Essays on Economic Vulnerability and Inequality

by

Gaëlle Simard-Duplain

BA (Honours), Concordia University, 2011

MA, University of British Columbia, 2013

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Economics)

The University of British Columbia
(Vancouver)

October 2019

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

Essays on Economic Vulnerability and Inequality

submitted by Gaëlle Simard-Duplain in partial fulfillment of the requirements for

the degree of Doctor of Philosophy

in Economics

Examining Committee:

Prof. David Green, Economics, University of British Columbia
Supervisor

Prof. Hiroyuki Kasahara, Economics, University of British Columbia
Supervisory Committee Member

Prof. Florian Hoffman, Economics, University of British Columbia
University Examiner

Prof. Sanghoon Lee, Sauder School of Business, University of British Columbia
University Examiner

Additional Supervisory Committee Members:

Prof. Giovanni Gallipoli, Economics, University of British Columbia
Supervisory Committee Member

Abstract

In the first chapter, I develop and estimate a model of life-cycle labour supply that incorporates the role of divorce. To do so, I set a collective model of household decision-making in an intertemporal context. The reduced-form literature has produced contradictory results on the effect of divorce and divorce risk on women's labour decisions. My model provides a unifying framework within which to view these findings. It also contributes to the structural literature, which has mostly studied how divorce alters bargaining power in marriage, and ignored women's insurance response to the risk of dissolution. I find that divorce risk decreases the lifetime expected income of married women, increasing their labour in all periods. However, this effect is mitigated over time, as women stay in and learn about their marriage. Among women who experience divorce, it exacerbates pre-existing differences across marriages.

In the second chapter, we investigate the role played by education in the intergenerational mobility of Canadian children. We use the Longitudinal and International Study of Adults (LISA), a panel of integrated survey and administrative data covering 1982 to 2013. We estimate that the education level of children accounts for one third to one half of the correlation between their income and their parents'. Furthermore, our estimates show comparable experiences of intergenerational mobility across individuals with differing levels of education. Both results suggest untapped opportunities in the education system to improve people's mobility.

In the third chapter, we exploit the Longitudinal Administrative Databank (LAD) to study the evolution of wealth inequality in Canada, from 1982 to 2011. Until now, research has relied largely on the Survey of Financial Security, which has sporadic coverage, under-coverage at the top of the wealth distribution, and a small sample size. We use the income capitalization method to take advantage of the higher-frequency data and large sample size provided by the LAD. Consistent with existing work, we find that the top 10% share was fairly constant over the period considered. However, our results differ in that we observe that the top 1% share grew moderately but steadily between 1995 and 2007.

Lay Summary

I present three essays on economic vulnerability and inequality in the Canadian context. In the first chapter, I study the impact of divorce and divorce risk on the labour supply of women. I find that women who are more likely to divorce work more while married. Divorce itself exacerbates differences across women that already existed during their marriage. In the second chapter, we study the role of education in intergenerational mobility. We find that education is associated with one third to one half of the correlation between parents' and children's incomes. Furthermore, people who differ in their educational attainment have a similar experience of mobility. Finally, in the third chapter, we outline the evolution of wealth inequality in Canada, from 1982 to 2011. Our results suggest that the wealth share of extremely wealthy individuals has grown moderately but steadily between 1995 and 2007.

Preface

Chapter 1 is original, unpublished and independent work by the author, Gaëlle Simard-Duplain. Chapter 2 constitutes joint work with Xavier St-Denis (Statistics Canada) and Chapter 3, with Pablo Gutiérrez-Cubillos. In both cases, respective co-authors were equally involved in all stages of the research project, including the identification of the research question, review of the literature, preparation and analysis of the data, and writing of the paper.

Table of Contents

Abstract . . . . **iii**

Lay Summary . . . . **iv**

Preface . . . . **v**

Table of Contents . . . . **vi**

List of Tables . . . . **ix**

List of Figures . . . . **x**

Acknowledgments . . . . **xi**

Introduction **1**

1 The effect of divorce on women's labour supply: A life-cycle perspective **7**

1.1 Introduction 7

1.2 Literature review 11

1.3 Economic model 15

1.3.1 Set-up 16

1.3.2 Labour supply equations 18

1.3.3 Effect of marital dissolution 20

1.4 Data 25

1.5 Empirical analysis 26

1.5.1 Unobserved variables: marginal utility of wealth 27

1.5.2 Unobserved variables: intra-household bargaining power, quality of time, and economies of scale in consumption 35

1.6 Endogeneity 38

1.6.1 Marital status and divorce probability 39

1.6.2 Endogenous switching (sample selection) 41

1.6.3 Endogenous regressors: divorce probability 43
1.6.4 Endogenous regressors: own and spousal wages 44
1.6.5 Final procedure 45
1.7 Results 47
1.7.1 Effect on married women 47
1.7.2 Effect on divorced women 52
1.7.3 Results and the reduced-form literature 55
1.8 Model with myopia 63
1.9 Conclusion 70

2 An exploration of intergenerational income mobility with the Longitudinal and International Study of Adults for Canada 72
2.1 Introduction 72
2.2 Literature review 74
2.3 Data 77
2.4 Child income and characteristics across the parental income distribution 79
2.4.1 Transition matrices 79
2.4.2 Characteristics of respondents by their parental income quintile 81
2.5 Methodology 83
2.5.1 Estimating the role of education in the intergenerational transmission of income 84
2.6 Results 87
2.6.1 The role of education in the intergenerational transmission of income 87
2.6.2 Horizontal inequality and differentiation by parental background among higher education graduates 90
2.6.3 Intergenerational mobility across groups 92
2.6.4 The mediating role of job characteristics among employed individuals 96
2.7 Conclusion 98

3 The evolution of wealth inequality in Canada between 1982 and 2011: Evidence from personal income tax records 100
3.1 Introduction 100
3.2 Concept of wealth 103
3.3 Data and methodology 103
3.3.1 National Balance Sheet Accounts 104
3.3.2 Longitudinal Administrative Databank 105
3.4 Results 107
3.4.1 Changes in wealth inequality in different groups 107
3.5 Accounting for non-taxable capital income 112
3.6 Conclusion 115

Conclusion 117

Bibliography 119

A Appendix to Chapter 1 129

A.1 Labour supply equations: derivations 129

A.1.1 Intertemporal decision-making of the divorced household 129

A.1.2 Intertemporal decision-making of the married household 130

A.2 Data 135

A.2.1 Development of the family files 135

A.2.2 Marital trajectories 136

A.3 Marginal utility of wealth: derivations 138

A.3.1 Marginal utility of wealth in the married problem, $v_{i,t}$ 138

A.3.2 Marginal utility of wealth in the divorced problem, $v_{1,i,t}$ 140

A.4 Labour supply with myopia: derivations 142

B Appendix to Chapter 2 145

List of Tables

Table 1.1 Effect of divorce risk on married women . . . 48
Table 1.2 Summary of effects on married women: own wage, spousal wage, and quality of time 49
Table 1.3 Summary of effects on married women: divorce probability . . . 51
Table 1.4 Effect of divorce on divorced women . . . 53
Table 1.5 Summary of effects on divorced women: own wage, spousal wage, and quality of time 53
Table 1.6 Summary of effects on divorced women: divorce probability . . . 54
Table 1.7 Effect of divorce on divorced women, with myopia . . . 68
Table 1.8 Summary of effects on divorced women, with myopia: own wage, spousal wage, and quality of time . . . 69
Table 1.9 Summary of effects on divorced women, with myopia: divorce probability . . . 69

Table 2.1 Sample size, linkage rate and mean parent age by respondent birth year . . . 80
Table 2.2 Child family income quintile conditional on parental family income quintile . . . 81
Table 2.3 Child education by parental family income quintile . . . 82
Table 2.4 Job quality by parental family income quintile, employed respondents . . . 84
Table 2.5 Decomposition of relative intergenerational mobility . . . 88
Table 2.6 Decomposition of relative intergenerational mobility by projection decomposition method . . . 90
Table 2.7 Decomposition of relative intergenerational mobility among university graduates . . 91
Table 2.8 Varying relative intergenerational mobility level by highest degree obtained . . . 94
Table 2.9 Group differences in relative and absolute intergenerational mobility, total income . 95
Table 2.10 Decomposition of relative mobility in employment income, employed respondents . 97

Table 3.1 Equivalence between asset classes and capital income flows . . . 106

Table B.1 Child outcomes by parental total family income quintile . . . 145
Table B.2 Possible values of skill intensity index score . . . 145
Table B.3 Decomposition of relative intergenerational mobility . . . 146
Table B.4 Decomposition of relative intergenerational mobility, employment income (employed respondents only) . . . 147

List of Figures

Figure 2.1 Job skill intensity by parental total family income decile 83

Figure 3.1 Wealth composition from 1970 to 2011 105
Figure 3.2 Overall trends in wealth concentration 108
Figure 3.3 Overall trends in wealth concentration, comparison of LAD and SFS estimates . . 108
Figure 3.4 Trends in wealth concentration, within sex 109
Figure 3.5 Trends in wealth concentration, within cohort: middle cohorts 110
Figure 3.6 Trends in wealth concentration, within cohort: older cohorts 110
Figure 3.7 Trends in wealth concentration, within cohort: younger cohorts 111
Figure 3.8 Trends in wealth concentration, pseudo life-cycle patterns 112
Figure 3.9 Trends in wealth concentration, within regions 113

Acknowledgments

I wish to thank David Green for his guidance and support throughout my years at UBC. I am deeply grateful for the expertise, time and energy he contributed to my development as an economist. I also want to thank Giovanni Gallipoli and Hiroyuki Kasahara, who were both extremely generous with their time and encouragement, and who helped me become a better-rounded researcher. I am grateful to Matilde Bombardini, Erwin Diewert, Mukesh Eswaran, Florian Hoffmann, Vadim Marmer, Kevin Milligan and Marit Rehavi for very useful suggestions and discussions at different stages of my doctoral studies. Finally, I wish to thank Catherine Douglas, for being both a mentor and a friend throughout the PhD, and for sharing her expertise and passion for community engagement without reserve.

I am greatly indebted to Statistics Canada's Longitudinal and International Study of Adults (LISA) team, for taking me in and for their enthusiasm to collaborate on the development of the family files. In particular, I wish to thank Andrew Heisz, for his continued vision and optimism, and Isabelle Thony, for her leadership. I am also extremely grateful to Statistics Canada's Income Processing for T1 Family Files (T1FF) team, for their assistance and patience in the construction of the LISA family files.

Of course, the second and third chapters of this thesis would not be without my co-authors, Xavier St-Denis and Pablo Gutiérrez-Cubillos. I have greatly appreciated their commitment to the respective projects and their willingness to share their own expertise. I also wish to thank my colleagues at UBC, who have offered both support and economic insight over the course of my PhD: Ashton Brown, Nathan Canen, Tom Cornwall, Alastair Fraser, Brad Hackinen, Melanie Kleinplatz, Denis Kojevnikov, Bert Kramer, Jiayi Liu, Neil Lloyd, Tímea Laura Molnár, Nouri Najjar, Rogerio Santarrosa, Iain Snoddy, and Dylis Xu.

I am grateful to Corey Lapell, Sunghee Lapell, Leon Bayer, Katie Bayer, and all of the CrossFit Empower community, for becoming home away from home, and for providing perspective when perspective was needed.

Finally, I wish to thank my family: my mother, for being an ever-present source of inspiration and strength; my father and Gisèle, for being there despite the distance; and Jocelyn and Caroline, for making Ottawa feel like home. I am also deeply grateful to Lanny, for the love and friendship, and for doing all of it, from economics to meal provision.

Introduction

In this thesis, I present three essays on economic vulnerability and inequality in the Canadian context. The first two chapters study channels through which family can contribute to economic vulnerability and to inequalities. More specifically, the first chapter looks at the effect of divorce risk on the labour supply of married women, and of divorce on those who experience it. As for the second chapter, it investigates the part played by education in explaining the correlation between parents' and children's incomes. Both papers also highlight the ways in which policies interact with families to influence this relationship. Importantly, existing research has found that the ability of the Canadian tax and transfer system to mitigate income inequalities has substantially decreased since the second half of the 1990s (Green et al., 2016). Another way in which differences across people may be expressed is through their ability to save, and thus to accumulate wealth. The third chapter focuses on the evolution of wealth inequality in Canada over the past three decades and on the persistence of wealth rankings for given individuals.

In the first chapter, I examine how divorce influences the labour supply of married women, from the time of marriage to the period following dissolution. Divorce is both widespread and associated with negative outcomes. In Canada, nearly 40% of marriages are expected to end in divorce [1], with first marriages making up a little over 80% of divorces (Statistics Canada, n.d.-f).[2] Furthermore, divorce is correlated with female entries into poverty (Curtis and Rybczynski, 2014), divorced individuals exhibit higher morbidity and mortality rates (Joung et al., 1997), and children of divorced parents are more likely to experience downward income mobility in childhood (Burton et al., 2014) and marital dissolution in adulthood (Corak, 2001). Whether these relationships are causal matters for policy. If divorce leads to poor outcomes, then governments should favour policies that support women through their transition out of marriage. Otherwise, these outcomes should be the focus of policy makers, irrespective of marital status.

There is little consensus on the effect of divorce on women's labour supply. Some researchers have found that divorce risk may increase married women's labour supply (eg. Bargain et al., 2012;

[1] More specifically, nearly 40% of marriages are expected to end in divorce before their 30th anniversary. This is the 30 year total divorce rate, or TDR-30. The TDR-30 and TDR-50 are standards used by Statistics Canada to express divorce rates.

[2] These statistics all correspond to numbers that were gathered for Canada in 2005. Vital statistics, including marriages and divorces are collected from the provinces, and detailed divorce data is currently only publicly available for 2004 and 2005.

Parkman, 1992; Peters, 1986), while others have concluded that it may have negligible or negative effects (eg. Gray, 1998; Johnson and Skinner, 1986). Similarly, research has both found that women's labour increases following divorce (eg. Johnson and Skinner, 1988) and doesn't change (eg. Mueller, 2005). This literature has provided little theoretical basis to make sense of its conflicting results. As a result, it is unclear what has been measured, how it informs the changes people undergo and, most importantly, how it should guide policy making.

To understand fully the implications of divorce for married women, I build a theoretical model in which spouses make decisions together, and they do so in an intertemporal setting. The theoretical model allows me to derive predictions for labour supply at different points in a woman's life, while accounting for the many aspects that characterize being in a couple, including economies of scale and intra-household bargaining power. Although the theoretical framework applies to both spouses in a couple, I focus on women. The model predicts that divorce affects them differently at different points of their marriage: at the time of marriage, throughout the union, and at divorce, if divorce does take place.

According to the theoretical model, the labour supply of married and divorced women depends on their own and their spouse's lifetime wage trajectories. Furthermore, labour supply depends on contemporary values and lifetime trajectories of intra-household bargaining power and economies of scale. Therefore, the data used must provide long income, employment and marital histories, include a rich set of individual and couple characteristics, and follow women and their ex-spouses over time. To estimate the model, I therefore exploit a new Canadian data set, the Longitudinal and International Study of Adults (LISA), which combines survey and administrative data. LISA is a biennial longitudinal survey, for which data was collected in 2012 and 2014. Each respondent was then linked to their tax records for 1982 to 2013, providing long income panels. Furthermore, I collaborated with Statistics Canada to extend the original LISA. For each respondent, I used personal income tax records to find their past and previous family members, including spouses, but also parents, children, and siblings. I then linked each family member to their own tax records, for years 1982 to 2013.

The estimates from the empirical analysis show that divorce and divorce risk impact the labour supply of women from the time they get married to the period following marital dissolution. A 10% increase in the lifetime divorce risk faced by women when they get married leads to a 4.34 to 5.88% increase in their labour supply throughout their marriage. This is similar in magnitude to the effect of a 10% increase in their lifetime wages. At the time of marriage, the expectation that they may get married in the future decreases women's lifetime income, increasing their labour supply in all periods. Over time, as women remain married, they re-adjust their labour supply downwards to reflect the fact that divorce risk has not materialized. The re-adjustment is greatest among women in couples where both spouses would face tighter budget constraints after divorce: it is approximately four times as large as in couples where spouses would face looser budget constraints. Furthermore, I find that the occurrence of divorce exacerbates pre-existing differences across marriages; that is, the effect of divorce on women's labour supply is conditional on the marriage they were in. Differences in spousal wage, quality of time spent with one's spouse and divorce risk while married have larger impacts on women's labour supply after than before divorce.

I make four contributions in the first chapter of this thesis. First, I demonstrate theoretically and empirically that marital dissolution affects women differently depending on the point in time that is considered. As such, I provide the first unifying framework for a literature that has produced largely contradictory estimates of the effect of divorce. Second, I constructed a new data set that allows me to estimate the full effect of divorce on women. In comparison, much of the literature has exploited changes to divorce law. In this chapter, I show that this produces estimates that ignore the baseline impact of divorce risk at the time of marriage, as well as subsequent adjustments. Third, by showing how divorce impacts women differently at different points in their life-cycle, I provide a stronger basis for policy. In the empirical portion of the paper, I estimate that the effect of divorce is conditional on the marriage that was terminated. That is, gaps in poverty levels or health outcomes across women of different marital status are not driven but rather amplified by divorce. Finally, I also contribute to the literature that has studied the mechanisms through which divorce impacts the labour supply of married women (eg. Chiappori et al., 2002; Voena, 2015). Until now, it has largely focused on the effect of changes to women's post-dissolution conditions change on their bargaining power during marriage. Alternatively, I account for the effect of bargaining power on labour supply, but my focus is on changes in labour supply that result from women insuring against divorce.

