)	-0.358*** (0.0415)	0.0620* (0.0199)	-0.110*** (0.0166)
Tenure Deviation				-0.00628*** (0.000633)
age group fixed effect	No	Yes	Yes	Yes
education fixed effect	No	Yes	Yes	Yes
industry fixed effect	No	No	Yes	Yes
occupation fixed effect	No	No	Yes	Yes
part-time status	No	No	No	Yes
N	19,776,939	19,776,939	19,776,939	19,776,939
R^2	0.001	0.005	0.008	0.011

Notes: The sample covers years from 1987 to 2019. The dependent variable is a binary variable equal to one if a worker becomes unemployed in the next month (rescaled to 100). ‘Female’ is a gender dummy variable. ‘Tenure Deviation’ denotes how many months a worker’s tenure is greater (or less) than the industry-occupation average tenure. All regressions control for province and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors, clustered at the provincial level, are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Bibliography

- Abowd, J. M. and Zellner, A.: 1985, Estimating gross labor-force flows, *Journal of Business and Economic Statistics* **3**(3), 254–283.
- Albagli, E., Contreras, G., Tapia, M. and Wlasiuk, J. M.: 2022, Earnings cyclicality of new and continuing jobs: The role of tenure and transition length, *Labour Economics* **78**, 102242. <https://doi.org/10.1016/j.labeco.2022.102242>.
- Albanesi, S. and Şahin, A.: 2018, The gender unemployment gap, *Review of Economic Dynamics* **30**, 47–67.
- Autor, D., Cho, D., Crane, L. D., Goldar, M., Lutz, B., Montes, J., Peterman, W. B., Ratner, D., Villar, D. and Yildirmaz, A.: 2022, An evaluation of the Paycheck Protection Program using administrative payroll microdata, *Journal of Public Economics* **211**, 104664.
- Azmat, G., Güell, M. and Manning, A.: 2006, Gender gaps in unemployment rates in OECD countries, *Journal of Labor Economics* **24**(1), 1–37.
- Baker, M. and Benjamin, D.: 1999, Early retirement provisions and the labor force behavior of older men: Evidence from Canada, *Journal of Labor Economics* **17**(4), 724–756.
- Baker, M., Corak, M. and Heisz, A.: 1998, The labour market dynamics of unemployment rates in Canada and the United States, *Canadian Public Policy* **24**(1), S72–S89.
- Bédard, N. and Michaud, P.-C.: 2021, Playing with fire? Household debt near retirement in Canada, *Canadian Public Policy* **47**(1), 56–71.
- Béland, L.-P., Brodeur, A., Mikola, D. and Wright, T.: 2022, The short-term economic consequences of COVID-19: Occupation tasks and mental health in Canada, *Canadian Journal of Economics* **55**(S1), 214–247.
- Blanchard, O. J. and Diamond, P.: 1990, The cyclical behavior of the gross flows of US workers, *Brookings Papers on Economic Activity* **1990**(2), 85–155.
- Blau, D. M.: 1998, Labor force dynamics of older married couples, *Journal of Labor Economics* **16**(3), 595–629.
- Blau, F. D. and Kahn, L. M.: 1994, Rising wage inequality and the US gender gap, *American Economic Review* **84**(2), 23–28.
- Blau, F. D. and Kahn, L. M.: 2000, Gender differences in pay, *Journal of Economic Perspectives* **14**(4), 75–100.

- Blau, F. D. and Kahn, L. M.: 2003, Understanding international differences in the gender pay gap, *Journal of Labor Economics* **21**(1), 106–144.
- Blau, F. D. and Kahn, L. M.: 2006, The US gender pay gap in the 1990s: Slowing convergence, *ILR Review* **60**(1), 45–66.
- Blau, F. D. and Kahn, L. M.: 2017, The gender wage gap: Extent, trends, and explanations, *Journal of Economic Literature* **55**(3), 789–865.
- Blau, F. D., Kahn, L. M., Boboshko, N. and Comey, M.: 2024, The impact of selection into the labor force on the gender wage gap, *Journal of Labor Economics* **42**(4), 1093–1133.
- Blit, J., Skuterud, M. and Veall, M. R.: 2020, The pandemic and short-run changes in output, hours worked and labour productivity: Canadian evidence by industry, *International Productivity Monitor* **39**(39), 16–32.
- Bocci, C. and Beaumont, J.-F.: 2004, Longitudinal hot-deck imputation for household non-response in the LFS. Household Survey Methods Division, Statistics Canada.
- Bredemeier, C. and Winkler, R.: 2017, The employment dynamics of different population groups over the business cycle, *Applied Economics* **49**(26), 2545–2562.
- Brochu, P.: 2021, A researcher’s guide to the Labour Force Survey: Its evolution and the choice of public use versus master files, *Canadian Public Policy* **47**(3), 335–357.
- Brochu, P., Créchet, J. and Deng, Z.: 2020, Labour market flows and worker trajectories in Canada during COVID-19. CLEF Working Paper Series, No. 32.
- Brochu, P. and Créchet, J.: 2022, Survey non-response in COVID-19 times: The case of the Labour Force Survey, *Canadian Public Policy* **48**(3), 451–472.
- Brodeur, A., Gray, D., Islam, A. and Bhuiyan, S.: 2021, A literature review of the economics of COVID-19, *Journal of Economic Surveys* **35**(4), 1007–1044.
- Campolieti, M.: 2011, The ins and outs of unemployment in Canada, 1976–2008, *Canadian Journal of Economics* **44**(4), 1331–1349.
- Cattapan, A., Acker-Verney, J. M., Dobrowolsky, A., Findlay, T. and Mandrona, A.: 2020, Community engagement in a time of confinement, *Canadian Public Policy* **46**, S287–S299.
- Clark, K. B. and Summers, L. H.: 1981, Demographic differences in cyclical employment variation, *Journal of Human Resources* **16**(1), 61–79.
- Cross, P. and Bergevin, P.: 2012, Turning points: Business cycles in Canada since 1926, *Commentary No. 366 - C.D. Howe Institute* pp. 1–24.
- Davis, O.: 2021, Employment and retirement among older workers during the COVID-19 pandemic, *Research note*, Schwartz Center for Economic Policy Analysis (SCEPA), Department of Economics, The New School for Social Research.

- Davis, S. J., Faberman, R. J. and Haltiwanger, J.: 2006, The flow approach to labor markets: new data sources and micro–macro links, *Journal of Economic Perspectives* **20**(3), 3–26.
- Elsby, M. W. L., Hobijn, B., Karahan, F., Kořar, G. and řahin, A.: 2019, Flow origins of labor force participation fluctuations, *AEA Papers and Proceedings* **109**, 461–464.
- Elsby, M. W. L., Hobijn, B. and řahin, A.: 2013, Unemployment dynamics in the OECD, *Review of Economics and Statistics* **95**(2), 530–548.
- Elsby, M. W. L., Hobijn, B. and řahin, A.: 2015, On the importance of the participation margin for labor market fluctuations, *Journal of Monetary Economics* **72**, 64–82.
- Elsby, M. W. L., Michaels, R. and Solon, G.: 2009, The ins and outs of cyclical unemployment, *American Economic Journal: Macroeconomics* **1**(1), 84–110.
- Finnie, R. and Gray, D.: 2018, How do older laid-off workers get by: Reemployment, early retirement, or social insurance benefits?, *Canadian Public Policy* **44**(2), 173–189.
- Forsythe, E.: 2022, Why don't firms hire young workers during recessions?, *Economic Journal* **132**(645), 1765–1789.
- Forsythe, E. and Wu, J.-C.: 2021, Explaining demographic heterogeneity in cyclical unemployment, *Labour Economics* **69**, 101955.
- Fortin, N. M.: 2019, Increasing earnings inequality and the gender pay gap in Canada: Prospects for convergence, *Canadian Journal of Economics* **52**(2), 407–440.
- Fortin, N. M., Bell, B. and Böhmer, M.: 2017, Top earnings inequality and the gender pay gap: Canada, Sweden, and the United Kingdom, *Labour Economics* **47**, 107–123.
- French, E.: 2005, The effects of health, wealth, and wages on labour supply and retirement behaviour, *Review of Economic Studies* **72**(2), 395–427.
- Fujita, S. and Ramey, G.: 2006, The cyclical nature of job loss and hiring. Federal Reserve Bank of Philadelphia.
- Fujita, S. and Ramey, G.: 2009, The cyclical nature of separation and job finding rates, *International Economic Review* **50**(2), 415–430.
- Gerbery, D. and Miklořovič, T.: 2020, Labour market transitions and their determinants in Slovakia: Path from crisis to recovery, *Ekonomický řasopis* **68**(7), 651–676.
- Goldin, C.: 1990, *Understanding the gender gap: An economic history of American women*, National Bureau of Economic Research.
- Goldin, C.: 2006, The quiet revolution that transformed women's employment, education, and family, *The American Economic Review* **96**(2), 1–21.
- Goldin, C.: 2014, A grand gender convergence: Its last chapter, *American Economic Review* **104**(4), 1091–1119.

- Goldin, C.: 2024, Nobel lecture: An evolving economic force, *American Economic Review* **114**(6), 1515–1539.
- Goldin, C., Kerr, S. P., Olivetti, C. and Barth, E.: 2017, The expanding gender earnings gap: Evidence from the LEHD-2000 census, *The American Economic Review* **107**(5), 110–114.
- Gower, D.: 1997, Measuring the age of retirement, *Perspectives on Labour and Income* **9**(2), 11–17.
- Gustman, A. L. and Steinmeier, T. L.: 1986, A structural retirement model, *Econometrica* **54**(3), 555–584.
- Gustman, A. L., Steinmeier, T. L. and Tabatabai, N.: 2010, What the stock market decline means for the financial security and retirement choices of the near-retirement population, *Journal of Economic Perspectives* **24**(1), 161–182.
- Hairault, J.-O., Le Barbanchon, T. and Sopraseuth, T.: 2015, The cyclicity of the separation and job finding rates in France, *European Economic Review* **76**, 60–84.
- Hamilton, J. D.: 2018, Why you should never use the Hodrick-Prescott filter, *Review of Economics and Statistics* **100**(5), 831–843.
- Hoynes, H., Miller, D. L. and Schaller, J.: 2012, Who suffers during recessions?, *Journal of Economic Perspectives* **26**(3), 27–48.
- Jaimovich, N. and Siu, H. E.: 2009, The young, the old, and the restless: Demographics and business cycle volatility, *American Economic Review* **99**(3), 804–826.
- Jones, S. R. G., Lange, F., Riddell, W. C. and Warman, C.: 2021, Canadian labour market dynamics during COVID-19. IZA Discussion Paper Series, No. 14588.
- Jones, S. R. G. and Riddell, W. C.: 1998, Gross flows of labour in Canada and the United States, *Canadian Public Policy* pp. S103–S120.
- Jones, S. R. G. and Riddell, W. C.: 1999, The measurement of unemployment: An empirical approach, *Econometrica* **67**(1), 147–162.
- Jones, S. R. G. and Riddell, W. C.: 2006, Unemployment and nonemployment: Heterogeneities in labor market states, *Review of Economics and Statistics* **88**(2), 314–323.
- Jones, S. R. G. and Riddell, W. C.: 2019, Unemployment, marginal attachment, and labor force participation in Canada and the United States, *Journal of Labor Economics* **37**(S2), S399–S441.
- Jung, P. and Kuhn, M.: 2019, Earnings losses and labor mobility over the life cycle, *Journal of the European Economic Association* **17**(3), 678–724.
- Koutentakis, F.: 2015, Gender unemployment dynamics: Evidence from ten advanced economies, *Labour* **29**(1), 15–31.

- Krolikowski, P.: 2017, Job ladders and earnings of displaced workers, *American Economic Journal: Macroeconomics* **9**(2), 1–31.
- Lauerova, J. S. and Terrell, K.: 2007, What drives gender differences in unemployment?, *Comparative Economic Studies* **49**, 128–155.
- Lemieux, T., Milligan, K., Schirle, T. and Skuterud, M.: 2020, Initial impacts of the COVID-19 pandemic on the Canadian labour market, *Canadian Public Policy* **46**(S1), S55–S65.
- Light, A. and McGarry, K.: 1998, Job change patterns and the wages of young men, *Review of Economics and Statistics* **80**(2), 276–286.
- Lydon, R. and Simmons, M.: 2020, The ins and outs of the gender unemployment gap in the OECD, *Working Paper: SSRN 4832546*.
- Madrian, B. C. and Lefgren, L. J.: 1999, A note on longitudinally matching Current Population Survey (CPS) respondents, *NBER Technical Working Paper No. T0247*.
- McFall, B. H.: 2011, Crash and wait? The impact of the Great Recession on the retirement plans of older Americans, *American Economic Review* **101**(3), 40–44.
- Milligan, K. S. and Schirle, T.: 2018a, The labor force participation of older men in Canada, *NBER Working Paper w24874*, National Bureau of Economic Research.
URL: <https://www.nber.org/papers/w24874>
- Milligan, K. and Schirle, T.: 2018b, Health and capacity to work of older Canadians: Gender and regional dimensions, *Canadian Public Policy* **44**(2), 159–172.
- Montes, J., Smith, C. and Dajon, J.: 2022, “The great retirement boom”: The pandemic-era surge in retirements and implications for future labor force participation, *Finance and Economics Discussion Series 2022-081*, Board of Governors of the Federal Reserve System, Washington.
URL: <https://doi.org/10.17016/FEDS.2022.081>
- Morissette, R., Deng, Z. and Messacar, D.: 2021, Working from home: Potential implications for public transit and greenhouse gas emission, *Economic and Social Reports*, Statistics Canada.
- Munnell, A. H., Sanzenbacher, G. T. and Rutledge, M. S.: 2018, What causes workers to retire before they plan?, *Journal of Retirement* **6**(2), 35–52.
- Myatt, A. and Murrell, D.: 1990, The female/male unemployment rate differential, *Canadian Journal of Economics* **23**(2), 312–322.
- Nakamura, A., Nakamura, E., Phong, K. and Steinsson, J.: 2019, New evidence on the cyclicity of employer-to-employer flows from Canada, *AEA Papers and Proceedings* **109**, 456–460.
- Noda, S. and Teramoto, K.: 2024, A dynamic model of rational “panic buying”, *Quantitative Economics* **15**(2), 489–521.
- Olivetti, C., Pan, J. and Petrongolo, B.: 2024, The evolution of gender in the labor market, *Handbook of Labor Economics*, Vol. 5, Elsevier, pp. 619–677.