In the second chapter, a co-author and I explore a different channel through which family may impact individuals' economic outcomes: the transmission of earning potential from parents to children. More specifically, we evaluate the extent to which Canadians' experience of intergenerational mobility is associated with their education. In Canada, a 10% increase in a man's income is associated with a 2.3 to 3.8% increase in their son's adult income, and with a 1.9 to 2.9% increase in their daughter's income (Chen et al., 2017). This is half as much as in the United States, but 30% higher than in Finland, Norway or Denmark (Corak, 2013). Everywhere, institutions and policies that affect investments in and returns to human capital play a significant part in determining the extent to which different circumstances early on lead to gaps in later-life outcomes.

The literature on intergenerational mobility has broadly followed two courses. On the one hand, the literature on intergenerational mobility has sought to refine the measurement of mobility, exploring the effect of data limitations (eg. Atkinson et al., 1983; Solon, 1992; Zimmerman, 1992; Chen et al., 2017), but also seeking to draw comparisons over time (eg. Connolly et al., 2019) and across places (eg. Chetty et al., 2014; Chetty et al., 2014; Corak, 2017; Connolly et al., 2019). On the other hand, the development of increasingly rich data sets has given impetus to a growing literature on the mechanisms that underlie some of the broader numbers. Bridging these two strands are a small number of papers that have tried to decompose measures of intergenerational mobility; i.e., that have tried to map these numbers back to some of the investigated drivers. In the United States and the United Kingdom, these have found that approximately one half of the observed intergenerational correlation in income is associated with education (respectively Bowles and Gintis, 2002, and Blanden et al., 2007). This relationship hinges on two links: first, parental income must be associated with education, and second, education must be associated with traits and skills that are valued on the labour market. The institutional distinctions Corak (2013) draws between Canada and the United States suggest that both of these links are weaker here than south of the border. Therefore, we not only expect that the correlation would be smaller in Canada, but also that education would explain a smaller portion of the correlation that we do see.

As for the first chapter, we take advantage of the Longitudinal and International Study of Adults (LISA), which provides detailed survey data on respondents, as well as a panel of administrative data covering 1982 to 2013 for both respondents and their parents. The link between survey and administrative data allows us to measure the part played by education in the intergenerational income correlation of Canadians. To do so, we use regressions of child income rank on parental income rank. We decompose the intergenerational income correlation coefficient and find that educational attainment accounts for 38.2 to 50.1% of the association, depending on the measures of child and parental income used. Furthermore, we find evidence that field of study strengthens income correlation. To better understand the role of education, we further decompose intergenerational mobility to account for the role played by a number of skills respondents report using at work (reading, writing, communication, mathematics, physical strength and dexterity) and specific indicators of job quality (unionization status, permanence of contract and authority over other employees). We find that almost half of the portion of intergenerational mobility that is associated with education is itself correlated with job skills. However, very little of the role played by education has to do with selection into higher-quality jobs, net of job skills.

Our contribution is two-fold. First, we show that despite greater mobility in Canada, the portion of the intergenerational correlation in incomes that is explained by education is similar to what is reported for the United States and United Kingdom (respectively Bowles and Gintis, 2002, and Blanden et al., 2007), if not higher. This suggests a role for educational policy to further its potential in leveling the playing field. Second, to the best of our knowledge, our paper provides the first evidence on the part played by job skills in intergenerational mobility. In particular, we find that half of the portion of income correlation that is associated with education is linked to job skills, including reading, writing, communication, mathematics, physical strength and dexterity. This is useful to understand the role of education and related institutions in Canada, but it also contributes to the broader literature on the relative role of cognitive skills. Research has generally focused on ability measured in childhood and adolescence (eg. Blanden et al., 2007). Our analysis speaks directly to the skills people use at work, thus more closely to the skills people are remunerated for.

In the first two chapters of this thesis, I study channels through which family may contribute to economic vulnerability and to inequalities. Both papers also highlight the ways in which policies interact with families to influence this relationship. Importantly, existing research has found that the ability of the Canadian tax and transfer system to mitigate income inequalities has substantially decreased since the second half of the 1990s (Green et al., 2016). Another way in which differences across people may be expressed is through their ability to save, and thus to accumulate wealth. In the third chapter, a co-author and I document the evolution of wealth inequality in Canada, from 1982 to 2011. Research on the evolution of wealth inequality has produced often broadly varying results, both quantitatively and qualitatively, depending on the data source used (eg. Kopczuk, 2015; Kopczuk and Saez, 2004; Saez and Zucman, 2016). In general, five data sources have been used to measure wealth inequality: surveys of household assets and liabilities, estate tax records, personal income tax records, lists of wealthy individuals, and wealth tax data. Kopczuk (2015) provides a thorough review of the first four data sources, along with an in-depth discussion of the sources of differences in the estimates they have produced.

In Canada like elsewhere, existing research has faced a substantial challenge in coming to a conclusion on the state and evolution of wealth inequality, especially in more recent periods. Relying on estate tax records, Davies and Di Matteo (2018) find that the share of wealth held by the top 1% decreased between 1946 and 1970, from 29.6 to 19.6 %. For the period between 1970 and 2012, Davies and Di Matteo (2018) combine household survey data and rich lists, and find that wealth concentration has remained approximately constant over the period, with the top 1% holding approximately 23 to 24%. Earlier work using the same (unadjusted) household survey found somewhat different results. Oja (1983) concludes that the period between 1970 and 1977 was characterized by a moderate decline in wealth inequality, as measured by the percentage of wealth held by different quintiles. That change was driven in part by a decrease in the concentration of financial wealth. For that period, Morissette and Zhang (2006) found similar results. Between 1977 and 1984, their results are in line with Davies and Di Matteo (2018). According to them, wealth inequality later increased, between 1984 and 1999, due not only to an increase in net worth at the top of the distribution, but also to a decrease at the bottom of the distribution; that is, certain groups became worse off both in relative and in absolute terms (Morissette et al., 2006; Morissette and Zhang, 2006).

In this context, we contribute to the overall portrait of wealth inequality in Canada, by bringing to bear a new source of data. More specifically, we present new evidence on wealth inequality by applying the income capitalization method to the personal income tax records of Canadians. As mentioned, most of the existing research on wealth inequality has relied mostly on the SFS, conducted in 1999, 2005, 2012 and 2016, and on its precursor, the asset and debt supplement to the Survey of Consumer Finances (SCF), collected every seven years between 1957 and 1984. Although an invaluable source of information on the wealth held by Canadian households, it is limited in its coverage of the wealthier segments of the population, in its sample size, and in the frequency of its data collection. According to Statistics Canada, there is also evidence that survey respondents under-report their assets and liabilities, in particular their financial assets. To the extent that we expect people at different points of the distribution to differ in the composition of their portfolio, this implies that any analysis that relies on the Survey of Financial Security is likely to under-estimate wealth inequality at any given point in time. Whether it would affect trends is harder to determine. In addition to contributing to the overall understanding of what has happened with wealth inequality in Canada over the past three decades, we add to existing work by providing the first set of high-frequency (yearly) estimate for that period. This allows us to consider more specifically what happened around key moments, like changes in the redistributive role of the tax and transfer system in the 1990's, which is known to have contributed to the increase in after-tax income inequality (Green et al., 2016), the dot-com bubble and the Great Recession. Indeed, the Survey of Financial Security's sporadic coverage makes it ill-suited to investigate these events. Furthermore, our paper presents the first estimates of trends in individual-level inequality in Canada. Among other things, this makes for more straightforward comparisons with the literature on income inequality.

Consistent with earlier results from Davies and Di Matteo (2018), we find that the top 10% share has been fairly stable over the period considered, with a very modest decrease from 1988 to 1999 and a modest increase thereafter. Our higher-frequency data allows us to determine that the latter arises from an increase in the top 10% share between 1999 and 2007, followed by an almost equal decrease between 2007 and 2011. Furthermore, whereas Davies and Di Matteo (2018) found that the top 1% share was fairly stable between 1984 and 2012, our results suggest that very wealthy people have in fact become relatively wealthier between 1995 and 2007. Their gains were subsequently halved between 2007 and 2011. In general, our estimates of inequality are higher than in existing work: we calculate that the wealth share held by the top 10% and the top 1% respectively peak at 76 and 37%, compared to 56 and 23% in Davies and Di Matteo (2018).

Chapter 1

The effect of divorce on women's labour supply: A life-cycle perspective

1.1 Introduction

In this paper, I examine how divorce influences the labour supply of married women, from the time of marriage to the period following dissolution. In Canada, nearly 40% of marriages are expected to end in divorce[1], with first marriages making up a little over 80% of divorces (Statistics Canada, n.d.-f).[2] Furthermore, divorce is correlated with multiple negative outcomes. Between 1993 and 2010, approximately a third of women's entries into poverty coincided with a change of the household's main earner, often following marital dissolution (Curtis and Rybczynski, 2014). Divorced individuals also exhibit higher morbidity and mortality rates (Joung et al., 1997), and their children are more likely to experience downward income mobility in childhood (Burton et al., 2014) and marital dissolution in adulthood (Corak, 2001). If these relationships denote causation, then governments should favour policies that support women through their transition out of marriage. Conversely, divorce may simply be correlated with factors that contribute to poverty and poor health, such as adverse family background or low human capital. In that case, these outcomes should be the focus of policy makers, irrespective of marital status.

The reduced-form literature has produced often conflicting results about the effect of divorce on women's labour supply.[3] There is a general consensus that divorce risk affects married women, starting at the very beginning of their union (eg. Stevenson, 2007). However, the direction of the effect is

[1]More specifically, nearly 40% of marriages are expected to end in divorce before their 30th anniversary. This is the 30 year total divorce rate, or TDR-30. The TDR-30 and TDR-50 are standards used by Statistics Canada to express divorce rates.

[2]These statistics all correspond to numbers that were gathered for Canada in 2005. Vital statistics, including marriages and divorces are collected from the provinces, and detailed divorce data is currently only publicly available for 2004 and 2005.

[3]Other outcomes have also been studied beyond labour supply, including investment in human capital, housing and children (eg. Reinhold et al., 2013; Stevenson, 2007), and domestic abuse, suicide and spousal suicide (Brassiolo, 2016; Stevenson and Wolfers, 2006).

ambiguous: some researchers have estimated that divorce risk increases labour supply (eg. Bargain et al., 2012; Parkman, 1992; Peters, 1986), while others have found negligible or negative effects (eg. Gray, 1998; Johnson and Skinner, 1986). Studies have also compared the labour supply of women before and after divorce, with some finding increases in participation and hours worked (eg. Johnson and Skinner, 1988) and others concluding that dissolution leads to no significant changes (eg. Mueller, 2005). This literature has provided little theoretical basis to make sense of its conflicting results. It has regularly lumped together married people at different points of their union, and compared women before and after dissolution, without much consideration for the time elapsed since the event. As a result, it is unclear what has been measured, and how it informs the changes people undergo. Again, this has important implications for policy making. If divorce has a causal effect on the labour supply of married women, differences in its effect over their life course speaks to the mechanisms they draw upon to insure against dissolution.

To understand fully the implications of divorce for married women, I build a life-cycle model of household decision-making that formalizes the way in which marital dissolution affects individuals differently over their lifetime. Although the theoretical framework applies to both spouses in a couple, I focus on women. The model's starting point is the collective model of household decision-making, which incorporates the presence of economies of scale in consumption, the differential value of leisure time spent alone and with one's spouse, and the role of intra-household bargaining power. The model is set in an intertemporal context. Furthermore, it formally accounts for the fact that changes in labour supply may also affect the probability of divorce. The model predicts that divorce affects married people in three different ways. First, at the time of marriage individuals expect that they may divorce in the future. This shifts their expected lifetime income down, and they increase their labour supply accordingly in every period. As years pass and they learn about their union, that initial shift is mitigated; i.e., they gradually reduce the initial insurance taken against divorce. Finally the model predicts no effect following divorce, conditional on the marriage that ended.

According to the theoretical model, the labour supply of married and divorced women depends on their own and their spouse's lifetime wage trajectories. Furthermore, labour supply depends on contemporary values and lifetime trajectories of intra-household bargaining power and economies of scale. Therefore, the data used for the analysis must provide measures for these couple-specific variables. Furthermore, women who divorce must be included in the data set for panels long enough that they are observed before and after marital dissolution. The necessity for long panels suggests that administrative records may be good candidates for the study of divorce. However, they inform a very limited number of individual characteristics and as such are not suitable for the empirical implementation of the model I develop. Instead, I exploit a new Canadian data set, the Longitudinal and International Study of Adults (LISA), which combines survey and administrative data. LISA is a biennial longitudinal survey, for which data was collected in 2012 and 2014. Each respondent was then linked to their tax records for 1982 to 2013, providing long income panels.

The research question and theoretical framework also require that the income of partners be observed beyond marital dissolution; i.e., not just when they are part of a respondent's family. To address this, it was necessary to extend the original LISA.[4] For this purpose, I collaborated with Statistics Canada to create the LISA family files. For each LISA respondents, I used personal income tax records to find their past and previous family members, including spouses, but also parents, children, and siblings. I then linked each family member to their own tax records, for years 1982 to 2013. The extended LISA, including the family files, therefore includes long income and marital trajectories for respondents and their family members, as well as rich variables documenting the characteristics of respondents.

The estimates from the empirical analysis show that divorce and divorce risk impact the labour supply of women from the time they get married to the period following marital dissolution. A 10% increase in the lifetime divorce risk faced by women when they get married leads to a 4.34 to 5.88% increase in their labour supply throughout their marriage. This is similar in magnitude to the effect of a 10% increase in their lifetime wages. At the time of marriage, the expectation that they may get divorced in the future decreases women's lifetime income, increasing their labour supply in all periods. Over time, as women remain married, they re-adjust their labour supply downwards to reflect the fact that divorce risk has not materialized. The re-adjustment is greatest among women in couples where both spouses would face tighter budget constraints after divorce: it is approximately four times as large as in couples where spouses would face looser budget constraints.

Furthermore, I find that the occurrence of divorce exacerbates pre-existing differences across marriages; that is, the effect of divorce on women's labour supply is conditional on the marriage they were in. Differences in spousal wage, quality of time spent with one's spouse and divorce risk while married have larger impacts on women's labour supply after than before divorce. For instance, an increase in spousal wage at time t while women are married increases their labour supply in that period by 1.76%. However, it has no effect in other periods, until after divorce; i.e., the cross-effect of changes in spousal wage is limited to periods after dissolution. Similarly, the cross-effect of an increase in the quality of time spent with one's spouse is one and a half time as large after divorce; and the cross-effect of an increase in divorce risk is twice as large after marital dissolution.

I make four contributions in this paper. First, I demonstrate how marital dissolution affects women differently depending on the point in time that is considered. To do so, I incorporate divorce in a model of life-cycle labour supply with household decision-making. The reduced-form literature has extensively documented the gradual effect of divorce on labour supply, starting in marriage and unfolding through the transition. However, it has produced largely contradictory estimates of the direction of this effect. I estimate the parameters from my model and use them to show that previous findings are consistent with divorce impacting women differently over their lifetime. As such, I provide the first unifying framework for the existing literature's contradictory findings.

Second, I exploit a source of variation that allows me to estimate the full effect of divorce on women. To make this possible, I constructed the LISA family files, a set of tax records for all past and

[4] Notably, a related issue also affects traditional longitudinal surveys, which run the risk of suffering attrition correlated with household breakdown.

present family members of LISA respondents. This data is essential to implement my theoretical model. Notably, both the reduced-form and the structural literatures have relied heavily on variation in the legal framework that regulates marital dissolution. However, the effect that is estimated using that variation necessarily ignores the baseline impact of divorce risk at the time of marriage, as well as the subsequent adjustment that comes as people learn about their marriage. Instead, studies using variation from divorce law conflate the two. Indeed, when legislation changes, people re-set their lifetime expectations of divorce risk, having already acquired a certain amount of information depending on the time elapsed since their marriage.

Third, by showing how divorce impacts women differently at different points in their life-cycle, I provide a stronger basis for policy. In the empirical portion of the paper, I estimate that divorce amplifies existing differences across women. Its effect is therefore conditional on the marriage that was terminated. This has two important implications for how we formulate policies to help people through marital transitions. On the one hand, this result confirms that divorce is a significant shock, more than just a correlate of poor outcomes. Although economic and other differences across women may not be driven by marital dissolution, their effect is more severe after divorce than during marriage. On the other hand, the fact that divorce impacts women by exacerbating pre-existing factors means that corresponding policies should at least in part be targeted earlier in life, and at factors that directly affect women's economic security. Indeed, there are vulnerabilities that may not be realized during marriage but matter if and when women experience divorce.