- Pue, K., Westlake, D. and Jansen, A.: 2021, Does the profit motive matter? COVID-19 prevention and management in Ontario long-term-care homes, *Canadian Public Policy* **47**(3), 421–438.
- Quinby, L. D., Rutledge, M. S. and Wettstein, G.: 2021, How has COVID-19 affected the labor force participation of older workers?, *Working Paper 2021-13*, Center for Retirement Research at Boston College.
- Schmieder, J. F., von Wachter, T. and Heining, J.: 2023, The costs of job displacement over the business cycle and its sources: Evidence from Germany, *American Economic Review* **113**(5), 1208–1254.
- Shimer, R.: 2005, The cyclicalities of hires, separations, and job-to-job transitions, *Review - Federal Reserve Bank of Saint Louis* **87**(4), 493.
- Shimer, R.: 2012, Reassessing the ins and outs of unemployment, *Review of Economic Dynamics* **15**(2), 127–148.
- Slade, M. E.: 2022, Many losers and a few winners: The impact of COVID-19 on Canadian industries and regions, *Canadian Journal of Economics* **55**(S1), 282–307.
- Smeaton, D., Di Rosa, M., Principi, A. and Butler, Z.: 2018, Reverse retirement — a mixed methods study of returning to work in England, Italy and the United States: Propensities, predictors and preferences, *International Journal of Ageing and Later life* **12**(1), 5–40.
- Topel, R. H. and Ward, M. P.: 1992, Job mobility and the careers of young men, *Quarterly Journal of Economics* **107**(2), 439–479.
- Usalcas, J. and Kinack, M.: 2017, History of the Canadian Labour Force Survey, 1945 to 2016, Labour Statistics Technical Papers. Catalogue No. 75-005-M.
- Villeneuve, H., Abdeen, A., Papineau, M., Simon, S., Cruickshank, C. and O'Brien, W.: 2021, New insights on the energy impacts of telework in Canada, *Canadian Public Policy* **19**(3), 1–33.
- Williams, N.: 2009, Seniority, experience, and wages in the UK, *Labour Economics* **16**(3), 272–283.
- Xu, H. and Couch, K. A.: 2017, The business cycle, labor market transitions by age, and the Great Recession, *Applied Economics* **49**(52), 5370–5396.
- Yashiv, E.: 2007, US labor market dynamics revisited, *The Scandinavian Journal of Economics* **109**(4), 779–806.

Appendix A

A.1 Panel Creation and Sample Restrictions

The Labour Force Survey (LFS) is not a panel per se. Although one can use a unique household ID (HHLDID) (and a unique individual ID within a household (LINE)), the household ID variable amalgamates information from several variables, including PROV, PROV1, PSEUDOUI, FRAME, STRAFRAM, TYPE, CLUST, ROTATION, LISTLINE and MULT. This has implications for the creation of mini-panels, which I discuss as follows.

To track new job matches and separations across two consecutive months, it is necessary to have complete labour market information of a respondent for two months. However, in the LFS, not all individuals have full two-month employment information. This incompleteness arises mainly due to three factors: one-sixth of the sample rotating out monthly as part of the survey's design, individuals moving out/in partway through the six-month survey window, and instances of survey non-response. For those who never respond, they are absent in both months of the two-month data panel. The LFS modifies its weightings to account for such instances. In cases where an individual responds in one month but not the next, the LFS employs a hot-deck imputation technique to impute their labour force information (Bocci and Beaumont, 2004). I choose not to omit individuals subject to Whole Record Imputation from my analysis. Based on empirical evidence, including or excluding these participants does not significantly alter the core results of the study. As a result of these various factors, typically, up to five-sixths of the original sample is available for creating balanced, two-month mini-panels.

Following the guidelines provided by Statistics Canada, I adapt the household identification variable HHLDID from January 1984 to December 2005. This involves standardizing 17-digit HHLDIDs to an 18-digit format by adding a zero at the end. The final digit of HHLDID, MULT, represents multi-dwelling structures and is left blank for single detached homes during this period. However, after this period, single detached homes were assigned a value of '0'. A notable point is that an individual could have different HHLDID values over two months, such as a 17-digit ID in December 2005 and an 18-digit ID in January 2006. This variation was particularly challenging during the transition phase when '0' values for MULT started being used. Once the HHLDID variable was standardized, this discrepancy was resolved. It is also important to highlight that, although an individual's ID may recur across various years in the LFS, it predominantly remains distinct for each specific month and year. This uniqueness is consistent, with only a few exceptions noted in 1977. In the two-month mini-panels constructed for this study, there are no instances of duplicate individual IDs within the same month and

year. Additionally, these IDs are not recycled within a six-month time frame, which guarantees that there are no complications in the creation of two-month mini-panels.

Lastly, in line with Madrian and Lefgren (1999), I remove individuals from the dataset who show inconsistencies in age and gender across two successive months. The LFS collects socio-demographic information, including gender, during the initial or birth interview, and this information is not updated in later interviews. Therefore, changes in gender are not expected. While the survey does update respondent ages according to their birthdays, discrepancies in age from one month to the next that are not consistent with a one-year increase could indicate a potential coding error. For instance, if a respondent's age differs from the previous month, and this difference does not align with an annual increment, it suggests an issue with the data recording process. To maintain data accuracy, I exclude such cases from the constructed two-month panels. This practice is standard in creating panel data from monthly surveys. This exclusion of mismatches in gender and age results in approximately 1% of the sample being omitted from the study.

A.2 Time Path and Steady States of Participation Rate

Expressing the participation rate at time t with its stock at time $t - 1$ and inflows and outflows between the two periods, one can obtain the law of motion of the participation rate as

$$pr_t = pr_{t-1} + (1 - pr_{t-1}) \cdot \lambda_t^{OI} - pr_{t-1} \cdot \lambda_t^{IO}.$$

Then, by solving the first-order difference equation, the time path of the participation rate is:¹

$$pr_t = (pr_0 - \frac{\lambda_t^{OI}}{\lambda_t^{IO} + \lambda_t^{OI}}) \cdot [1 - (\lambda_t^{IO} + \lambda_t^{OI})]^t + \frac{\lambda_t^{OI}}{\lambda_t^{IO} + \lambda_t^{OI}}. \quad (\text{A.1})$$

The first item on the right-hand side of equation (A.1) determines how fast the participation rate converges to its steady state. As the activity level of the labour market at the participation margin is low most times, i.e. magnitudes of λ^{IO} and λ^{OI} are small, it takes longer for the participation rate to converge to its steady state. Therefore, this study does not use a steady-state decomposition of the participation rate.

In contrast, unemployment is a relatively transitory state; the probabilities of transiting from unemployment to either employment or non-participation are relatively high. Thus, the unemployment rate can converge to its steady state faster, so that its steady state can be used as an approximation of the unemployment rate. With the same logic, steady-state approximation can be applied to the marginal-attachment-to-population rate.

¹As $[1 - (\lambda_t^{IO} + \lambda_t^{OI})] < 1$, the steady state participation rate at time t exists and can be estimated as:

$$\hat{pr}_t^* = \frac{\hat{\lambda}_t^{OI}}{\hat{\lambda}_t^{IO} + \hat{\lambda}_t^{OI}}.$$

Appendix B

B.1 Appendix Figures and Tables

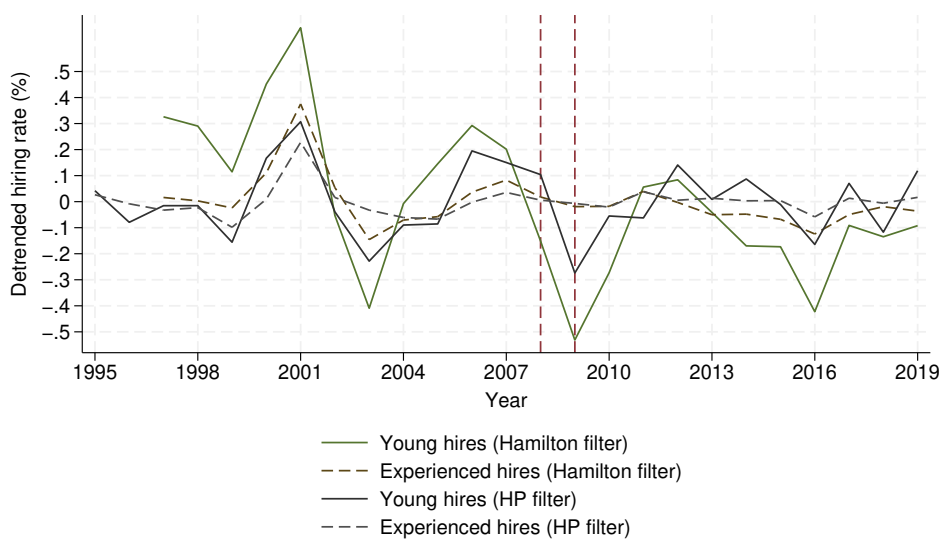


FIGURE B.1: Detrended Hiring Rate Over the Business Cycle, Comparison between HP filter and Hamilton filter. Each series is seasonally adjusted by removing month-fixed effects, detrended by the HP/Hamilton filter, and then averaged across the years. Red dashed vertical lines indicate recession dates in Canada.

TABLE B.1: Joint Distribution of Experience and Tenure Groups: Employed Individuals

	Tenure ≥ 3	Tenure < 3	Total
PE > 10	51.9% (9,767,260)	19.3% (3,580,274)	71.2% (13,347,534)
PE ≤ 10	8.6% (1,428,610)	20.1% (3,414,603)	28.8% (4,843,213)
Total	60.5% (11,195,870)	39.5% (6,994,877)	100.0% (18,190,747)

Notes: Each cell reports the percentage of the total weighted sample of currently employed individuals, aged 15-69, from the 1995-2019 LFS public-use files. Unweighted sample counts are reported in parentheses below each cell. Potential experience (PE) is defined as age minus years of schooling minus 6. Because the public-use file includes only age groups, individual age is imputed using the median of each age group. Tenure refers to years with the current employer. The tabulation includes all employed individuals, not just those who experienced job transitions. Weights correspond to LFS sampling weights.

TABLE B.2: Joint Distribution of Experience and Tenure Groups: Non-Employed but Worked in the Past Twelve Months

	Tenure ≥ 3	Tenure < 3	Total
PE > 10	23.1% (594,146)	26.7% (776,059)	49.9% (1,370,205)
PE ≤ 10	4.9% (113,983)	45.2% (1,123,389)	50.1% (1,237,372)
Total	28.0% (708,129)	71.9% (1,899,448)	100.0% (2,607,577)

Notes: Each cell reports the percentage of the total weighted sample of non-employed individuals aged 15-69 who had worked in the past 12 months, based on the 1995-2019 LFS public-use files. Unweighted sample counts are reported in parentheses below each cell. Potential experience (PE) is defined as age minus years of schooling minus 6. Since only age groups are available in the public-use file, age is imputed using the group median. Tenure refers to tenure with the previous employer. Weights correspond to LFS sampling weights.

TABLE B.3: Hiring over the Business Cycle: With and Without Controls (Provincial)

Outcome: Pr(Hired) \times 100	(1)	(2)	(3)	(4)
<i>Panel A: Aggregate effect</i>				
U. rate	-0.00849 (0.0406)	-0.0964*** (0.0200)	-0.119*** (0.0200)	-0.0578* (0.0204)
R^2	0.000	0.000	0.002	0.004
<i>Panel B: Disaggregate by potential experience</i>				
PE \leq 10	4.764*** (0.223)	4.731*** (0.236)	4.551*** (0.224)	4.561*** (0.225)
PE \leq 10 \times U. rate	-0.0975 (0.0516)	-0.195*** (0.0347)	-0.206*** (0.0345)	-0.149*** (0.0249)
PE $>$ 10 \times U. rate	0.0383 (0.0347)	-0.0653** (0.0175)	-0.0858** (0.0186)	-0.0272 (0.0188)
R^2	0.009	0.010	0.011	0.012
Prov. fixed effect	No	Yes	Yes	Yes
Demographic fixed effect	No	No	Yes	Yes
Month-year fixed effect	No	No	No	Yes
N	21,253,704	21,253,704	21,253,704	21,253,704

Notes: The dependent variable is a binary variable which refers to a worker beginning a job at a new firm. ‘U. rate’ refers to the provincial unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE B.4: Hiring over the Business Cycle: Young and Experienced (Provincial)