Finally, I also contribute to the structural literature that has studied the mechanisms through which divorce impacts the labour supply of married women (eg. Chiappori et al., 2002; Voena, 2015). Until now, the focus has largely been on the effect that operates through women's bargaining power in marriage: as women's post-dissolution conditions change, for instance through legal or welfare reforms, so do their outside option, and thus their bargaining power. Alternatively, I control for the effect of bargaining power on labour supply, but my focus is on changes in labour supply that result from women insuring against divorce. In that context, I map the effect of divorce to women's observable characteristics. This allows me to identify women who may have a more limited access to the insurance mechanism, and to formulate more concrete targets for policy-making.

The rest of the paper proceeds as follows. In Section 1.2, I discuss the reduced-form and structural literatures that have studied the effect of divorce on women's labour supply. I then present my theoretical model in Section 1.3. In Section 1.4, I introduce the data used for the analysis. In Section 1.5 and Section 1.6, I respectively discuss the empirical strategy, and the approach used to deal with the various sources of endogeneity involved. I present results in Section 1.7. Finally, I develop an extension of the baseline economic model in Section 1.8, allowing for the possibly that spouses do not foresee the possibility of divorce. I conclude in Section 1.9.

1.2 Literature review

The reduced-form literature on divorce has extensively documented the ongoing nature of the divorce process, showing that its effects start several years prior to dissolution, and affect people well after the event. Johnson and Skinner (1988) found that women increased their hours worked following marital dissolution. They used data from the 1968 to 1982 waves of the PSID, including three years before and two years after divorce. Conversely, Mueller (2005) compared the labour supply of Canadian women one year before and one year after dissolution, and found no difference.

Research into the labour supply consequences of divorce risk on married individuals yielded even more equivocal results. Using the 1979 wave of the CPS, Peters (1986) and Parkman (1992) compared states with mutual consent divorce to states that had recently allowed spouses to unilaterally divorce each other. Both found higher labour force participation among married women in unilateral divorce states. Also using the CPS, Gray (1998) compared the changes in female labour force participation in both sets of states, between 1968, when no legislative changes had taken place, and 1979. He found no difference, leading him to conclude that previous results picked up the endogenous introduction of the legal changes. Using a simultaneous equations model to account for selection into divorce at the individual level, Johnson and Skinner (1986) estimated that living in a no-fault divorce state decreased hours worked. In Ireland, Bargain et al. (2012) found that female labour force participation increased following the legalization of divorce in 1996, but that hours didn't change. None of these studies distinguished women who differed in terms of the time elapsed since their marriage or the length of their exposure to the new legal framework. For instance, Johnson and Skinner (1986) compared the 1972 labour supply of women who divorce at some point in the following six years, with no consideration for the year in which they were married or the year in which their home state changed legislation. A notable exception is Stevenson (2007), who looked at the effect of the legislative changes to grounds for divorce on newly formed couples. She estimated that women in couples who married after the reforms were more likely to work in the first two years of marriage, compared to women married before the changes.

In addition to divorce risk, the reduced-form literature has investigated the effect of post-divorce conditions on married women's labour supply, focusing on the role of human capital, property distribution rules at divorce, and the (re-)marriage market. Parkman (1992) reproduced the work of Peters (1986), allowing for the effect of unilateral divorce to vary with women's characteristics. He found larger increases in labour force participation for younger, white women, which he attributed to greater potential losses in human capital. Gray (1998) found that the introduction of unilateral divorce increased female labour force participation in states with community-property division rules, matched by a decrease in home production. However, participation decreased in title-based states. Conversely, Stevenson (2007) found that married women were more likely to be employed in unilateral divorce states, regardless of property division rules. Both studies used data from the same time periods, but Stevenson (2007) focused on the first two years of marriage while Gray (1998) did not distinguish based on time since marriage. The model I develop in this paper allows divorce risk and post-divorce conditions to affect people differently at various points in their union. Notably, Johnson and Skinner (1988) found that controlling for time-invariant heterogeneity reversed the estimated effect of being divorced on labour supply, from negative to positive. In other words, the labour supply of couples who end up divorcing differs from that of continuously married couples, in a way that confounds the time-varying effect of divorce. To provide an interpretation for these time-invariant differences, I use theory to parametrize the fixed effect and map it back to observables.

A growing body of work in structural economics has informed the effect of divorce through its role as the outside option of marriage. Although not exclusively, much of this work has built on Chiappori (1988, 1992)'s collective model of decision-making, which incorporates the impact of intra-household bargaining power and economies of scale on individuals' labour and consumption decisions. Several of these contributions have sought to estimate key parameters of household decision-making, including intra-household bargaining power, economies of scale, and preference parameters, in a variety of contexts. In that vein, Chiappori et al. (2002) and Voena (2015) have used changes to divorce legislation as sources of exogenous variation in the intra-household bargaining power of spouses. Chiappori et al. (2002) focused on the contemporary effect of divorce legislation and the pool of potential new spouses as exogenous shifters of intra-household bargaining power. Using the 1988 wave of the PSID, they found that married women's labour supply is lower in states or counties where their bargaining power is higher. Notably, the factors that Chiappori et al. (2002) used as shifters of intra-household bargaining power were hypothesized to impact the latter by improving the value of being divorced. As a result, the labour supply they estimated likely captures a combination of changes in both bargaining power and divorce risk.[5]

As for Voena (2015), she explicitly set out to study the effect of changes to the legal framework of divorce, on the portion of household resources allocated to women as well as on their labour supply. Unlike Chiappori et al. (2002), she set her model in an intertemporal context. It predicts that conditions in divorce affect resource allocation in marriage only when divorce can be initiated unilaterally. In turn, the passage to unilateral divorce only affects couples to the extent that it re-distributes resources across spouses; i.e., it can have no effect in states with title-based marital property division. Using

[5] Chiappori et al. (2002) use four aspects of divorce legislation in their analysis: grounds for divorce, property division, enforcement of support orders, and treatment of spousal degrees as assets in divorce proceedings. Their empirical implementation treats the following as factors that increase women's bargaining power in marriage: mutual consent divorce, community property, payment of support orders directly to court officers, and treatment of degrees as divisible property at divorce. In their main estimation, they construct an index that sums the four dummy variables, and find that women work less in states with higher index values. While they attribute this effect to an increase in women's bargaining power, two things are worth noting. First, the hypothesized mechanism is that these characteristics impact women's bargaining power by improving their outcome in the case of divorce; i.e., by improving the value of being divorced. As a result, both bargaining power and divorce risk are changing. Second, it is not fully straightforward to determine the direction in which bargaining power and divorce risk are changing. On the one hand Voena (2015) shows that it is the interaction between grounds for divorce and property division laws that determine the impact on bargaining power. The index used by Chiappori et al. (2002) is not constructed to pick up these interactions. On the other hand, although the literature has widely accepted unilateral divorce as increasing divorce risk, the discussion in Voena (2015) shows this may not be true in all property division regimes. Furthermore, even under community property, adjustments in intra-household resource allocation may imply that divorce risk does not systematically increase under unilateral divorce.

the PSID (1968 to 1993), the NLSMW (1967 to 1999) and the NLSYW (1968 to 1999), Voena (2015) showed that the data matched these predictions: married female employment decreased and household assets increased following the introduction of unilateral divorce, but only with equal property division rules. In turn, Voena (2015) used these reduced-form results to back out an estimate of intra-household bargaining power. In the structural estimation, she found that women who enjoyed a relatively small share of resources in marriage benefited from common property division rules, and correspondingly decreased their labour supply. On the other hand, women in couples with fairly equal resource distribution were actually worse off after the introduction of these rules.

Like in this paper, Voena (2015) demonstrated that divorce risk and conditions in divorce interact in their influence on married women. However, I answer a different question and correspondingly use a different source of variation. Indeed, Voena (2015) studied the effect of legislative changes pertaining to divorce on married couples. Her contribution informs the effect of a one-time change in divorce risk and post-dissolution conditions. As such, the effect estimated is necessarily net of the baseline impact of divorce. To the best of my knowledge, my paper is the first to evaluate the full effect of divorce on married women, starting from the beginning of their union, and allowing for that effect to change as years pass. The variation used by Voena (2015) corresponds to a shift in the remaining lifetime resources women expect to receive. The income effect of that change affects labour supply in all subsequent periods. However, I show in this paper that the shift is mitigated over the course of people's marriages, as they learn about their marriage. Although Voena (2015)'s model is intertemporal in allowing current decisions to depend on future outcomes, it is similar to the reduced-form literature in that it assumes that that effect is the same across all married people, irrespective of the point at which they are in their marriage.

Low et al. (2017) and Mazzocco et al. (2013) also looked at divorce as the outside option of married people, setting their respective models in an intertemporal context. Low et al. (2017) focused on the introduction of time limits to welfare benefits in the United States. They modelled these changes as affecting married people's post-dissolution budget constraint, and thus their bargaining power in marriage. Low et al. (2017) found that welfare use and divorce decreased as a result, while labour supply increased. These results held for couples already married at the time of the reform, as well as for couples who were married after the changes. While their analysis recognized that different mechanisms may be at play for the two groups, it ultimately provided little structure to compare the two. In Mazzocco et al. (2013), the key mechanism was intra-household specialization; that is, the extent to which men's and women's labour supply and household production change with respect to one another before, during and after marriage. In their model, specialization resulted from differential wage processes, driven in part by human capital accumulation, fertility, which was hypothesized to exacerbate existing specialization, and marital surplus. Importantly, the role of marital surplus was through the effect of divorce risk on labour supply: decreases in marital surplus increased divorce risk, thus increasing incentives for human capital accumulation during marriage (and vice versa). However, the structural parameters that Mazzocco et al. (2013) estimated do not directly inform the impact of changes in divorce risk on the labour supply of women. Using the PSID from 1984 to 1996, they showed that their model successfully replicated changes in male and female labour supply and household production before, during and after marriage. Interestingly, their model was least successful in replicating divorce hazard at different points in marriage: the simulated hazard was too high in early marriage and (much) too low in later years. From their model's standpoint, this implies that actual female labour supply was inexplicably high in early marriage, and inexplicably low in later years. This is consistent with the fact that their model does not explicitly account for the impact of divorce risk on labour supply. That being said, more work would be needed to understand the disparity. This paper constitutes one step in that direction, since I directly estimate the effect of divorce risk, while accounting for several of the intra-household factors that Mazzocco et al. (2013) incorporate. Future research should extend the model presented here to include the impact of the human capital accumulation incentive that is considered by Mazzocco et al. (2013), to clarify how it relates with divorce risk.

The starting point for the theoretical framework I use in this study is Chiappori (1988, 1992)'s collective model of decision-making. That is, I use a model of household decision-making that incorporates the role of intra-household bargaining, economies of scale, and the quality of the relationship between spouses in determining the labour supply of spouses. To understand the effect of divorce on women, it is necessary to use a theoretical framework that allows to analyse individuals both when they're part of a broader household, that is when they're married, and when they're divorced; in other words, to follow them as they transition from the time that they're married and making decisions with their spouse, to the time after divorce. This is not possible if married households are modelled according to the unitary framework. However, my contribution in this paper differs from the large portion of the collective model literature which has focused on the estimation of the parameters that characterize the household problem, including the intra-household bargaining power and/or resources allocation. Although I explicitly recognize the importance of these factors by proxying for them in the estimation of labour supply and divorce risk, my goal is not to estimate them.

More recently, the structural literature pertaining to the collective model has furthered its investigation of bargaining power and other household characteristics, by studying their determination in the context of endogenous marriage formation (eg. Chiappori et al., 2012; Goussé et al., 2017). This literature has also explored the general equilibrium effects of changes that impact marriage formation, both within households and beyond, in the labour market. Unlike this line of work, I take existing marriages as given, in both the theoretical and the empirical sections of the paper. As mentioned, one of this paper's contributions is to highlight the fact that women respond to divorce risk from the very beginning of their marriage. Ignoring the endogenous formation of marriage implies that I cannot distinguish between the effect of differences in lifetime divorce risk profiles across women from the effect of other trade-offs, that may have been made for them (and their spouse) to enter low- or high-divorce-risk marriages. For instance, marrying someone with weaker or more unstable labour market attachment may result in a higher-divorce-risk marriage. People may be more likely to enter these types of marriages if they have more affinity with their spouse; i.e., if the quality of time they spend together is greater.

Without formally incorporating marriage formation in the theoretical framework, it's impossible to distinguish the effect of divorce risk from the effect of these trade-offs. In other words, the effect I estimate is necessarily a partial equilibrium effect; it doesn't account for the fact that other household characteristics may change to convince people to enter high-divorce-risk marriages. More broadly, it doesn't account for marriage-market-level factors that underlie and respond to changes in the divorce risk profiles faced by individuals.

1.3 Economic model

The starting point for the model developed in this study is Chiappori (1988, 1992)'s collective model of decision-making. Married households are made up of two individuals, each with their own preferences, whose collective choices are characterized by Pareto efficient outcomes. Like Voena (2015), I set this model in an intertemporal framework. The possibility of divorce constitutes a downward shift in women's expectation about their lifetime income, and further affects them as they acquire more information throughout their marriage. The transfer of resources that takes place at dissolution determines their post-divorce trajectory.

The collective model of decision making allows me to account explicitly for the role of intra-household bargaining power, as well as of for the pecuniary and non-pecuniary gains to marriage. These are crucial for assessing the effect of family transitions like divorce. By affecting the composition of the household, they alter not only the resources available to each spouse, but also the prices they face, through economies of scale. Over the years, the theory for the collective model has evolved to account for different types of preferences (Chiappori, 1992; Chiappori and Ekeland, 2006); for public goods and economies of scale (Blundell et al., 2005; Browning et al., 2013); for exogenous determinants of intra-household bargaining power (Browning and Chiappori, 1998); for non-participation in the labour market (Blundell et al., 2007); and for limited commitment issues in intertemporal settings (Chiappori and Mazzocco, 2017; Mazzocco, 2004, 2007).[6]

Early contributions to the literature have documented the success of the collective model in explaining the observed behaviour of households (eg. Chiappori, 1988, 1992). As mentioned previously, a number of papers have studied the effect of changes in intra-household bargaining power on married couples' labour supply (eg. Chiappori et al., 2002; Voena, 2015).[7] Conversely, family transitions themselves have been used as sources of identifying variation to estimate couple-specific parameters, such as complementarities in spouses' leisure (eg. Couprie, 2007; Michaud and Vermeulen, 2011). In this study, I do not seek to identify intra-household bargaining power or the extent of economies of scale. As such, I don't need observations on consumption or detailed time use data. Rather, I rely on

[6] Theoretical work has also studied the implications of the collective model for investment in education (Chiappori et al., 2009), children and child support payments (Chiappori and Weiss, 2006, 2007); and the effect of different divorce laws on the divorce rate (Chiappori and Weiss, 2007).

[7] It has also been used empirically to estimate economies of scale and resource sharing within households, and to investigate factors or events that affect these parameters (eg. Alessie et al., 2006; Bargain and Donni, 2012; Bargain et al., 2014; Browning et al., 2013; Cherchye et al., 2012; Couprie et al., 2010; Dunbar et al., 2013 ; Lewbel and Pendakur, 2008; Lise and Seitz, 2011).

the insight from the collective model to inform the role these variables play in determining the effect of divorce.

The model is set in an intertemporal framework following MaCurdy (1981, 1985). In that context, labour supply depends on contemporary variables, including wages and preferences, as well as on a combination of past outcomes and expectations over future outcomes. In any given period, all past and future variables are fully captured by the current value of the marginal utility of wealth; i.e., the shadow price of relaxing the budget constraint. In a setting with uncertainty, MaCurdy (1985) showed that the marginal utility of wealth can be expressed as a random walk with drift; i.e., as a function of initial conditions and a linear time trend, with an error. As will be further discussed in Section 1.5, this allows for a clear mapping to observables, making full use of the rich data at hand. Key to the intertemporal labour literature is the distinction between differences in labour supply that arise across people who have different lifetime wage profiles, and changes in labour supply that arise as a person moves along their wage profile throughout their life. In the model, divorce risk affects married individuals' labour supply through those same channels: it shifts expected lifetime income down, and further alters people's choices as they acquire more information about their marriage every period.

As discussed in Section 1.2, the intertemporal framework is necessary to provide structure and discipline to the empirical estimation. This allows me to formulate predictions about the effect of divorce at different stages of women's life-cycle. It yields parameters with clear interpretation, that can then serve as the basis to think about the potential role of policy with respect to divorce. In turn, this cannot be done without acknowledging more recent contributions from the collective model literature, which have shown that the probability and consequences of divorce interact with intra-household factors. The model yields the following key predictions:

1. At the time of marriage, expectations over future divorce risk shift lifetime income down. This increases the labour supply of women in all periods.

2. As couples progress through their married years together, that initial income effect is mitigated; that is, the decrease (increase) in labour supply becomes smaller in magnitude every period. Intuitively, every period where spouses remain married, previous uncertainty about their marital status is resolved; the incentive for asset accumulation to insure against divorce decreases with the time since marriage.