	(1)	(2)	(3)	(4)
Outcome: Pr(Hired) \times 100	All	EE	UE	NE
PE \leq 10	4.561*** (0.225)	2.373*** (0.205)	11.43*** (1.181)	9.390*** (0.896)
PE \leq 10 \times U. rate	-0.149*** (0.0249)	-0.123*** (0.0219)	-2.450*** (0.271)	-0.438*** (0.0729)
PE $>$ 10 \times U. rate	-0.0272 (0.0188)	-0.0432*** (0.00506)	-1.782*** (0.181)	-0.0734 (0.0489)
Sample	All	Employed	Unemployed	NILF
R^2	0.0122	0.00729	0.0289	0.0267
Wald test (p-value)	0.000230	0.00877	0.000157	0.00209
N	21,253,704	14,499,841	1,165,829	5,588,034

Notes: The dependent variable is a binary variable which refers to beginning a job at a new firm. ‘All’ refers to all new hires. ‘EE’ refers to new hires from another job. ‘UE’ refers to new hires from unemployment. ‘NE’ refers to new hires from non-participation. ‘U. rate’ refers to the provincial unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE B.5: Exits and Other Flows (Provincial)

	(1)	(2)	(3)	(4)	(5)	(6)
	E exits	EU	EE	EN	NU	UN
PE \leq 10	4.898*** (0.286)	0.874*** (0.106)	2.385*** (0.161)	1.639*** (0.234)	6.629*** (0.895)	4.261*** (0.454)
PE \leq 10 \times U. rate	0.0893 (0.0456)	0.136*** (0.0223)	-0.116*** (0.0233)	0.0690 (0.0340)	0.281** (0.0805)	0.134 (0.111)
PE $>$ 10 \times U. rate	0.102* (0.0355)	0.121*** (0.0122)	-0.0514*** (0.00959)	0.0322 (0.0218)	0.406*** (0.0346)	0.104 (0.149)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	0.754	0.302	0.0170	0.259	0.250	0.596
R^2	0.0188	0.00540	0.00744	0.0107	0.0175	0.0227
N	14,499,841	14,499,841	14,499,841	14,499,841	5,588,034	1,165,829

Notes: 'E exists' is a binary variable which refers to a job separation. 'EU' refers to a job separation to unemployment. 'EE' refers to a job separation ending in a new job. 'EN' refers to a job separation to non-participation. 'NU' and 'UN' refer to transitions between non-participation and unemployment. 'U. rate' refers to the provincial unemployment rate. 'PE' refers to potential experience, defined as (age - education - 6). Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE B.6: Involuntary and Voluntary Separations to Unemployment (Provincial)

	(1) Pr(Involuntary) \times 100	(2) Pr(Voluntary) \times 100
PE \leq 10	0.162 (0.137)	0.504*** (0.0353)
PE \leq 10 \times U. rate	0.152*** (0.0216)	-0.00531 (0.00451)
PE $>$ 10 \times U. rate	0.109*** (0.00798)	0.00667 (0.00360)
Wald test (p-values)	0.0621	0.0174
R^2	0.00360	0.00174
N	14,499,841	14,499,841

Notes: ‘Involuntary’ is a binary variable which refers to involuntarily separate from a job to unemployment. ‘Voluntary’ refers to voluntarily separating from a job to unemployment. ‘U. rate’ refers to the provincial unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). Estimates include main effects and the provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct.

TABLE B.7: Hiring over the Business Cycle: Age Versus Education (Provincial)

Outcome: Pr(Hired) \times 100	(1)	(2)	(3)	(4)
College	-0.744* (0.291)	-0.858** (0.258)	-1.670*** (0.343)	-1.560** (0.357)
No College \times U. rate	-0.0127 (0.0452)	-0.105*** (0.0217)	-0.116*** (0.0216)	-0.0536* (0.0209)
College \times U. rate	-0.0476** (0.0141)	-0.121** (0.0286)	-0.129** (0.0302)	-0.0778* (0.0342)
Province fixed effect	No	Yes	Yes	Yes
Demographic fixed effect	No	No	Yes	Yes
Month-year fixed effect	No	No	No	Yes
Wald test (p-value)	0.340	0.593	0.700	0.468
R^2	0.000538	0.000828	0.00224	0.00362
N	21,253,704	21,253,704	21,253,704	21,253,704

Notes: ‘Hired’ is a binary variable which refers to beginning a job at a new firm. ‘U. rate’ refers to the provincial unemployment rate. ‘College’ refers to having a university degree. Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the Col. \times U. rate and No col. \times U. rate coefficients are statistically distinct.

TABLE B.8: Job-to-Job Transition over the Business Cycle: Experience and Tenure (Provincial)

Outcome: Pr(EE) \times 100	(1)	(2)	(3)
PE \leq 10	2.373*** (0.205)		1.782*** (0.156)
Tenure $<$ 3		1.982*** (0.192)	1.312*** (0.155)
PE \leq 10 \times U. rate	-0.123*** (0.0219)		-0.104*** (0.0144)
PE $>$ 10 \times U. rate	-0.0432*** (0.00506)		-0.0331** (0.00763)
Tenure $<$ 3 \times U. rate		-0.0813*** (0.0163)	-0.0122 (0.0150)
Tenure \geq 3 \times U. rate		-0.0394*** (0.00527)	
R^2	0.00729	0.00755	0.00984
Wald test (p-value)	0.00877	0.0622	0.00504
N	14,499,841	14,499,841	14,499,841

Notes: The dependent variable is a binary variable which denotes the happening of a job-to-job transition, re-scaled to 100. ‘U. rate’ refers to the provincial unemployment rate. ‘PE’ refers to potential experience, defined as (age - education - 6). ‘Tenure’ refers to job tenure (in years). Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald tests for columns (1) (3) indicate whether the PE \leq 10 \times U. rate and PE $>$ 10 \times U. rate coefficients are statistically distinct. The Wald test for column (2) indicates whether the Tenure $<$ 3 \times U. rate and Tenure \geq 3 \times U. rate coefficients are statistically distinct.