3. Conditional on the assets inherited from the marriage, divorce has no impact on the divorced after the initial adjustment at dissolution.

1.3.1 Set-up

Household members make decisions collectively, given their individual preferences. In any given period, each divides his or her time across three activities: private leisure (eg. reading a book); public leisure, or time spent with their spouse (eg. hiking); and work. In addition, they consume a private and a public good, and choose how much assets to carry over to the next period.

Let i denote households, where $i = 1,\ldots,I$; t denotes time, where $t = 0,\ldots,T$; and s denotes spouses, where $s = 1,2$. Assume that each spouse knows with certainty that they will live for $T+1$ periods. The utility function of spouse s in couple i at time t is $U_{s,i,t}(l_{s,i,t}, L_{s,i,t}, c_{s,i,t}, C_{s,i,t})$, where $l_{s,i,t}$ represents their private leisure, $L_{s,i,t}$, their public leisure, $c_{s,i,t}$, their consumption of a private good, and $C_{s,i,t}$, their consumption of a public good. The labour supply of spouse s is $N_{s,i,t} = 1 - l_{s,i,t} - L_{s,i,t}$. Finally, $\mathrm{w}_{s,i,t}$, for $s = 1,2$, p_t and $Y_{i,t}$ are wages faced by the household, the price of the public good, and non-labour income. The price of the private consumption good is normalized to 1.

As discussed in Section 1.4, the main disadvantage of using personal income tax records is that wages are not observed, only earnings. Consequently, while I start from a theoretical model where individuals choose hours worked, the empirical strategy uses labour force participation as the dependent variable. Provided that preference shocks are well behaved over time, it can be shown that this produces estimates of the parameters of interest.

In what follows, I assume a Stone-Geary utility function,

$$U_{s,i,t} = \gamma_{l,s,i,t} \ln(l_{s,i,t} - \bar{l}_{s,i,t}) + \gamma_{L,s,i,t} \ln(L_{s,i,t} - \bar{L}_{s,i,t}) + \gamma_{c,s,i,t} \ln(c_{s,i,t} - \bar{c}_{s,i,t}) + \gamma_{C,s,i,t} \ln(C_{s,i,t} - \bar{C}_{s,i,t})$$

where $\gamma_{l,s,i,t}$, $\gamma_{L,s,i,t}$, $\gamma_{c,s,i,t}$ and $\gamma_{C,s,i,t}$ characterize preferences for the private and public leisure, and for the private and public goods; and $\bar{l}_{s,i,t}$, $\bar{L}_{s,i,t}$, $\bar{c}_{s,i,t}$ and $\bar{C}_{s,i,t}$ are corresponding basic needs levels. Note that $\bar{N}_{s,i,t} = 1 - \bar{l}_{s,i,t} - \bar{L}_{s,i,t}$ is the maximum amount of work spouse s can supply.

Married household members may enjoy consumption and leisure privately or with their spouse. Single individuals also consume the goods and take part in the activities that can be shared with a spouse, but they do so alone. For instance, housing is a public good for couples, but is consumed privately by single individuals; similarly, singles may travel or go to the movies on their own. Following Browning et al. (2013), I use a consumption technology function to relate the amount of a public good that is purchased by the household to its private-good equivalent; that is, to the total amount enjoyed by each household member. This allows to draw comparisons between households that face different prices based on their composition. Let $K_{i,t}$ be the amount of a public good purchased by the household, and $C_{i,t} = C_{1,i,t} + C_{2,i,t}$, the private-good equivalent. Then the consumption technology function, $F_K(\cdot)$, determines the conversion between the amount of public good purchased and the consumption it affords each household member: $K_{i,t} = F_K(C_{i,t})$. The function reflects the fact that a two-member household only needs to spend $p_t K_{i,t}$ on a public good for its individual members to enjoy respectively $C_{1,i,t}$ and $C_{2,i,t}$ of the good. If the two household members were to live separately, the sum of their private expenses on the public good would be $p_t C_{i,t}$. If the good is perfectly public, then $K_{i,t} = C_{1,i,t} = C_{2,i,t}$. However, the consumption technology function F_K also allows for goods that cannot be fully shared. For instance, if household members 1 and 2 were to live separately, they would buy amounts $C_{1,i,t}$ and $C_{2,i,t}$ of housing. By living together, they only buy $K_{i,t} \le C_{1,i,t} + C_{2,i,t}$, but it is unlikely that each enjoys $K_{i,t}$ of the housing. In other words, the consumption technology function determines the economies of scale from living in a married household. For now, I assume that there is only one private good and only one public good, although the model can easily be extended to allow for n goods, private or public, or partially private and public (Browning et al., 2013).

Similarly, I define a leisure technology function. Let $\Lambda_{i,t}$ be the amount of public leisure purchased by the household, and note that the price of that leisure is $\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}$, as each household member must contribute one unit of their own time. If they were to spend that time alone, $\mathrm{w}_{s,i,t}$ would buy them $l_{s,i,t}$; when the time is spent together, they each enjoy the private-leisure equivalent $L_{s,i,t}$. Hence, the leisure technology function is $\Lambda_{i,t} = F_{\Lambda}(L_{i,t})$, where $L_{i,t} = L_{1,i,t} + L_{2,i,t}$. In this context, $\Lambda_{i,t}$ is the total amount of public leisure purchased by the household, while $L_{i,t}$ is the total private-leisure equivalent enjoyed by the household.[8] The greater the private-leisure equivalent of a couple's public leisure, the greater the quality of the time they spend together.

1.3.2 Labour supply equations

The married household's problem depends in part on the situation faced by each spouse if divorce takes places. Therefore, I start by evaluating the divorced individual's problem. Let $s = 1$ denote the wife, and $Q_{1,i,t} = \{l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}, A_{i,t+1}\}$, the set of her choice variables at time t. After marital dissolution, the problem she faces can be expressed recursively as

$$V_{1,i}^{D}(A_{i,t}) = \max_{Q_{1,i,t}} U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) + \beta E_t\left[V_1^{D}(A_{i,t+1})\right]$$

subject to the budget constraint

$$\mathrm{w}_{1,i,t}(l_{1,i,t} + L_{1,i,t}) + c_{1,i,t} + p_t C_{1,i,t} + A_{i,t+1} = \mathrm{w}_{1,i,t} + Y_{i,t} + (1 + r_t)A_{i,t}$$

In the single household, note that the difference between the private and public leisure depends only on how they respectively enter the utility function; i.e., spouse 1 may have different preferences for different types of activities (eg. reading vs. hiking), but in a single household they have the same

[8]The leisure technology function is not the same as a production function where each spouse spends an amount of time $L_{s,i,t}$ to produce a public good $\Lambda_{i,t}$. Importantly, while the private-good equivalents consumed by the household sum up to $L_{1,i,t} + L_{2,i,t}$, this is not to say that the time of household members 1 and 2 are perfect substitutes in the context of public leisure. Examples of actual home production functions in the context of the collective model can be found in Couprie (2007) and Donni (2008). The leisure technology function is more restrictive. It assumes full complementarity between spouses' leisure time in some activities (eg. hiking). However, it does so in a very specific way; i.e., spouses get utility from the time spent together, $\Lambda_{i,t}$, purchased at price $\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}$. A consequence of this modelling choice is that the substitution between private-leisure equivalents, $L_{1,i,t}$ and $L_{2,i,t}$, depends only on preferences, not prices. That is, $\Lambda_{i,t}$ is chosen by the household based on relative preferences for the public leisure; relative wages have no impact. Browning et al. (2013) provide a useful way to think about public goods (and public leisure), by representing the household problem as a two-stage problem: first, the amount of public good (and public leisure) is chosen, along with the intra-household allocation of remaining resources; and second, household members make private consumption and leisure decisions individually, given the resources they received after the public purchases.

price. In Section A.1, I show that the labour supply of spouse 1 at time t is

$$\begin{aligned} N_{1,i,t} &= 1 - l_{1,i,t} - L_{1,i,t} \\ &= 1 - \bar{l}_{1,i,t} - \bar{L}_{1,i,t} - \frac{\gamma_{l,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} - \frac{\gamma_{L,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} \\ &= \bar{N}_{1,i,t} - \frac{\gamma_{l,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} - \frac{\gamma_{L,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} \end{aligned} \tag{1.1}$$

where $\nu_{1,i,t}$ is the Lagrange multiplier on the budget constraint; i.e., it is the marginal utility of wealth for the divorced household. Hence, the divorced individual's labour supply at time t can be expressed as a function of their maximum labour supply, their taste for the two types of leisure, their wage, and the Lagrange multiplier. In Section 1.3.3, I show that the latter is a function of initial assets, as well as initial expectations over wages, prices and interest rates, leisure and consumption needs, and preferences. Equation 1.1 is referred to as the marginal-utility-of-wealth-constant (MUWC) labour supply for the divorced individual (MaCurdy, 1981).

Next, I turn to the problem faced by the married household. Let $\lambda_{i,t}$ denote the resource share of individual 1 in the household utility function; and $\omega_{i,t+1}$, the probability that the couples remains together in the following period. Like Browning et al. (2013), I assume a linear consumption technology function, $K_{i,t} = F_K(C_{i,t}) = g_{K,i,t}C_{i,t} - \kappa_{K,i,t}$; and similarly for leisure, $\Lambda_{i,t} = F_\Lambda(L_{i,t}) = g_{\Lambda,i,t}L_{i,t} - \kappa_{\Lambda,i,t}$. In this paper, I refer to $g_{\Lambda,i,t}$ as the quality of time parameter. Let

$$Q_{i,t} = \{l_{1,i,t}, l_{2,i,t}, L_{1,i,t}, L_{2,i,t}, c_{1,i,t}, c_{2,i,t}, C_{1,i,t}, C_{2,i,t}, A_{i,t+1}\},$$

the set of choice variables for the married household. The problem is then expressed as follows

$$\begin{aligned} V_i^M(A_{i,t}) = \max_{Q_{i,t}} & \lambda_{i,t} U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) \\ & + \lambda_{i,t}\beta E\left[\omega_{i,t+1}V_{1,i}^M(A_{i,t+1}) + (1-\omega_{i,t+1})V_{1,i,t}^D(A_{i,t+1})\right] \\ & + (1-\lambda_{i,t})U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) \\ & + (1-\lambda_{i,t})\beta E\left[\omega_{i,t+1}V_{2,i}^M(A_{i,t+1}) + (1-\omega_{i,t+1})V_{2,i}^D(A_{i,t+1})\right] \end{aligned}$$

subject to the budget constraint

$$\begin{aligned} \mathrm{w}_{1,i,t}l_{1,i,t} + \mathrm{w}_{2,i,t}l_{2,i,t} &+ (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t})g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t})\kappa_{\Lambda,i,t} \\ &+ c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t\kappa_{K,i,t} + A_{i,t+1} \\ &= \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1+r_t)A_{i,t} \end{aligned}$$

where

$$V^M_{s,i}(A_{i,t}) = \max_{Q_{i,t}} U_{s,i,t}(l_{s,i,t}, L_{s,i,t}, c_{s,i,t}, C_{s,i,t}) + \beta E\left(\omega_{i,t+1} V^M_{s,i}(A_{i,t+1}) + (1-\omega_{i,t+1}) V^D_{s,i}(A_{i,t+1})\right)$$

for $s = 1, 2$. Note that $V^D_1(A_{i,t})$ and $V^D_2(A_{i,t})$ are the problems faced respectively by household members 1 and 2 if they experience marital dissolution. As shown in Section A.1, the MUWC labour supply of household member 1 when they are married is

$$N_{1,i,t} = \bar{N}_{1,i,t} - \frac{\lambda_{i,t}\gamma_{l,1,i,t}}{\nu_{i,t} w_{1,i,t}} - \frac{\lambda_{i,t}\gamma_{L,1,i,t}}{\nu_{i,t} g_{\Lambda,i,t}(w_{1,i,t} + w_{2,i,t})} \tag{1.2}$$

where $\nu_{i,t}$ is the shadow price of relaxing the married household's budget constraint. As for the divorced case, it is a function of initial assets and initial expectations over wages, prices and interest rates, leisure and consumption needs, and preferences. In addition, the married household's marginal utility of wealth also depends on initial expectations over intra-household bargaining power and economies of scale.

1.3.3 Effect of marital dissolution

Equation 1.2 and Equation 1.1 respectively express the labour supply of women when married and if divorced. As discussed in Section 1.2, the reduced-form literature on divorce has largely occupied itself with various comparisons of these two equations; including the difference between divorced and married labour supply for a given person, differences between the married labour supply of people who end up divorcing and of those who don't, differences between the married labour supply of people who face different divorce laws, etc.

Conversely, I use Equation 1.2 and Equation 1.1 to characterize the effect of divorce and divorce risk over women's life-cycle. Two things are worth noting. First, Equation 1.2 expresses the labour supply of married women at time t as a function of contemporary variables, such as own and spousal wages, and of the marginal utility of wealth, $\nu_{i,t}$. The latter is the only channel through which future divorce risk may affect married women.[9] Second, in the transition between Equation 1.2 and Equation 1.1, women face three changes: a decrease in the household's (potential) income; full access over household resources, as well as sole responsibility for household needs; and the loss of economies of scale in both consumption and leisure. Although these changes are experienced at the time of divorce, they are expected to impact married women prior to and even in the absence of dissolution, to the extent that the possibility of divorce is at least partly foreseeable. Indeed, individuals can smooth leisure and consumption over their lifetime, to minimize the shock of marital dissolution.

At the time of divorce, the change in labour supply therefore reflects any forecast error people may have made with respect to the timing of divorce, and to the magnitude and direction of its impact. After dissolution, divorce should have no effect, conditional on the budget constraint that ex-spouses inherited

[9] Alternatively, contemporary variables such as $\lambda_{i,t}$ could be allowed to depend on post-divorce outcomes, as in Voena (2015). The current specification ignores this channel to focus on the role of asset accumulation.

from the dissolution. Again, the effect of a past dissolution on the labour supply of married individuals is expected to operate mainly through the marginal utility of wealth of the divorced household.[10]

I formalize this intuition by relying on a result from MaCurdy (1985), who provides an expression for the marginal utility of wealth in an intertemporal framework with uncertainty. Extending it to account for divorce risk allows me to describe formally the channels through which divorce and divorce risk affect women, and to formulate predictions. All corresponding derivations are found in Section A.3.

Effect of divorce on married women

First, consider the lifetime budget constraint of the married household, and recall that $\omega_{i,t}$ denotes the probability that the couple remains together in period t. Furthermore, let $t = 0$ denote time of marriage. In addition to uncertainty over future relationship status, the household faces uncertainty about future wages and prices, preferences and needs, and intra-household bargaining power and economies of scale. Then,

$$\begin{aligned}
&A_{i,0} + E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \mathrm{w}_{1,i,t} + \sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \omega_{i,t}\mathrm{w}_{2,i,t}\right] \\
&\quad = E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t (1-\omega_{i,t})\left[\mathrm{w}_{1,i,t}(l_{1,i,t}+L_{1,i,t}) + c_{1,i,t} + p_t C_{1,i,t}\right]\right. \\
&\quad + \sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \{\omega_{i,t}\left[\mathrm{w}_{1,i,t}l_{1,i,t} + \mathrm{w}_{2,i,t}l_{2,i,t} + (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}(L_{1,i,t}+L_{2,i,t}) - (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})\kappa_{\Lambda,i,t}\right. \\
&\qquad\qquad \left.\left. + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t}+C_{2,i,t}) - p_t\kappa_{K,i,t}\right]\}\right]
\end{aligned} \tag{1.3}$$

where

$A_{i,0}$ is assets at the time of marriage;

p_t is the price of the public consumption good;

$l_{s,i,t}$, $L_{s,i,t}$, $c_{s,i,t}$ and $C_{s,i,t}$ are respectively the private and public leisure and consumption of spouse s.

[10] Alternative channels have also been explored in the literature. For instance, Mazzocco et al. (2013) models labour supply decisions during marriage as determinants of human capital, and therefore of the wages divorced individuals can rely on after dissolution. This study focuses on the asset channel, but this does not preclude further refinements of the model.

Substituting the MUWC functions determining leisure and consumption in Equation 1.3 yields

$$
\begin{aligned}
&A_{i,0} + E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \mathrm{w}_{1,i,t} + \sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \omega_{i,t}\mathrm{w}_{2,i,t}\right] \\
&= E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t (1-\omega_{i,t})\left(\mathrm{w}_{1,i,t}(\bar{l}_{1,i,t} + \bar{L}_{1,i,t}) + \frac{\gamma_{l,i,1,t} + \gamma_{L,i,1,t}}{\nu_{1,i,t}}\right)\right] \\
&+ E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t (1-\omega_{i,t})\left(\bar{c}_{1,i,t} + p_t\bar{C}_{1,i,t} + \frac{\gamma_{c,i,1,t} + \gamma_{C,i,1,t}}{\nu_{1,i,t}}\right)\right] \\
&+ E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \omega_{i,t}\left(\mathrm{w}_{1,i,t}\bar{l}_{1,i,t} + \mathrm{w}_{2,i,t}\bar{l}_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t})g_{\Lambda,i,t}(\bar{L}_{1,i,t} + \bar{L}_{2,i,t})\right.\right. \\
&\qquad\qquad \left.\left. + \frac{\lambda_{i,t}(\gamma_{l,i,1,t} + \gamma_{L,i,1,t}) + (1-\lambda_{i,t})(\gamma_{l,i,2,t} + \gamma_{L,i,2,t})}{\nu_{i,t}}\right)\right] \\
&+ E_0\left[\sum_{t=0}^{T}\left(\frac{1}{1+r_t}\right)^t \omega_{i,t}\left(\bar{c}_{1,i,t} + \bar{c}_{2,i,t} + p_t g_{K,i,t}(\bar{C}_{1,i,t} + \bar{C}_{2,i,t})\right.\right. \\
&\qquad\qquad \left.\left. + \frac{\lambda_{i,t}(\gamma_{c,i,1,t} + \gamma_{C,i,1,t}) + (1-\lambda_{i,t})(\gamma_{c,i,2,t} + \gamma_{C,i,2,t})}{\nu_{i,t}}\right)\right]
\end{aligned}
\tag{1.4}
$$

While the presence of uncertainty makes it impossible to derive a closed form solution for $\nu_{i,t}$, it is evident from this equation that $\nu_{i,t}$ is a complex function of preferences and needs, wages and prices, and intra-household bargaining power and economies of scale in all periods. Furthermore, it depends on divorce probability, $(1-\omega_{i,t})$, for $t = 0,...,T$.

With this in mind, let $\ln \nu_{i,t} = E_{t-1}(\ln \nu^*_{i,t}) + \varepsilon^*_{i,t}$; that is, the marginal utility of wealth at time t is equal to its expectation in the previous period, plus a one-period forecast error, $\varepsilon^*_{i,t}$. Equation 1.4 implies that this expectation is defined over all future variables, including wages, prices, intra-household bargaining power, divorce probability, etc., and so is the forecast error is defined with respect to all of these variables. In Section A.3, I show that the marginal utility of wealth of the married household at time t, $\nu_{i,t}$, can be expressed as

$$\nu_{i,t} = \nu_{i,0}\frac{1}{\prod_{j=0}^{t}\exp(b^*_{i,j})}\frac{1}{\prod_{j=0}^{t}\Omega_{i,j}}\prod_{j=0}^{t}\exp(\varepsilon^*_{i,j}) \tag{1.5}$$

where

$\nu_{i,0}$ is the marginal utility of wealth at time 0;

$\exp(b^*_{i,t}) = \beta(1+r_t)E_{t-1}(\exp(\varepsilon^*_{i,t}))$;

$\Omega_{i,t}$ is a term that arises from the probability of divorce at time t;

$\exp(\varepsilon^*_{i,t})$ is the one-period forecast error.

Recall that married women's labour supply is increasing in the couple's marginal utility of wealth, and consider what happens as we shift the various determinants of $\nu_{i,t}$. To simplify exposition, consider the two-period version of Equation 1.5,

$$\nu_{i,1} = \nu_{i,0} \frac{1}{\beta(1+r_1)E_0(\exp(\varepsilon^*_{i,1}))} \frac{1}{\Omega_{i,1}} \exp(\varepsilon^*_{i,1}) \tag{1.6}$$

First, $\nu_{i,0}$ denotes time-invariant differences between couples. Therefore, we expect couples with a higher initial marginal utility of wealth to work more in every period. A key insight of the life-cycle model of labour supply is that individuals who expect to earn higher (lower) wages throughout their lifetime will decrease (increase) their labour supply in every period (eg. MaCurdy, 1981). That is, shifts in $\nu_{i,0}$ have an income effect on labour supply. Among married couples, Equation 1.4 suggests that $\nu_{i,0}$ depends not only on the own wage profile, but also on the spouse's wage profile, and on the lifetime paths of intra-household bargaining power, economies of scale, and divorce probability. Second, a positive unanticipated shock at time 1, $\varepsilon^*_{i,1} > 0$, implies that $\nu_{i,t}$ grows between periods 0 and 1. For instance, a negative shock to current or expected future wages would translate into a negative income shock, thus an increase in $\nu_{i,t}$. Third, suppose that instead of an actual positive unanticipated shock at time 1, the couple had the expectation at time 0 that they would face such a shock at time 1; that is, $E_0[\exp(\varepsilon^*_{i,1})] > 1$. If the couple expects an increase in their marginal utility of wealth between periods 0 and 1, they will choose leisure, consumption and savings at time 0 to smooth out the transition.

Finally, consider the last term, $\Omega_{i,t}$. Denote $\phi_{1,i,0} = \frac{E_0(\nu_{1,i,1})}{E_0(\nu_{i,1})}$ and $\phi_{2,i,0} = \frac{E_0(\nu_{2,i,1})}{E_0(\nu_{i,1})}$; and recall that $\omega_{i,1}$ and $\lambda_{i,0}$ respectively refer to the probability that the couple remains together in period 1, and the intra-household bargaining power of the woman in period 0. Then,

$$\Omega_{i,1} = \omega_{i,1} + (1-\omega_{i,1})[\lambda_{i,0}\phi_{1,i,0} + (1-\lambda_{i,0})\phi_{2,i,0}] \tag{1.7}$$

As $\omega_{i,1}$ departs from 1 (as the probability of divorce increases), the effect on $\nu_{i,1}$ depends on $\lambda_{i,0}$, $\phi_{1,i,0}$ and $\phi_{2,i,0}$. Specifically, there are four possible scenarios:

1. $E_0(\nu_{1,i,1}) > E_0(\nu_{i,1})$ and $E_0(\nu_{2,i,1}) > E_0(\nu_{i,1})$; that is, both spouses face a tighter budget constraint after divorce.

   Then, $\Omega_{i,1} > 1$, and spouses set $\nu_{i,1}$ pre-emptively lower to make up for the possibility that it might increase in the following period.

2. $E_0(\nu_{1,i,1}) < E_0(\nu_{i,1})$ and $E_0(\nu_{2,i,1}) < E_0(\nu_{i,1})$; that is, both spouses face a looser budget constraint after divorce.

   Then, $\Omega_{i,1} < 1$, and spouses set $\nu_{i,1}$ pre-emptively higher to make up for the possibility that it might decrease in the following period.

3. $E_0(\nu_{1,i,1}) > E_0(\nu_{i,1})$ and $E_0(\nu_{2,i,1}) < E_0(\nu_{i,1})$; that is, the wife faces a tighter budget constraint after divorce, and the opposite holds for the husband.

Then, the size of $\Omega_{i,1}$ and the effect on $\nu_{i,1}$ depend on $\lambda_{i,0}$. The larger the intra-household bargaining power of the wife at time 0, the more the couple will make their decisions to account for her contingency; and vice versa.

4. $E_0(\nu_{1,i,1}) < E_0(\nu_{i,1})$ and $E_0(\nu_{2,i,1}) > E_0(\nu_{i,1})$; that is, the wife faces a looser budget constraint after divorce, and the opposite holds for the husband.

   Again, the size of $\Omega_{i,1}$ and the effect on $\nu_{i,1}$ depend on $\lambda_{i,0}$.

To see what these scenarios imply in terms of the effect of divorce on the labour supply of married women, consider the first case. If $\omega_{1,i,1} \in (0,1)$, then $\nu_{i,1}$ is lower than it would have been in the absence of divorce risk. Therefore, once period 1 is realized and the value of $\nu_{i,1}$ is set, spouses will work less, compared to a counterfactual scenario in which they didn't face divorce risk in period 0. Alternatively, this implies that couples over-saved in period 0; divorce risk causes them to shift consumption from the earlier periods to the later periods, to make up for the fact that they may face tighter constraints in those periods. As these periods are realized (and the couple remains together), the spouses decrease their labour supply to reflect the fact that they over-saved in earlier periods.

To summarize, divorce risk affects married women at two points in time. First, it shifts $\nu_{i,0}$, as evident from Equation 1.4. Second, as people progress through their marriage, they re-adjust their labour supply to account for the fact that they haven't divorced yet. If lifetime divorce-risk resulted in an increase of their labour supply in all periods, they revise their labour supply down as they progress through their marriage.

Effect of divorce on divorced women

Equation 1.1 shows that the effect of divorce on the divorced is fully captured by $\nu_{1,i,t}$, the marginal utility of wealth of the divorced household. As before, the result from MaCurdy (1981, 1985) can be extended to provide an expression for $\nu_{1,i,t}$. In Section A.3, I show that

$$\nu_{1,i,t} = \nu_{i,0}\left[\frac{1}{\prod_{j=0}^{t^d-1}\exp(b^*_{i,j})}\frac{1}{\prod_{j=0}^{t^d-1}\Omega_{i,j}}\prod_{j=0}^{t^d-1}\exp(\varepsilon^*_{i,j})\right]\frac{1}{\prod_{j=t^d}^{t}\exp(b_{1,i,j})}\prod_{j=t^d}^{t}\exp(\varepsilon_{1,i,j}) \tag{1.8}$$

where the term in square brackets corresponds to the evolution in the marginal utility of wealth of the married household, up to the period prior to dissolution; and

$\nu_{i,0}$ was the married couple's marginal utility of wealth at the time of marriage;

$\exp(b_{1,i,t}) = \beta(1+r_t)E_{t-1}(\exp(\varepsilon_{1,i,t}))$;

$\exp(\varepsilon_{1,i,t})$ is the one-period forecast error for the divorced household.

Therefore, divorce is expected to have no effect on the divorced, once the marginal utility of wealth inherited from marriage is accounted for.

1.4 Data

The theoretical model provides useful predictions to understand the role divorce plays throughout women's lives. However, its estimation imposes heavy demands on data. It must include own and spousal wage trajectories, as well as measures of intra-household bargaining power and economies of scale, in consumption and in leisure. Furthermore, all these must be available for panels sufficiently long that women who divorce are observed before and after dissolution. Even more so, their soon-to-be-ex-partners must also be observed before and after the transition. This last feature is especially challenging, since traditional longitudinal surveys run the risk of suffering attrition correlated with household breakdown. Finally, the data set used for the analysis must also include variables informative of preferences and other time-invariant characteristics that enter in the theoretical model. As such, administrative records that provide long panels are not suitable, as they inform a very limited number of sample member characteristics.

To bring the theoretical model to the data, I use the Longitudinal and International Study of Adults (LISA). While the LISA is a Canadian Survey, its name derives in part from the fact that its Wave 1 sample was partly built from the Canadian sample for the Program for International Assessment of Adult Skills (PIAAC). Demographic and individual income variables are drawn from the LISA. Specifically, I use the first two waves of the survey: Wave 1 was collected from November 2011 to June 2012, and Wave 2, from January to June 2014. The survey incorporates questions on core subjects and feature modules which change from one wave to the next; in addition, a portion of its content refers to long-passed events. As a result, the LISA provides detailed and long-term coverage of key outcomes, including family, education, and labour experiences, while recording less frequently measured variables likely to play an important role in moderating the effects of adverse events. Since the first wave included a number of people who participated in the PIAAC, it therefore counts with an assessment of literacy, numeracy, and problem-solving skills for approximately half of the sample. As for the second wave, it includes topics such as marital history, life events and life satisfaction, non-cognitive skills, pensions, and children.

The LISA sample was designed to be representative of the population of the Canadian provinces at the time of the first wave, in 2012. It is not limited to the working age population. Children of original LISA responding households are added to the sample of potential respondents for future waves of the survey as they turn 15 years old, and individuals who join households are also included through the recording of the household roster. Wave 1 of the LISA included approximately $32,000$ respondents and non-respondents, of which $25,500$ continued on to Wave 2, or approximately 80% of the initial sample. As a comparison, the Youth in Transition Survey (YITS) experienced 22% attrition between its first two cycles.

In addition to the two survey waves, the LISA data set has been merged with T1 Family Files (T1FF) for 1982 to 2013, with T4 files and Pension Plans in Canada (PPIC) files, both for 2000 to 2013, and with the Immigrant Longitudinal Database (IMDB) for 1980-2013. This linkage was initially done for LISA respondents only. For the purpose of this paper however, it was also necessary to access data that would allow on the one hand to determine family structure and family transitions, and on the other hand to follow past spouses after they had left respondents' households. To do so, I supplemented the initial administrative linkages by linking respondents to their past and current family members, including spouses and common-law partners, parents, children, and siblings. The data development methodology is outlined in Section A.2. The LISA and the administrative files have been available in Canada's network of Research Data Centres, or RDCs, since 2015. They now include the administrative records for respondents' family members, henceforth referred to as the family files. Together, the administrative files provide detailed economic and marital trajectories for LISA survey participants and their past and current family members.

1.5 Empirical analysis

The theoretical model yields the following expressions for women's labour supply, respectively when they are married and when they are divorced,

$$N^M_{1,i,t} = \bar{N}_{1,i,t} - \frac{\lambda_{i,t}\gamma_{l,1,i,t}}{\nu_{i,t}\mathrm{w}_{1,i,t}} - \frac{\lambda_{i,t}\gamma_{L,1,i,t}}{\nu_{i,t}g_{\Lambda,i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})} \tag{1.9}$$

$$N^D_{1,i,t} = \bar{N}_{1,i,t} - \frac{\gamma_{l,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} - \frac{\gamma_{L,1,i,t}}{\nu_{1,i,t}\mathrm{w}_{1,i,t}} \tag{1.10}$$

In turn, these suggest the following linear approximations as the basis for the empirical model

$$\ln N^M_{1,i,t} = F^M_i + \ln\nu_{i,t} + \beta^M_N \ln\bar{N}_{1,i,t} + \beta_\lambda \ln\lambda_{i,t} + \beta_g \ln g_{\Lambda,i,t} + \beta^M_1 \ln \mathrm{w}_{1,i,t} + \beta_2 \ln \mathrm{w}_{2,i,t} + u^M_{1,i,t} \tag{1.11}$$

$$\ln N^D_{1,i,t} = F^D_i + \ln\nu_{1,i,t} + \beta^M_N \ln\bar{N}_{1,i,t} + \beta^D_1 \ln \mathrm{w}_{1,i,t} + u^D_{1,i,t} \tag{1.12}$$

where

$$F^M_i = \ln\gamma_{l,1,i} + \ln\gamma_{L,1,i} \tag{1.13}$$

$$F^D_i = \ln\gamma_{l,1,i} + \ln\gamma_{L,1,i} \tag{1.14}$$

$$u^M_{1,i,t} = \eta_{l,1,i,t} + \eta_{L,1,i,t} \tag{1.15}$$

$$u^D_{1,i,t} = \eta_{l,1,i,t} + \eta_{L,1,i,t} \tag{1.16}$$

and

$\mathrm{w}_{s,i,t}$, $s = 1,2$, is the wage of spouse s;

$\lambda_{i,t}$ is the wife's intra-household bargaining power of spouse 1;

$g_{\Lambda,i,t}$ characterizes the economies of scale in leisure, or the quality of time spent with one's spouse;

$\nu_{i,t}$ and $\nu_{1,i,t}$ respectively denote the marginal utility of wealth of the married and divorced households at time t;

$\gamma_{l,1,i}$ and $\gamma_{L,1,i}$ respectively denote preferences for private and public leisure, and $\eta_{l,1,i,t}$ and $\eta_{L,1,i,t}$ are randomly distributed shocks to these preferences;

$\bar{N}_{1,i,t}$ is the maximum hours spouse 1 can work in period t;

t and t^d respectively denote time since marriage and time of marital dissolution.

The empirical model described by Equation 1.11 and Equation 1.12 presents two significant challenges. On the one hand, several of the right-hand-side variables are unobserved: the marginal utility of wealth for married and divorced households, $\nu_{i,t}$ and $\nu_{1,i,t}$, and the intra-household bargaining power and quality of time terms, $\lambda_{i,t}$ and $g_{\Lambda,i,t}$. On the other hand, even if these variables were observed, there are three sources of endogeneity that complicate the estimation. First, the married and divorced samples are unlikely to be determined randomly, so that sample selection must be accounted for in estimating the two equations. Second, $\Omega_{i,t}$ may be correlated with $u^M_{1,i,t}$, the error in Equation 1.11, either because labour supply also affects the probability of divorce, or because there are unobserved determinants common to both $N^M_{1,i,t}$ and $\Omega_{i,t}$. Finally, own and spousal wages are also susceptible to endogeneity. I address the strategy to deal with each source of endogeneity in Section 1.6.