TABLE B.9: Exits and Other Flows: Tenure Groups (Provincial)

	(1) E exits	(2) EU	(3) EE	(4) EN	(5) NU	(6) UN
Tenure < 3	2.509*** (0.302)	0.0608 (0.109)	2.021*** (0.147)	0.427 (0.225)	4.922*** (0.728)	3.855*** (0.339)
Tenure < 3 × U. rate	0.397*** (0.0464)	0.298*** (0.0129)	-0.0755** (0.0180)	0.174*** (0.0334)	0.613*** (0.119)	0.0709 (0.128)
Tenure ≥ 3 × U. rate	-0.0474 (0.0413)	0.0320 (0.0191)	-0.0503*** (0.0101)	-0.0291 (0.0244)	0.756*** (0.0656)	0.275 (0.143)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	0.00000477	4.92e-08	0.150	0.000195	0.114	0.000114
R^2	0.0264	0.0110	0.00829	0.0118	0.0184	0.0158
N	14,499,841	14,499,841	14,499,841	14,499,841	1,208,085	820,675

Notes: ‘E exists’ is a binary variable which refers to a job separation. ‘EU’ refers to a job separation to unemployment. ‘EE’ refers to a job separation ending in a new job. ‘EN’ refers to a job separation to non-participation. ‘NU’ and ‘UN’ refer to transitions between non-participation and unemployment. ‘U. rate’ refers to the provincial unemployment rate. ‘Tenure’ refers to the tenure of the previous job. Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the Tenure < 3 × U. rate and Tenure ≥ 3 × U. rate coefficients are statistically distinct.

TABLE B.10: Exits and Other Flows: Experience & Tenure Groups (Provincial)

	(1) E exits	(2) EU	(3) EE	(4) EN	(5) NU	(6) UN
PE \leq 10	4.374*** (0.287)	0.932*** (0.125)	1.756*** (0.131)	1.686*** (0.216)	2.095 (1.189)	5.573*** (0.566)
Tenure < 3	0.905* (0.309)	-0.258 (0.131)	1.366*** (0.129)	-0.203 (0.211)	4.354** (1.164)	1.685*** (0.233)
PE \leq 10 \times U. rate	-0.255*** (0.0474)	-0.0648 (0.0314)	-0.108*** (0.0189)	-0.0828* (0.0313)	0.222 (0.105)	0.162 (0.115)
PE > 10 \times U. rate	-0.0273 (0.0413)	0.0444* (0.0177)	-0.0462** (0.0104)	-0.0255 (0.0243)	0.889*** (0.0843)	0.247 (0.137)
Tenure < 3 \times U. rate	0.533*** (0.0416)	0.303*** (0.0181)	0.00189 (0.0124)	0.228*** (0.0311)	0.0998 (0.149)	-0.114** (0.0312)
Sample	Employed	Employed	Employed	Employed	NILF	Unemployed
Wald test (p-values)	0.0000438	0.0000712	0.0101	0.0489	0.00168	0.192
R^2	0.0288	0.0111	0.0104	0.0129	0.0213	0.0199
N	14,499,841	14,499,841	14,499,841	14,499,841	1,208,085	820,675

Notes: 'E exists' is a binary variable which refers to a job separation. 'EU' refers to a job separation to unemployment. 'EE' refers to a job separation ending in a new job. 'EN' refers to a job separation to non-participation. 'NU' and 'UN' refer to transitions between non-participation and unemployment. 'U. rate' refers to the provincial unemployment rate. 'PE' refers to potential experience, defined as (age - education - 6). 'Tenure' refers to the tenure of the previous job. Estimates include main effects and provincial, demographic, and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors are clustered at the provincial level and reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The Wald test indicates whether the PE \leq 10 \times U. rate and PE > 10 \times U. rate coefficients are statistically distinct.

Appendix C

C.1 Marginal Contribution of Each Flow on Steady-State Unemployment Rate

Equation (3.7) shows that the gender unemployment rate gap can be decomposed into the contributions of the six transition rates, and equation (3.6) shows that the contribution of each transition rate depends on its marginal impact on the steady-state unemployment rate and the gender difference in this transition rate. To have a better understanding of how each transition rate contributes to the gender unemployment gap, here I show the marginal impact of each flow on the steady-state unemployment rate.

Denote $a = \lambda_t^{NE} \lambda_t^{EU} + \lambda_t^{EN} \lambda_t^{NU} + \lambda_t^{EU} \lambda_t^{NU}$ and $b = \lambda_t^{UE} \lambda_t^{NE} + \lambda_t^{UN} \lambda_t^{NE} + \lambda_t^{NU} \lambda_t^{UE}$. According to equation (3.4), one can write the steady-state unemployment rate as $u_t^{ss} = \frac{a}{b+a}$. Note that the steady-state unemployment (u_t^{ss}) has a similar structure as the unemployment rate (u_t), where a denotes all possible inflows to unemployment (N to E to U, E to N to U, E to U and N to U), while b denotes all possible inflows to employment (U to N to E, N to U to E, U to E and N to E).

The marginal impact of transition rate from employment to unemployment (λ^{EU}) at time t is:

$$\begin{aligned} \frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{EU}} &= \frac{a+b}{a} \cdot \frac{(\lambda_t^{NE} + \lambda_t^{NU})(a+b) - a \cdot (\lambda_t^{NE} + \lambda_t^{NU})}{(a+b)^2} \\ &= \frac{b \cdot (\lambda_t^{NE} + \lambda_t^{NU})}{a \cdot (a+b)} \end{aligned}$$

As $a > 0$ and $b > 0$, the marginal impact of λ^{EU} on the steady-state unemployment rate is positive.

The contribution of λ_t^{EU} on gender unemployment rate gap $F_t^{EU} = \left. \frac{\partial \ln u_t^f}{\partial \lambda_t^{EU}} \right|_{\Lambda_t = \Lambda_t^m} \cdot (\lambda_t^{EU,f} - \lambda_t^{EU,m})$. Since women have lower employment-to-unemployment transition rates ($\lambda_t^{EU,f} - \lambda_t^{EU,m} < 0$), and the marginal impact of λ^{EU} on steady-state unemployment rate is positive, F_t^{EU} is negative. The employment-to-unemployment transition rate is a negative contributor to the gender unemployment rate gap.

Similarly, the marginal impact of the employment-to-non-participation rate (λ^{EN}) on the steady-state unemployment rate is positive:

$$\begin{aligned}\frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{EN}} &= \frac{a+b}{a} \cdot \frac{\lambda_t^{NU}(a+b) - \lambda_t^{NU}a}{(a+b)^2} \\ &= \frac{b \cdot \lambda_t^{NU}}{a \cdot (a+b)} > 0\end{aligned}$$

Since women have higher employment-to-non-participation rates than men ($\lambda_t^{EN,f} - \lambda_t^{EN,m} > 0$), $F_t^{EN} > 0$. λ^{EN} is a positive contributor to the gender unemployment rate gap.

The marginal impact of the unemployment-to-employment rate (λ^{UE}) on the steady-state unemployment rate is negative:

$$\begin{aligned}\frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{UE}} &= \frac{a+b}{a} \cdot \frac{-a \cdot (\lambda_t^{NU} + \lambda_t^{NE})}{(a+b)^2} \\ &= -\frac{\lambda_t^{NU} + \lambda_t^{NE}}{a+b} < 0\end{aligned}$$

As women have slightly lower unemployment-to-employment transition rates than men ($\lambda_t^{UE,f} - \lambda_t^{UE,m} < 0$), $F_t^{UE} > 0$. λ^{UE} is a positive contributor to the gender unemployment gap while its contribution is small ($\lambda_t^{UE,f} - \lambda_t^{UE,m}$ is close to zero).