1.5.1 Unobserved variables: marginal utility of wealth

For now, consider the approach to deal with, $\nu_{i,t}$ and $\nu_{1,i,t}$. From Section 1.3.3 and Section 1.3.3, we have that

$$\ln \nu_{i,t} = \ln \nu_{i,0} - \sum_{j=1}^{t} b^M_{i,j} - \sum_{j=1}^{t} \ln \Omega_{i,j} + \sum_{j=1}^{t} \varepsilon_{i,j} \tag{1.17}$$

and

$$\ln \nu_{1,i,t} = \ln \nu_{i,0} - \sum_{j=1}^{t^d-1} b^M_{i,j} - \sum_{j=1}^{t^d-1} \ln \Omega_{i,j} + \sum_{j=1}^{t^d-1} \varepsilon_{i,j} - \sum_{t^d}^{t} b^D_{1,i,t} + \sum_{j=t^d}^{t} \varepsilon_{1,i,j} \tag{1.18}$$

First, consider $\ln \nu_{i,0}$, which matters for both married and divorced women. From Equation 1.4, the marginal utility of wealth at time 0 depends on initial assets, as well as expectations at time 0 over own and spousal wages, intra-household bargaining power, economies of scale, prices, needs and preferences in all periods. Therefore, it is necessarily correlated with contemporary values of these variables, which belong on the right-hand side of Equation 1.11, and with any exogenous determinants that could be used to instrument for these variables. In brief, $\ln \nu_{i,0}$ cannot be left in the error term. One option is to use a fixed effect estimator to difference out the problem. However, this incurs an important loss of information. As mentioned before, part of the effect of divorce on married individuals is expected to act through shifts in $\ln \nu_{i,0}$. This alone is not insurmountable. MaCurdy (1981, 1985) estimates a simpler version of Equation 1.11 by first differences, then obtains an estimate of the fixed effect, which he uses as the dependent variable in regressions on time-invariant individual characteristics. However, although tax data offers the advantage of long panels with great coverage, hours worked are not observed. As a result, I use labour force participation as the dependent variable in the empirical analysis. Therefore, Equation 1.11 and Equation 1.12 specify a non-linear model, which complicates the estimation of the fixed effects. Generally, models with fixed effects face the incidental parameter problem, as described by Neyman and Scott (1948); that is, the number of individual parameters to be estimated grows with sample size, making their estimation inconsistent. With some exceptions, this inconsistency is then transmitted to the estimation of the common parameters. These exceptions broadly fall into three categories. First, some models allow for functional differentiation, where a function of the common parameters that does not depend on the individual effects is maximized. First difference transformations of linear models fall in this category. Second, the fixed effects may be parametrized with respect to observables, effectively reducing the problem to a random effects one. Third, under large-T asymptotics, the incidental parameter problem reduces to a bias problem, and the estimator (or moment condition, or concentrated likelihood) may be adjusted to correct for that bias. The particular problem at hand lands itself particularly well to the second approach, given that we're interested in the factors that shift $\ln \nu_{i,0}$, including expectations over future divorce probabilities. Hence, I incorporate the parametrization of the fixed effect directly into the estimation of Equation 1.11. Based on Equation 1.4, I assume that $\ln \nu_{i,0}$ can be approximated as follows

$$\begin{aligned}\ln \nu_{i,0} = Z_i\phi + \alpha^A A_{i,0} + \sum_{t=0}^{T}\alpha_t^{\mathrm{w}_1} E_0 \ln \mathrm{w}_{1,i,\underline{t}_1+t} + \sum_{t=0}^{T}\alpha_t^{\mathrm{w}_2}\left(E_0\ln\omega_{i,t} + E_0 \ln \mathrm{w}_{2,i,\underline{t}_2+t}\right) \\ + \sum_{t=0}^{T}\alpha_t^{g_\Lambda}\left(E_0\ln\omega_{i,t} + E_0\ln g_{\Lambda,i,t}\right) + a_{i,0}\end{aligned} \tag{1.19}$$

where Z_i is a vector of spouses' time-invariant characteristics that determine preferences and needs for leisure and consumption; $A_{i,0}$ denotes assets at the time of marriage; and $a_{i,0}$ is an error term. Note that the own and spousal wage profiles are necessarily defined with respect to the respective spouses' age, t_1 and t_2, while couple-specific variables are defined with respect to time since marriage, t. To bring them all on the same timeline, define $\underline{t}_s$, $s = 1,2$, as the age at which spouse s gets married. Then, $t_1 = \underline{t}_1 + t$ and $t_2 = \underline{t}_2 + t$. Finally, note that if $\omega_{i,t} = 1$, divorce is not a possibility and the $E_0 \ln \omega_{i,t}$ terms disappear.

With the exception of Z_i, the variables on the right-hand side of Equation 1.19 are not directly observed. To see how I address this, consider first $E_0 \ln \mathrm{w}_{1,i,t}$. This is the expectation that the woman in couple i formulates at time 0 about her lifetime wage path. Again I follow MaCurdy (1981, 1985), and take advantage of the data. Suppose that wages follow a quadratic path over the life-cycle; that is, a person's lifetime wage profile can be fully described by three parameters: its intercept, slope and curvature. Although the exact parameters are not observed, the goal is to estimate them for each woman in the sample. These must be estimated individually; i.e., it is not sufficient to regress wages on time. Indeed, the term $\ln v_{i,0}$ drives differences in labour supply choices across women. In that context, the parameters of a person's lifetime wage profile are individual-specific factors, which we eventually want to relate back to observable individual characteristics. Therefore, it would not be appropriate to simply regress wages on time, since it would give the same π_j^h parameters for everyone.

Given sufficiently many observations for each individual, MaCurdy (1981) shows that these parameters can be estimated from observed wage trajectories. Let $t_1 = age - 25$. I assume that a woman's own wage profile can be expressed as follows

$$\begin{aligned} E_0 \ln \mathrm{w}_{1,i,t_1} &= \ln \mathrm{w}_{1,i,t_1} + V_{i,t_1}^{\mathrm{w}_1} \\ &= \pi_{0,i}^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} t_1 + \pi_{2,i}^{\mathrm{w}_1} t_1^2 + \mathrm{v}_{i,t_1}^{\mathrm{w}_1} + V_{i,t_1}^{\mathrm{w}_1} \\ &= \pi_{0,i}^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} t_1 + \pi_{2,i}^{\mathrm{w}_1} t_1^2 + \mathrm{v}_{i,t_1}^{\mathrm{w}_1 *} \end{aligned} \tag{1.20}$$

where $V_{i,t_1}^{\mathrm{w}_1}$ denotes a forecast error. Let τ denote the number of sample periods used for the analysis. In this paper, I use a minimum of ten periods. Denote $t_1(j)$, the respondent's age in sample period j. Furthermore, define the difference operator D_k, such that $D_k \ln \mathrm{w}_{1,i,j} = \ln \mathrm{w}_{1,i,j} - \ln \mathrm{w}_{1,i,j-k}$. MaCurdy (1981) shows that estimates of $\pi_{2,i}^{\mathrm{w}_1}$, $\pi_{1,i}^{\mathrm{w}_1}$ and $\pi_{0,i}^{\mathrm{w}_1}$ can be obtained as follows

$$\tilde{\pi}_{2,i}^{\mathrm{w}_1} = \frac{1}{\tau-2} \sum_{j=1}^{\tau-2} \left[\frac{D_{j+1} \ln \mathrm{w}_{1,i,j+2}}{j+1} - D_1 \ln \mathrm{w}_{1,i,2} \right] \tag{1.21}$$

$$\tilde{\pi}_{1,i}^{\mathrm{w}_1} = \frac{1}{\tau-1} \sum_{j=1}^{\tau-1} \left[\frac{D_j \ln \mathrm{w}_{1,i,j+1}}{j} - \tilde{\pi}_{2,i}^{\mathrm{w}_1} (2t_1(j+1) - j) \right] \tag{1.22}$$

$$\tilde{\pi}_{0,i}^{\mathrm{w}_1} = \frac{1}{\tau} \sum_{j=1}^{\tau} \left[\ln \mathrm{w}_{1,i,j} - \tilde{\pi}_{1,i}^{\mathrm{w}_1} t_1(j) - \tilde{\pi}_{2,i}^{\mathrm{w}_1} (t_1(j))^2 \right] \tag{1.23}$$

The term $\tilde{\pi}_{2,i}^{\mathrm{w}_1}$ is meant to capture the curvature of spouse 1's lifetime wage trajectory. To see how this is the case, consider the two differences that make up the term in parentheses: $D_1 \ln \mathrm{w}_{1,i,2}$ and $\frac{D_{j+1} \ln \mathrm{w}_{1,i,j+2}}{j+1}$. The first one captures the slope of the wage profile between periods 1 and 2, and the second one captures the slope of the wage profile between period 1 and any period $j > 2$. If the profile was linear rather than quadratic, the two would be equal and the curvature parameter would equal 0. Any difference between the two parameters arises as a result of curvature. Equation 1.21 traces out the average curvature of the wage profile, by comparing the change between periods 1 and 2 to the change between period 1 and periods that are increasingly further in time. In turn, to estimate the slope of the wage profile, $\tilde{\pi}_{1,i}^{w_1}$ must take into account the curvature. Because the wage profile is not a straight line, the difference between any two periods, $D_j \ln w_{1,i,j+1}$, captures a mix of the slope and the curvature. Moreover, the contribution of the curvature is increasingly large the further way we are from the early years of the profile. Therefore, Equation 1.22 takes the difference $D_j \ln w_{1,i,j+1}$ and uses the estimate $\tilde{\pi}_{2,i}^{w_1}$ to correct it for curvature. This correction is proportional to the age at which wage is measured, $t_1(j+1)$. Finally, the intercept is recovered by using the estimates of $\tilde{\pi}_{1,i}^{w_1}$ and $\tilde{\pi}_{2,i}^{w_1}$ obtained in previous steps.

The parameters used in the estimation are obtained by projecting the $\tilde{\pi}_{j,i}^{w_1}$'s on individual time-invariant characteristics, $S_i^{w_1}$, such as education and family background. That is, I run the following regression,

$$\tilde{\pi}_{j,i}^{w_1} = S_i^{w_1} q_j^{w_1}, \ \ j = 0, 1, 2 \tag{1.24}$$

and compute $S_i^{w_1} \hat{q}_j^{w_1}$ to obtain $\hat{\pi}_{j,i}^{w_1}$, $j = 0, 1, 2$. This is done for three reasons. First and foremost, just as contemporary wages, lifetime wage profiles are likely to be endogenous determinants of labour supply. Second, doing so also provides a straightforward mapping between observables and labour supply. This gives results on the impact of individual characteristics on labour supply, allowing me to tie the results to something that can be observed and targeted in policy making. Third, it also reduces the sensitivity of the estimates of $\pi_{j,i}^{w_1}$, $j = 0, 1, 2$, to the period used for their estimation. In practice, I use wage trajectories of people who have worked at least ten periods, not necessarily consecutively. However, the $\tilde{\pi}_{j,i}^{w_1}$'s change based on when these periods occur in the lifetime of individuals. For instance, if someone is observed working from 25 to 34 years old, their curvature estimate will be relatively low, compared to someone who is observed working from 45 to 54 years old. Using the projection of the $\tilde{\pi}_{j,i}^{w_1}$'s on time-invariant characteristics is necessary to address endogeneity, but it also offers the advantage of restricting the $\pi_{j,i}^{w_1}$ estimates to well-behaved ranges. Essentially, I use the $\tilde{\pi}_{j,i}^{w_1}$'s of people with similar characteristics to extrapolate outside the range observed for each individual person. For people for whom I don't observe a minimum of ten periods, I impute the $\hat{\pi}_{j,i}^{w_1}$'s based on their characteristics.[11]

Having estimated the parameters that characterize the own wage profile, the same approach can be applied to the remaining profiles.[12] Then,

[11] For now, I ignore the possibility that there is non-random selection into the sample of women who are observed working for ten periods or more. To see what this implies, consider the regression of $\tilde{\pi}_{j,i}^{w_1}$ on own education, Ed_i, and father's and mother's education, respectively $ParedF_i$ and $ParedM_i$

$$\tilde{\pi}_{j,i}^{w_1} = \beta_{j,0} + \beta_{j,1} Ed_i + \beta_{j,2} ParedF_i + \beta_{j,3} ParedM_i + e_{j,i}$$

I assume that $E(e_{j,i}|work\ 10+) = 0$, and for all women I use

$$\hat{\pi}_{j,i}^{w_1} = \hat{\beta}_{j,0} + \hat{\beta}_{j,1} Ed_i + \hat{\beta}_{j,2} ParedF_i + \hat{\beta}_{j,3} ParedM_i$$

[12] Note that the $\pi_{j,i}^{h}$'s are not indexed by the spouse, $s = 1, 2$. I assume that within a couple spouses have the same expectation

$$E_0 \ln \mathrm{w}_{2,i,t_2} = \pi_{0,i}^{\mathrm{w}_2} + \pi_{1,i}^{\mathrm{w}_2} t_2 + \pi_{2,i}^{\mathrm{w}_2} t_2^2 + \mathrm{v}_{i,t_2}^{\mathrm{w}_2 *} \tag{1.25}$$

$$E_0 \ln g_{\Lambda,i,t} = \pi_{0,i}^{g_\Lambda} + \pi_{1,i}^{g_\Lambda} t + \pi_{2,i}^{g_\Lambda} t^2 + \mathrm{v}_{i,t}^{g_\Lambda *} \tag{1.26}$$

$$E_0 \ln \omega_{i,t} = \pi_{0,i}^{\omega} + \pi_{1,i}^{\omega} t + \pi_{2,i}^{\omega} t^2 + \mathrm{v}_{i,t}^{\omega *} \tag{1.27}$$

Substituting the expressions for all four profiles into Equation 1.19 and using the estimated values of $\pi_{j,i}^h$'s yields

$$\ln \nu_{i,0} = Z_i \phi + \alpha^A A_{i,0} + \hat{\Pi}_{0,i}' D + a_{i,0}^* \tag{1.28}$$

where

$$\hat{\Pi}_{0,i} = (\hat{\pi}_{0,i}^{\mathrm{w}_1}, \hat{\pi}_{1,i}^{\mathrm{w}_1}, \hat{\pi}_{2,i}^{\mathrm{w}_1}, \hat{\pi}_{0,i}^{\mathrm{w}_2}, \hat{\pi}_{1,i}^{\mathrm{w}_2}, \hat{\pi}_{2,i}^{\mathrm{w}_2}, \hat{\pi}_{0,i}^{g_\Lambda}, \hat{\pi}_{1,i}^{g_\Lambda}, \hat{\pi}_{2,i}^{g_\Lambda}, \hat{\pi}_{0,i}^{\omega}, \hat{\pi}_{1,i}^{\omega}, \hat{\pi}_{2,i}^{\omega})',$$

$$D = (\alpha_0^{\mathrm{w}_1}, \alpha_1^{\mathrm{w}_1}, \alpha_2^{\mathrm{w}_1}, \alpha_0^{\mathrm{w}_2}, \alpha_1^{\mathrm{w}_2}, \alpha_2^{\mathrm{w}_2}, \alpha_0^{g_\Lambda}, \alpha_1^{g_\Lambda}, \alpha_2^{g_\Lambda}, \alpha_0^{\omega}, \alpha_1^{\omega}, \alpha_2^{\omega})',$$

$$\alpha_j^h = \sum_{t=0}^{T} \alpha_t^h t^j, \quad j = 0, 1, 2, \;\; h \in \{\mathrm{w}_1, \mathrm{w}_2, g_\Lambda, \omega\},$$

$$a_{i,0}^* = a_{i,0} + \textstyle\sum_{t=0}^{T} \left[\alpha_t^{\mathrm{w}_1} \mathrm{v}_{i,t_1}^{\mathrm{w}_1 *}\right] + \sum_{t=0}^{T} \left[\alpha_t^{\mathrm{w}_2} (\mathrm{v}_{i,t}^{\omega *} + \mathrm{v}_{i,t_2}^{\mathrm{w}_2 *})\right] + \sum_{t=0}^{T} \left[\alpha_t^{g_\Lambda} (\mathrm{v}_{i,t}^{\omega *} + \mathrm{v}_{i,t}^{g_\Lambda *})\right]$$

The estimates of α_0^h, α_1^h and α_2^h, for $h \in \{\mathrm{w}_1, \mathrm{w}_2, g_\Lambda, \omega\}$, correspond to the effects of changes in the parameters of the respective profiles. In turn, these changes in these parameters constitute shifts in the present value of a woman's lifetime wealth; that is, they have an income effect on her labour supply. To gain some intuition for the parameters contained in D, consider again the profile for own wages, $E_0 \ln \mathrm{w}_{1,i,t}$. An increase in $\pi_{0,i}^{\mathrm{w}_1}$ corresponds to an upward parallel shift of the entire wage profile, it increases lifetime income. As a result, it is expected to decrease labour supply in all periods. An increase in $\pi_{1,i}^{\mathrm{w}_1}$ increases the slope of the profile, and also has a positive income effect on labour supply. Finally, an increase in $\pi_{2,i}^{\mathrm{w}_1}$ flattens the profile upwards. For later periods, when the passage of time has a greater negative then positive effect on wage, this flattening means a smaller downturn, and therefore a positive income effect. In brief, all three parameters are expected to have a positive effect on lifetime income, and thus to decrease labour supply in all periods. Similar to the parameters of the own wage profile, I expect spousal wage profile and quality of time parameters to decrease labour supply in all periods. Conversely, lifetime divorce probability should have a negative effect on income, and thus increase labour supply in all periods.

Ultimately, because the $\pi_{j,i}^h$'s are mapped back to observables, the $\alpha_{0,j}^h$'s also have an interpretation with respect to the effect of individual characteristics on labour supply. The time-invariant characteris-

over each other's wage profiles and the profiles of their couple-specific variables.

tics used depend on the profile being estimated. For own and spouse wages, I use the wife's education, as well as her father's and mother's education[13]. For the quality of time spent together, I use the age gap between the two spouses, and a measure of the language gap derived from their tax records. Finally, for divorce risk I use the wife's education, as well as indicators for whether she lived with both parents at birth and at 15 years old.

There is one more unobserved term in Equation 1.28, $A_{i,0}$, the couple's assets at the time of marriage. Indeed, although tax data includes measures of capital income, it does not include measures of baseline assets. I derive an estimate for $A_{i,0}$ following the same procedure as for the profile parameters. Let

$$y_{i,t} = \pi^{y}_{0,i} + \pi^{y}_{1,i}t + \pi^{y}_{2,i}t^2 + \mathrm{v}^{y*}_{i,t}$$

and

$$\tilde{\pi}^{y}_{j,i} = S^{y}_{i} q^{y}_{j}, \;\; j = 0, 1, 2$$

where $y_{i,t}$ and $A_{i,t}$ denote capital income and assets at time t, and $y_{i,t} = rA_{i,t}$. S^{y}_{i} is a vector of time-invariant characteristics that impact the lifetime profile of assets, including wife's education, her father's and mother's education and the spouses' age gap. At time of marriage, $y_{i,0} = \pi^{y}_{0,i} = r_0 A_{i,0}$, so that $A_{i,0} = \frac{\pi^{y}_{i,0}}{r}$. Substituting $\hat{A}_{i,0} = \frac{\hat{\pi}^{y}_{0,i}}{r}$ in Equation 1.28 gives

$$\ln \nu_{i,0} = Z_i \phi + \alpha^A \hat{A}_{i,0} + \hat{\Pi}'_{0,i} D + a^{*}_{i,0} \tag{1.29}$$

This is the expression for $\ln \nu_{i,0}$ that I use in the estimation of the model. Recall that $\ln \nu_{i,0}$ is the first of three unobserved terms that make up the marginal utility of wealth for the married household, $\ln \nu_{i,t}$, and similarly for the divorced household. Before proceeding to the remaining terms, it is worth unpacking the assumptions implicit in the substitution I've described.

Consider again the example for own wage, where the expected profile is assumed to take the form

$$E_0 \ln \mathrm{w}_{1,i,t_1} = \pi^{\mathrm{w}_1}_{0,i} + \pi^{\mathrm{w}_1}_{1,i} t_1 + \pi^{\mathrm{w}_1}_{2,i} t_1^2 + \mathrm{v}^{\mathrm{w}_1 *}_{i,t_1}$$

where $\mathrm{v}^{\mathrm{w}_1 *}_{i,t} = \mathrm{v}^{\mathrm{w}_1}_{i,t} + V^{\mathrm{w}_1}_{i,t}$ and $V^{\mathrm{w}_1}_{i,t}$ is a forecast error. Importantly, this forecast error is the difference between what is expected at time 0 and what actually happens at time t. In other words, it is the sum of

[13] Spousal education is not observed. Based on the assortative mating literature, I use own education in the spousal wage equation instead. In the LISA sample, 25 to 30 % of the variation in married men's education can be explained by their wife's and her parents' education. An alternative would be to use education credits to impute education level, but it only allows to classify people in a few categories; i.e., the estimation would differ for own and spousal wage.

one-period forecast errors between time 0 and time t. For instance, for own wages,

$$V_{i,t}^{w_1} = \sum_{j=1}^{t} \varepsilon_{i,j}^{w_1} \tag{1.30}$$

where $\varepsilon_{i,j}^{w_1}$ denotes the one-period forecast error on own wages. With this in mind, note that Equation 1.21 and Equation 1.22 can be re-written as

$$\tilde{\pi}_{2,i}^{w_1} = \pi_{2,i}^{w_1} + \frac{1}{\tau-2}\sum_{j=1}^{\tau-2}\frac{1}{j}\left[\frac{D_{j+1}\mathrm{v}_{i,j+2}^{w_1 *}}{j+1} - D_1\mathrm{v}_{i,2}^{w_1 *}\right] \tag{1.31}$$

$$\tilde{\pi}_{1,i}^{w_1} = \pi_{1,i}^{w_1} + \frac{1}{\tau-1}\sum_{j=1}^{\tau-1}\left[\frac{D_j\mathrm{v}_{i,j+1}^{w_1 *}}{j}\right] \tag{1.32}$$

Therefore, the second term on the right-hand side of each equation must be equal to zero. To see why and what this implies, consider first what would happen if I used the $\tilde{\pi}_{j,i}^{w_1}$'s directly in the estimation equation. To simplify exposition, suppose there are only five periods, $\tau = 5$, and focus on the implications with respect to forecast errors. For the second term of Equation 1.32 to be equal to zero, it must be that

$$\frac{1}{4}\left[\frac{V_{i,2}^{w_1} - V_{i,1}^{w_1}}{1} + \frac{V_{i,3}^{w_1} - V_{i,1}^{w_1}}{2} + \frac{V_{i,4}^{w_1} - V_{i,1}^{w_1}}{3} + \frac{V_{i,5}^{w_1} - V_{i,1}^{w_1}}{4}\right] = 0$$

$$\frac{1}{4}\left[\frac{\left(\sum_{j=1}^{2}\varepsilon_{i,j}^{w_1}\right) - \varepsilon_{i,1}^{w_1}}{1} + \frac{\left(\sum_{j=1}^{3}\varepsilon_{i,j}^{w_1}\right) - \varepsilon_{i,1}^{w_1}}{2} + \frac{\left(\sum_{j=1}^{4}\varepsilon_{i,j}^{w_1}\right) - \varepsilon_{i,1}^{w_1}}{3} + \frac{\left(\sum_{j=1}^{5}\varepsilon_{i,j}^{w_1}\right) - \varepsilon_{i,1}^{w_1}}{4}\right] = 0$$

$$\frac{1}{4}\left[\frac{\sum_{j=2}^{2}\varepsilon_{i,j}^{w_1}}{1} + \frac{\sum_{j=2}^{3}\varepsilon_{i,j}^{w_1}}{2} + \frac{\sum_{j=2}^{4}\varepsilon_{i,j}^{w_1}}{3} + \frac{\sum_{j=2}^{5}\varepsilon_{i,j}^{w_1}}{4}\right] = 0$$

As for the requirement on the second term of Equation 1.31, it implies that

$$\frac{1}{3}\left[\frac{\sum_{j=3}^{3}\varepsilon_{i,j}^{w_1}}{2} + \frac{\sum_{j=3}^{4}\varepsilon_{i,j}^{w_1}}{3} + \frac{\sum_{j=3}^{5}\varepsilon_{i,j}^{w_1}}{4}\right] = 0$$

Hence, using the $\pi_{j,i}^{w_1}$'s requires some balance over time in the shocks. This might be violated if a big negative (positive) shock in early periods is followed by a series of smaller positive (negative) shocks. For instance, if a woman marries a man whom she was wildly mistaken about, experiences a large negative shock early on in marriage, which she then spends several years recovering from. This scenario is very particular, but not impossible. In fact, we might be most worried about these types of cases with respect to divorce. If I was using the $\tilde{\pi}_{j,i}^{w_1}$'s directly in the estimation equation, then women who have this experience would introduce correlation between the covariates and the error term.

However, I use variation from $\tilde{\pi}^{w_1}_{j,i}$'s that is picked up by individual characteristics, namely own and parents' education. In fact then, that type of scenario is not problematic unless more educated women or women with more educated parents are also more or less likely to experience it. It seems reasonable to assume that any unbalanced pattern of forecast error is independent of education or parents' education.

Having established how to implement $\ln \nu_{i,0}$ empirically, and what that implementation implies, I consider the second term in Equation 1.17, $-\sum_{j=1}^{t} b^M_{i,j}$. MaCurdy (1985) shows that this term is constant over time and across individuals provided that the one-period forecast error is identically distributed over time and across individuals. Therefore, it can be controlled for with a trend with respect to the time since marriage, $b^M t$. Similarly, I replace $-\sum_{j=t^d}^{t} b^D_{1,i,j}$ in Equation 1.18 with $b^D(t-t^d)$.

Finally, we turn to the last part of Equation 1.17, $-\sum_{j=1}^{t} \beta_{\Omega,j} \ln \Omega_{i,j}$, where

$$\Omega_{i,j} = \omega_{i,j} + (1-\omega_{i,j})[\lambda_{i,t-1}\phi_{1,i,t-1} + (1-\lambda_{i,t-1})\phi_{2,i,t-1}]$$

This term corresponds to the time-varying effect of divorce risk on the labour supply of married women. Recall that $\phi_{1,i,t-1} = \frac{E_{t-1}(\nu_{1,i,t})}{E_{t-1}(\nu_{i,t})}$ and $\phi_{2,i,t-1} = \frac{E_{t-1}(\nu_{2,i,t})}{E_{t-1}(\nu_{i,t})}$. That is, $\phi_{s,i,t}$ determines the comparison between the expected value of the marginal utility of wealth in the married couple, and the marginal utility of wealth of spouse s if they were to be divorced. Although those factors are not directly observed, we know that they are determined by the extent to which the budget constraint faced by spouses when divorced is tighter or looser (in expectation) than that faced when married. Specifically, if $\phi_{s,i,t} > 1$, such that $E_{t-1}(\nu_{s,i,t}) > E_{t-1}(\nu_{i,t})$, then spouse s is expected to have a tighter budget constraint when divorced. Conversely, if $\phi_{s,i,t} \in (0,1)$, spouse s is expected to have a looser budget constraint when divorced.

At any time t, $\ln \Omega_{i,t}$ decreases $\ln \nu_{i,t}$, and thus labour supply, if $\Omega_{i,t} > 1$; conversely, it increases labour supply if $\Omega_{i,t} < 1$. From Section 1.3.3, recall that $\Omega_{i,t} > 1$ if both spouses are worse off when divorced; and $\Omega_{i,t} < 1$ if both spouses are better off divorced. If there's a substantial disparity in their situation in divorce, then the value of $\Omega_{i,t}$ depends on the interaction between that difference and resource allocation within the married household. Focus for now on the first case, where $\phi_{1,i,t-1} > 1$ and $\phi_{2,i,t-1} > 1$. This implies that as more time passes since marriage, labour supply will decrease as a response to divorce risk, in households where both spouses would be worse off in the occurrence of divorce. This insight makes more sense when we look at the entire expression for $\ln \nu_{i,t}$. Denote $G_{i,j} = \lambda_{i,j-1}\phi_{1,i,j-1} + (1-\lambda_{i,j-1})\phi_{2,i,j-1}$; then, the total effect of divorce risk on labour supply at time t is

$$\sum_{j=0}^{T} \alpha_{0,j} E_0 \ln \omega_{i,j} - \sum_{j=1}^{t} \beta_{\Omega,j} \ln[\omega_{i,j} + (1-\omega_{i,j})G_{i,j}]$$

which can be re-written as

$$\sum_{j=1}^{t} \left\{ \alpha_{0,j} E_0 \ln \omega_{i,j} - \beta_{\Omega,j} \ln[\omega_{i,j} + (1-\omega_{i,j})G_{i,j}] \right\} + \sum_{j+1}^{T} \alpha_{0,j} E_0 \ln \omega_{i,j}$$

where I've used the fact that $E_0 \ln \omega_{i,0} = 0$. This suggests that as people move through their married years, each additional period where they don't divorce cancels some of the original labour supply shift caused by the expectation of divorce risk. For people who are worse off if divorce occurs, the original effect is expected to shift labour supply up in each period, and the adjustment should shift labour supply back down. For people who are better off if divorce occurs, the original effect is expected to shift labour supply down in each period, and the adjustment should shift labour supply back up.

As mentioned previously, $E_{t-1}\nu_{i,t}$, $E_{t-1}\nu_{1,i,t}$ and $E_{t-1}\nu_{2,i,t}$ are not observed. Therefore, to implement these adjustments empirically, I broadly categorize individuals based on their estimated permanent income, where permanent income is computed as the wage profile estimated earlier, integrated from ages 25 to 65. I assume that if a women has permanent income below the median, she has $\phi_{1,i,t-1} > 1$ for all t, and similarly, similarly for her husband. Finally, I categorize couples into four groups: those where both are at the bottom of the permanent income distribution ($\phi_{1,i} > 1$ and $\phi_{2,i} > 1$); those where the husband is at the top but not the wife ($\phi_{1,i} > 1$ and $\phi_{2,i} < 1$); those where both are at the top; and those where the wife is at the top but not the husband ($\phi_{1,i} < 1$ and $\phi_{2,i} > 1$). Couples where spouses have lower permanent income benefit more from the economies of scale of living together, and are therefore expected to face a tighter budget constraint when divorced than when married.

Although the theoretical model predicts that the trajectory of $\Omega_{i,j}$ until time t enters in the determination of $\ln\nu_{i,t}$, I only include the contemporary value in the empirical specification. Correspondingly, I replace the term $-\sum_{j=1}^{t^d-1} b^M_{i,j} - \sum_{j=1}^{t^d-1}\ln\Omega_{i,j}$ in Equation 1.18 by t^d, the time of dissolution.

Finally, the sum of forecast errors, $\sum_{j=1}^{t}\varepsilon_{i,j}$ is part of the error term in the married labour supply equation, and $\sum_{j=1}^{t^d-1}\varepsilon_{i,j} + \sum_{j=t^d}^{t}\varepsilon_{1,i,j}$ is part of the error term in the divorced labour supply equation.

1.5.2 Unobserved variables: intra-household bargaining power, quality of time, and economies of scale in consumption

Before addressing the fact that $\lambda_{i,t}$ and $g_{\Lambda,i,t}$ are not observed in the data set used for this study, I discuss potential endogeneity problems that may arise if we were using the real variables. It provides a useful benchmark to discuss how to proxy for these variables. Again, recall that the error term for Equation 1.11 is

$$\eta_{l,1,i,t} + \eta_{L,1,i,t} + \sum_{j=1}^{t}\varepsilon_{i,j}$$

where $\eta_{l,1,i,t}$ and $\eta_{L,1,i,t}$ are shocks to preferences for private and public leisure. Therefore, the question is whether there are reasons to believe that either $\lambda_{i,t}$ or $g_{\Lambda,i,t}$, or both, may be correlated with either one of these terms. $g_{\Lambda,i,t}$ is defined as the equivalence between the amount of public good purchased by a household and the amount of private-good equivalent consumed by that same household, so that $\Lambda_{i,t} = g_{\Lambda,i,t}L_{i,t} - \kappa_\Lambda$; it corresponds to economies of scale in leisure. Consider first the relationship with $\eta_{l,1,i,t}$ and $\eta_{L,1,i,t}$. For $g_{\Lambda,i,t}$ to be exogenous, it must be that the extent to which couples enjoy time spent together more than time spent alone is independent from their enjoyment of those activities that can be shared with a spouse (and those that can't). For instance, it precludes the possibility that a positive shock to preferences for hiking impacts the extent to which hikes are enjoyed differently with a spouse. Furthermore, $g_{\Lambda,i,t}$ cannot be correlated with $\sum_{j=1}^{t}\varepsilon_{i,j}$; that is, relative enjoyment of time spent with one's spouse cannot change with different trajectories in one-period forecast errors. Remember that these forecast errors pertain to profiles of own wages, spousal wages, intra-household bargaining power, and economies of scale in leisure. One can imagine that couples who've been especially unsuccessful at forecasting poor profiles (several large, negative shocks) may have a strained relationship. I assume that this is not the case.

As for $\lambda_{i,t}$, we need to assume that intra-household bargaining power is uncorrelated with shocks to preferences for leisure. All else equal, the collective model literature typically finds that leisure is increasing in intra-household bargaining power, through the income effect of being allocated a greater portion of household resources. Hence, this assumption precludes the possibility that the allocation of resources induces shocks to leisure preferences; for instance, if knowing that one has less access to leisure makes leisure more desirable. Conversely, it also forbids scenarios where a lower (higher) taste for leisure increases (decreases) intra-household bargaining power directly; i.e., where taste for leisure affects bargaining power other than through its (potential) effect on wages. Finally, contemporary bargaining power and the cumulative set of forecast errors must also be uncorrelated. For instance, unexpectedly poor realizations of one spouse's wages cannot affect their resource allocation, other than through the wage effect; i.e., wages may impact intra-household bargaining power, but the forecast error may not.

Since $\lambda_{i,t}$ and $g_{\Lambda,i,t}$ are not observed, I proxy for them. Recall the general requirements for a proxy variable. Suppose the following generic relationship,

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \gamma q + v \tag{1.33}$$

where q is unobserved. For a variable z to be a suitable proxy, it must be that (1) $E(y|x_1,x_2,q,z) = E(y|x_1,x_2,q)$, i.e., the proxy variable does not belong in Equation 1.33; and (2) if $q = \theta_0 + \theta_1 z + r$, then $Cov(x_j,r) = 0$, for $j = 1,2$.