The marginal impact of the unemployment-to-non-participation rate (λ^{UN}) on the steady-state unemployment rate is negative:

$$\begin{aligned}\frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{UN}} &= \frac{a+b}{a} \cdot \frac{-a \cdot \lambda_t^{NE}}{(a+b)^2} \\ &= -\frac{\lambda_t^{NE}}{a+b} < 0\end{aligned}$$

As women have higher unemployment-to-non-participation transition rates than men ($\lambda_t^{UN,f} - \lambda_t^{UN,m} > 0$), $F_t^{UN} < 0$. λ^{UN} is a negative contributor to the gender unemployment gap.

While the marginal effect of λ^{EU} and λ^{EN} on the steady-state unemployment rate are clearly positive, and the marginal effect of λ^{UE} and λ^{UN} are clearly negative; the marginal impact of λ^{NE} and λ^{NU} on the steady-state unemployment rate is not obvious.

The marginal impact of λ^{NE} on steady-state unemployment rate is

$$\begin{aligned}\frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{NE}} &= \frac{a+b}{a} \cdot \frac{\lambda_t^{EU}(a+b) - (\lambda_t^{EU} + \lambda_t^{UE} + \lambda_t^{UN})a}{(a+b)^2} \\ &= \frac{b \cdot \lambda_t^{EU} - a \cdot (\lambda_t^{UE} + \lambda_t^{UN})}{a \cdot (a+b)}\end{aligned}$$

The sign of this marginal impact depends on the sign of $b \cdot \lambda_t^{EU} - a \cdot (\lambda_t^{UE} + \lambda_t^{UN})$. However, empirically, the magnitudes of λ^{EU} , λ^{EN} , λ^{UE} , λ^{UN} , λ^{NE} and λ^{NU} are on average around .020, .023, .284, .231, .034 and .043, respectively, and the magnitudes of a and b are on average

around .0026 and .0301, respectively. The sign of $b \cdot \lambda_t^{EU} - a \cdot (\lambda_t^{UE} + \lambda_t^{UN})$ is empirically always negative, meaning λ^{NE} has a negative marginal impact on the unemployment rate.

This makes sense intuitively, as the transition from non-participation to employment would increase the stock of the labour force while the stock of unemployment remains unchanged. Thus, the unemployment rate would decrease when λ^{NE} increases.

As women have a lower non-participation-to-employment rate than men ($\lambda_t^{NE,f} - \lambda_t^{NE,m} < 0$), $F_t^{NE} > 0$. λ^{NE} is a positive contributor to the gender unemployment rate gap.

The marginal impact of λ^{NU} on steady-state unemployment rate is


$$\begin{aligned} \frac{\partial \ln u_t^{ss}}{\partial \lambda_t^{NU}} &= \frac{a+b}{a} \cdot \frac{(\lambda_t^{EU} + \lambda_t^{EN})(a+b) - a \cdot (\lambda_t^{EU} + \lambda_t^{EN} + \lambda_t^{UE})}{(a+b)^2} \\ &= \frac{b \cdot (\lambda_t^{EU} + \lambda_t^{EN}) - a \cdot \lambda_t^{UE}}{a \cdot (a+b)} \end{aligned}$$

The sign of this marginal effect depends on the sign of the term $b \cdot (\lambda_t^{EU} + \lambda_t^{EN}) - a \cdot \lambda_t^{UE}$. An increase in λ_t^{EU} raises the unemployment stock, while an increase in λ_t^{EN} reduces the stock of the labour force. Both effects exert upward pressure on the unemployment rate, represented by the term $b \cdot (\lambda_t^{EU} + \lambda_t^{EN})$. Conversely, an increase in λ_t^{UE} reduces the unemployment stock, exerting downward pressure on the unemployment rate, captured by the term $a \cdot \lambda_t^{UE}$.


Empirically, the positive contribution from $b \cdot (\lambda_t^{EU} + \lambda_t^{EN})$ outweighs the negative contribution from $a \cdot \lambda_t^{UE}$. Therefore, the marginal effect of λ^{NU} on the steady-state unemployment rate is positive.

Since women have lower non-participation-to-unemployment transition rates than men ($\lambda_t^{NU,f} - \lambda_t^{NU,m} < 0$), $F_t^{NU} < 0$. λ^{NU} is a negative contributor to the gender unemployment rate gap.

C.2 Appendix Figures and Tables




(A)




(B)

FIGURE C.1: Actual Gender Unemployment Rate Gap and the Proportion Explained by Individual Flow. Panel (a) is the gender unemployment gap explained by the employment-to-non-participation flow, and panel (b) is the gender unemployment gap explained by the non-participation-to-employment flow. The Solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the proportion of the gender unemployment gap that is explained by individual flows. Data Source: the LFS master files.




(A)




(B)

FIGURE C.2: Actual Gender Unemployment Rate Gap and the Proportion Explained by Individual Flow. Panel (a) is the gender unemployment gap explained by the unemployment-to-non-participation flow, and panel (b) is the gender unemployment gap explained by the non-participation-to-unemployment flow. The Solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the proportion of the gender unemployment gap that is explained by individual flows. Data Source: the LFS master files.



(A)



(B)

FIGURE C.3: Actual Gender Unemployment Rate Gap and the Proportion Explained by Individual Flow. Panel (a) is the gender unemployment gap explained by employment-to-unemployment flow, and panel (b) is the gender unemployment gap explained by unemployment-to-employment flow. The Solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the proportion of the gender unemployment gap that is explained by individual flows. Data Source: the LFS master files.

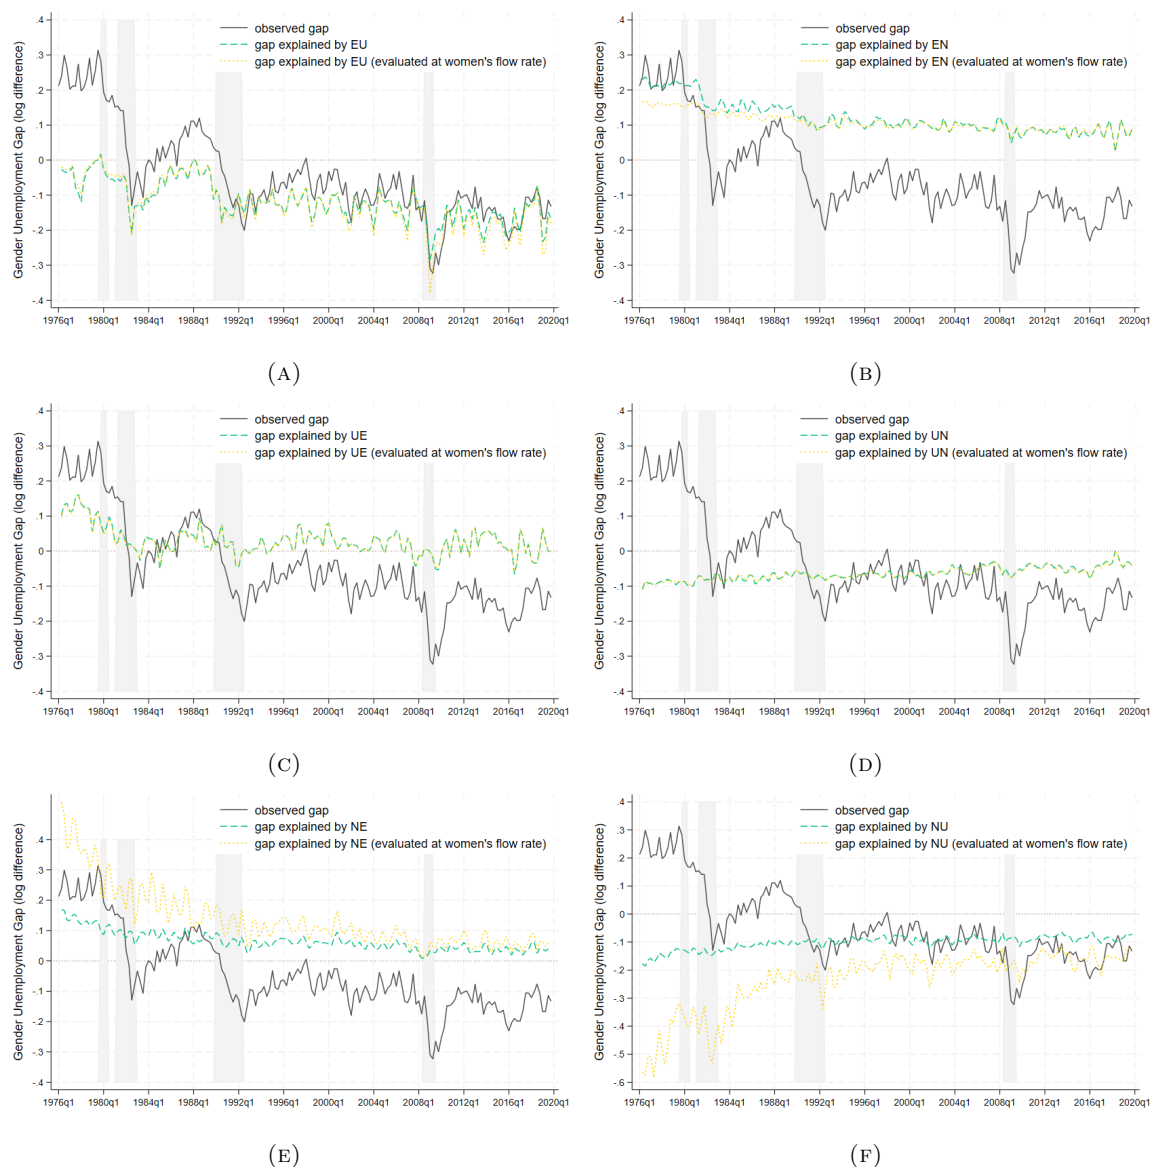


FIGURE C.4: Actual Gender Unemployment Rate Gap and the Proportion Explained by Each Individual Flow, with the partial derivative of steady-state unemployment rate with respect to each flow evaluated at women's flow rates. Panel (a) is the gender unemployment gap explained by employment-to-unemployment flow, panel (b) is the gender unemployment gap explained by employment-to-non-participation flow, panel (c) is the gender unemployment gap explained by unemployment-to-employment flow, panel (d) is the gender unemployment gap explained by unemployment-to-non-participation flow, panel (e) is the gender unemployment gap explained by non-participation-to-employment flow, and panel (f) is the gender unemployment gap explained by non-participation-to-unemployment flow. The Solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the proportion of the gender unemployment gap that is explained by each flow. The dotted line represents the proportion of the gender unemployment gap that is explained by each flow, with the partial derivative of the steady-state unemployment rate evaluated at women's flow rates. Data Source: the LFS master files.



FIGURE C.5: Actual Gender Unemployment Rate Gap and the Proportion Explained by Each Individual Flow, with the partial derivative of steady-state unemployment rate with respect to each flow evaluated at the mid-point flow rates. Panel (a) is the gender unemployment gap explained by employment-to-unemployment flow, panel (b) is the gender unemployment gap explained by employment-to-non-participation flow, panel (c) is the gender unemployment gap explained by unemployment-to-employment flow, panel (d) is the gender unemployment gap explained by unemployment-to-non-participation flow, panel (e) is the gender unemployment gap explained by non-participation-to-employment flow, and panel (f) is the gender unemployment gap explained by non-participation-to-unemployment flow. The Solid line represents the observed unemployment gap (log difference between women and men unemployment rate), and the dashed line represents the proportion of the gender unemployment gap that is explained by each flow. The dotted line represents the proportion of the gender unemployment gap that is explained by each flow, with the partial derivative of the steady-state unemployment rate evaluated at the mid-point flow rates. Data Source: the LFS master files.

TABLE C.1: Gender Difference in Employment-to-Unemployment Transition Rate - Without Occupation

	(1)	(2)	(3)	(4)
Female	-0.343*** (0.0503)	-0.355*** (0.0439)	-0.0670** (0.0197)	-0.258*** (0.0287)
Tenure Deviation				-0.00668*** (0.000643)
age group fixed effect	No	Yes	Yes	Yes
education fixed effect	No	Yes	Yes	Yes
industry fixed effect	No	No	Yes	Yes
part-time status	No	No	No	Yes
N	26,491,173	26,491,173	26,491,173	26,491,173
R^2	0.001	0.005	0.008	0.011

Notes: The sample covers years from 1976 to 2019. The dependent variable is a binary variable equal to one if a worker becomes unemployed in the next month (rescaled to 100). ‘Female’ is a gender dummy variable. ‘Tenure Deviation’ denotes how many months a worker’s tenure is greater (or less) than the industry-occupation average tenure. All regressions control for province and month-year fixed effects. Regressions are weighted using the LFS sampling weights. Standard errors, clustered at the provincial level, are shown in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

TABLE C.2: Proportion of Gender Unemployment Gap Explained by Each Individual Flow

	1976-2019	1976-1989	1990-2019
EU	1.052	.015	1.487
EN	-.517	.844	-1.090
UE	-.107	.279	-.270
UN	.355	-.351	.652
NE	-.274	.446	-.577
NU	.531	-.537	.979

Notes: The proportion of the gender unemployment gap explained by each flow is calculated by the contribution of a certain flow divided by the gender unemployment gap. The contribution of each flow is calculated by equation (3.6).