Intra-household bargaining power To proxy for intra-household bargaining power, I follow insight from the empirical collective model literature and use the local sex ratio (eg. Chiappori et al., 2002). Specifically, I use the ratio of single men to single women in the (female) respondent's age category and province.[14] The rationale is that greater numbers of men, relative to the number of women, imply greater chances for women to find an alternative spouse if divorce occurs. By improving their outside options, their share of resources in marriage increases. For this proxy to satisfy the first requirement, it must be that local sex ratios do not influence female labour supply, other than through the intra-household bargaining power. As for the second requirement, it implies that idiosyncratic determi-

[14] Population counts of single men and women for the years 1982 to 2013 are obtained from CANSIM table 17-10-0060-01 (Statistics Canada, n.d.-a).

nants of intra-household bargaining power are not correlated with the other right-hand-side variables in Equation 1.11. Perhaps the most obvious challenge to this is if unobserved love determines both intra-household bargaining power and the economies of scale in leisure. As discussed below, I proxy for quality of time spent with one's spouse in a way that account for this.

Quality of time To proxy for quality of time spent with one's spouse, I use one of the T1FF processing variables, the family flag variable, *fflag*, discussed in Section A.2. It is coded in a way that characterizes the quality of the match that was achieved when T1FF families were created. In any given year, a match between spouses may have been created if both spouses reported each other's SINs; if one spouse reported the other person's SIN and they were matched based on their address; or if one spouse reported the other person's SIN and the match was done on demographic characteristics. Alternatively, it is possible that spouses were unmatched in a given year. However, because the data allows me to follow families longitudinally, I can tell when individuals were married in a given year based on their status in contiguous years. Therefore, I code spouses that are unmatched in a given year as having the poorest match quality. Presumably, couples where spouses did not successfully report each other's SIN, where addresses didn't match, or where discrepancies in characteristics were large enough to be challenging for T1FF processing are couples where spouses are less effective at communicating or coordinating, or have less in common.[15]

For administrative match quality to be an adequate proxy for a couple's match quality, it mustn't matter for the determination of labour supply. An obvious concern is if people fail to file their taxes in years where they don't work. In that case, they wouldn't be matched with their spouse in that particular year, and would be given a low value of administrative match quality. However, tax filing rates have been very high, approximately 96%, since 1993, following the introduction of the GST and the conversion of the Child Tax Credit into the Child Tax Benefit. Both changes raised tax filing incentives for low-income families (Frenette et al., 2004; Frenette et al., 2007). In the data, there are in practice very few non-filing years among sample members and their spouses, once they appear in the data set. As for the second requirement, it is mostly unconcerning. Although unobserved love would be a problem if only one of bargaining power or quality of time was proxied for, it is not problematic in the case where both are proxied for.

It may be that poor administrative match quality is picking up disorganized couples; i.e., couples that are not particularly skilled at filing their taxes. I cannot discard this possibility. However, I note two things. First, if this is the case, I would argue that disorganized couples probably deal less effectively with the many demands on their time. If couples are overwhelmed with other obligations, they are unlikely to enjoy spending time together as much. Second, I could alternatively control for this by using deviation from couple-specific average administrative match quality; i.e., net out couples' average

[15] I only code unmatched individuals as spouses if they have appeared previously as married, and appear again as married later on. Therefore, it is possible that they were separated in the unmatched period, which would be in line with very poor quality of time spent together.

level of disorganization, and use only deviations from that level as the proxy for economies of scale in leisure.

Economies of scale in consumption A third variable required for the empirical analysis is unobserved, economies of scale in consumption, which appear as a determinant of marital status. I proxy for economies of scale in consumption by using the local relative price of public and private goods. In years where public goods are substantially cheaper than private goods, the economies of scale from living in a two-person households are larger. Specifically, I use the ratio of the price level of food purchased in restaurants to the price level of food purchased for consumption in the home. Statistics Canada used to produce an inter-city price index that compared the price of various goods across Canada's major cities within a given year. Although that series has since been discontinued, I use city- and good-specific price indices to expand the 2002 inter-city index to cover the period from 1982 to 2013.[16]

It is unlikely that local relative prices determine labour supply, so that the first requirement for the price ratio to proxy for economies of scale is not an issue. As for the second one, note that the collective model literature has found that several individual characteristics contribute to the economies of scale that couples face. That is because preferences for different goods translate into different savings for couples, depending on the bundles they purchase. I assume that these are second-order effects.

1.6 Endogeneity

As mentioned previously, three sources of endogeneity must be addressed to estimate the effects outlined in the previous section. First, the married and divorced samples are unlikely to be determined randomly, so that sample selection must be accounted for in estimating the two equations. Indeed, women who get divorced may have had a different series of one-period forecast errors than women who stay married. If that is the case, then the coefficients on assets at the time of marriage and on other components of $\ln v_{i,t}$ are expected to be biased. Second, $\Omega_{i,t}$ may be correlated with $u^M_{1,i,t}$, the error in Equation 1.11, either because labour supply also affects the probability of divorce, or because there are unobserved determinants common to both $N^M_{1,i,t}$ and $\Omega_{i,t}$. Finally, own and spousal wages are also susceptible to endogeneity.

The first two sources are related to the determination of marital status. In the next section I present a simple model of marital status, that can be related to observables and used to address these two issues. At the beginning of each period, spouses evaluate the married and the divorced problems. If no intra-household (re-)allocation of resources makes both spouses better off when married, then the marriage is dissolved. In Section 1.6.2, I address how the sample of married women is chosen to avoid bias from

[16] The inter-city price index comes from CANSIM table 18-10-0003-01 (Statistics Canada, n.d.-b) and the city- and good-specific price indices come from CANSIM table 18-10-0004-01 (Statistics Canada, n.d.-c). For the most part, the inter-city index uses data from each province's main city. Pendakur (2002) points out that as a result it is a better representation of the relative prices (or economies of scale) faced by people living in urban rather than rural areas.

sample selection, and I present the estimation equation for the divorced sample, accounting explicitly for selection. In Section 1.6.3 and Section 1.6.4 I discuss the estimation of the divorce probability that is included as part of the covariates in the married equation and the instrumental variable strategy for own and spousal wages. Finally, I present the final empirical specification, which integrates all of these considerations, in Section 1.6.5. For the estimation of Equation 1.11, I use a structural probit which allows me to account for the fact that wages are only observed for individuals who work. As for the divorced sample, whose labour supply is determined by Equation 1.12, I use a similar strategy, but allow for labour supply and marital status to be jointly determined, by using a bivariate probit.

1.6.1 Marital status and divorce probability

So far, I have discussed the labour supply of married and divorced women separately, without specifying how they transition from one group to the other. Doing so requires assumptions about the timing of decisions. Suppose that married couples enter period t with assets $A_{i,t}$ (and marginal utility of wealth $\nu_{i,t}$), determined in period $t-1$. At the beginning of the period, uncertainty about wages and prices, including economies of scale, is resolved. Based on this new information, couples decide whether to stay together or split. To do so, they evaluate the two corresponding problems, the married and the divorced problems presented in Section 1.3.

To formalize the divorce decision, denote $M_{i,t}$, an indicator variable equal to 1 if spouses are married. Furthermore, recall that $Q_{i,t}$ and $Q_{s,i,t}$, $s=1,2$, respectively denote the sets of choice variables of the married and the divorced households. Selection into the married or the divorced sample at time t is determined by

$$\begin{aligned} M_{i,t} = 1\big\{ & \lambda_{i,t}[U_{1,i,t}(Q^*_{i,t}) + \beta E(\omega_{i,t+1}V^M_{1,i}(A^*_{i,t+1}) + (1-\omega_{i,t+1})V^D_{1,i}(\Delta_{i,t+1}A^*_{i,t+1}))] \\ & + (1-\lambda_{i,t})[U_{2,i,t}(Q^*_{i,t}) + \beta E(\omega_{i,t+1}V^M_{2,i}(A^*_{i,t+1}) + (1-\omega_{i,t+1})V^D_{2,i}((1-\Delta_{i,t+1})A^*_{i,t+1}))] \\ & - U_{1,i,t}(Q^*_{1,i,t}) - \beta E V^D_{1,i}(A^*_{1,i,t+1})) - U_{2,i,t}(Q^*_{2,i,t}) - \beta E V^D_{2,i}(A^*_{2,i,t+1})) \geq 0 \big\} \end{aligned} \tag{1.34}$$

where $V^M_{s,i,t}$ is the value to the couple from remaining married and $V^D_{s,i,t}$ is the value to spouse s of being divorced. The factor $\Delta_{s,i,t}$, where $0 \leq \Delta_{s,i,t} \leq 1$ and $\Delta_{1,i,t} + \Delta_{2,i,t} \leq 1$, reflects the fact that a given spouse only retains a portion of the assets if the marriage breaks up. It may be a function of intra-household bargaining power, $\lambda_{i,t}$, or of factors beyond the control of the spouses such as the legal environment. Spouses must both be better off married for the couple to remain together. Note that the specification in Equation 1.34 allows for transfers between spouses to take place through $\lambda_{i,t}$, where $0 < \lambda_{i,t} < 1$. If one of the two spouses has a higher value of being divorced, than the household reallocates resources so that their utility of being married matches their utility of being divorced. If that is not possible, or if both spouses have a higher value of being divorced, then the marriage is dissolved.[17]

The spouses' net utilities from marriage are not directly observed. However, since they reflect

[17] Note that this is more flexible than transferable utility, which would require that both spouses get the same utility from a given amount of money.

expectations over optimal choices in the future, they are functions of expectations over future prices and incomes. I assume that $M_{i,t}$ can be approximated as follows:

$$
\begin{aligned}
M_{i,t} = 1\{\theta_0 + \theta^A A_{i,t} &+ \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} E_t \ln \mathrm{w}_{1,i,j} + \sum_{j=t}^{T} \theta_j^{\mathrm{w}_2} E_t \ln \mathrm{w}_{2,i,j} \\
&+ \sum_{j=t}^{T} \theta_j^{\lambda} E_t \ln \lambda_{i,j} + \sum_{j=t}^{T} \theta_j^{g_\Lambda} E_t \ln g_{\Lambda,i,j} + \sum_{j=t}^{T} \theta_j^{g_K} E_t \ln g_{K,i,j} + \iota_{i,t} \geq 0\}
\end{aligned} \tag{1.35}
$$

Similar to Section 1.5, I approximate for the expectation at time t over future values of each variable with a quadratic in time. That is,

$$E_t \ln \mathrm{w}_{1,i,j} = \pi_{0,i}^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} j + \pi_{2,i}^{\mathrm{w}_1} j^2 + \mathrm{v}_{i,j}^{\mathrm{w}_1 *} \tag{1.36}$$

$$E_t \ln \mathrm{w}_{2,i,j} = \pi_{0,i}^{\mathrm{w}_2} + \pi_{1,i}^{\mathrm{w}_2} j + \pi_{2,i}^{\mathrm{w}_2} j^2 + \mathrm{v}_{i,j}^{\mathrm{w}_2 *} \tag{1.37}$$

$$E_t \ln \lambda_{i,j} = \pi_{0,i}^{\lambda} + \pi_{1,i}^{\lambda} j + \pi_{2,i}^{\lambda} j^2 + \mathrm{v}_{i,j}^{\lambda *} \tag{1.38}$$

$$E_t \ln g_{\Lambda,i,j} = \pi_{0,i}^{g_\Lambda} + \pi_{1,i}^{g_\Lambda} j + \pi_{2,i}^{g_\Lambda} j^2 + \mathrm{v}_{i,j}^{g_\Lambda *} \tag{1.39}$$

$$E_t \ln g_{K,i,j} = \pi_{0,i}^{g_K} + \pi_{1,i}^{g_K} j + \pi_{2,i}^{g_K} j^2 + \mathrm{v}_{i,j}^{g_K *} \tag{1.40}$$

To see the implication of using an approximation of expectations formulated at time t, consider the case of own wage

$$
\begin{aligned}
\sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} E_t \ln \mathrm{w}_{1,i,j} &= \pi_{0,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} j + \pi_{2,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} j^2 + \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} \mathrm{v}_{i,j}^{\mathrm{w}_1 *} \\
&= \pi_{0,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} j + \pi_{2,i}^{\mathrm{w}_1} \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} j^2 + \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} \mathrm{v}_{i,j}^{\mathrm{w}_1 *} \\
&= \pi_{0,i}^{\mathrm{w}_1} \theta_0^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} \left[t \theta_0^{\mathrm{w}_1} + \theta_1^{\mathrm{w}_1} \right] + \pi_{2,i}^{\mathrm{w}_1} \left[t^2 \theta_0^{\mathrm{w}_1} + 2t \theta_1^{\mathrm{w}_1} + \theta_2^{\mathrm{w}_1} \right] + \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} \mathrm{v}_{i,j}^{\mathrm{w}_1 *}
\end{aligned}
$$

where $\theta_0^{\mathrm{w}_1} = \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1}$, $\theta_1^{\mathrm{w}_1} = \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} (j-t)$, and $\theta_2^{\mathrm{w}_1} = \sum_{j=t}^{T} \theta_j^{\mathrm{w}_1} (j-t)^2$. Then, Equation 1.35 can be expressed as follows

$$M_{i,t} = 1\{\theta_0 + \theta^A A_{i,t} + \hat{\Pi}_{0,i}^{M\prime} \Theta_0 + \hat{\Pi}_{1,i}^{M\prime} \Theta_1 + \hat{\Pi}_{2,i}^{M\prime} \Theta_2 + \iota_{i,t} \geq 0\}$$

where

$$\hat{\Pi}_{0,i}^{M} = (\hat{\pi}_{0,i}^{\mathrm{w}_1}, \hat{\pi}_{1,i}^{\mathrm{w}_1}, \hat{\pi}_{2,i}^{\mathrm{w}_1}, \hat{\pi}_{0,i}^{\mathrm{w}_2}, \hat{\pi}_{1,i}^{\mathrm{w}_2}, \hat{\pi}_{2,i}^{\mathrm{w}_2}, \hat{\pi}_{0,i}^{\lambda}, \hat{\pi}_{1,i}^{\lambda}, \hat{\pi}_{2,i}^{\lambda}, \hat{\pi}_{0,i}^{g_\Lambda}, \hat{\pi}_{1,i}^{g_\Lambda}, \hat{\pi}_{2,i}^{g_\Lambda}, \hat{\pi}_{0,i}^{g_K}, \hat{\pi}_{1,i}^{g_K}, \hat{\pi}_{2,i}^{g_K})',$$

$$\hat{\Pi}^M_{1,i} = (\hat{\pi}^{w_1}_{1,i} \times t, \hat{\pi}^{w_1}_{2,i} \times t, \hat{\pi}^{w_2}_{1,i} \times t, \hat{\pi}^{w_2}_{2,i} \times t, \hat{\pi}^{\lambda}_{1,i} \times t, \hat{\pi}^{\lambda}_{2,i} \times t, \hat{\pi}^{g_\Lambda}_{1,i} \times t, \hat{\pi}^{g_\Lambda}_{2,i} \times t, \hat{\pi}^{g_K}_{1,i} \times t, \hat{\pi}^{g_K}_{2,i} \times t)',$$

$$\hat{\Pi}^M_{2,i} = (\hat{\pi}^{w_1}_{2,i} \times t^2, \hat{\pi}^{w_2}_{2,i} \times t^2, \hat{\pi}^{\lambda}_{2,i} \times t^2, \hat{\pi}^{g_\Lambda}_{2,i} \times t^2, \hat{\pi}^{g_K}_{2,i} \times t^2)',$$

$$\Theta_0 = (\theta^{w_1}_0, \theta^{w_1}_1, \theta^{w_1}_2, \theta^{w_2}_0, \theta^{w_2}_1, \theta^{w_2}_2, \theta^{\lambda}_0, \theta^{\lambda}_1, \theta^{\lambda}_2, \theta^{g_\Lambda}_0, \theta^{g_\Lambda}_1, \theta^{g_\Lambda}_2, \theta^{g_K}_0, \theta^{g_K}_1, \theta^{g_K}_2)',$$

$$\theta^h_0 = \sum_{j=t}^{T} \theta^h_j, \quad h \in \{w_1, w_2, \lambda, g_\Lambda, g_K\},$$

$$\theta^h_1 = \sum_{j=t}^{T} \theta^h_j (j-t), \quad h \in \{w_1, w_2, \lambda, g_\Lambda, g_K\},$$

$$\theta^h_2 = \sum_{j=t}^{T} \theta^h_j (j-t)^2, \quad h \in \{w_1, w_2, \lambda, g_\Lambda, g_K\},$$

$$\Theta_1 = (\theta^{w_1}_0, \theta^{w_1}_1, \theta^{w_2}_0, \theta^{w_2}_1, \theta^{\lambda}_0, \theta^{\lambda}_1, \theta^{g_\Lambda}_0, \theta^{g_\Lambda}_1, \theta^{g_K}_0, \theta^{g_K}_1)',$$

$$\Theta_2 = (\theta^{w_1}_0, \theta^{w_2}_0, \theta^{\lambda}_0, \theta^{g_\Lambda}_0, \theta^{g_K}_0)',$$

Also, because marital status is determined at the couple level, $M_{1,i,t} = M_{2,i,t} = M_{i,t}$. Having established that, I first turn to the two issues that pertain to the endogeneity of marital status: sample selection and the endogeneity of $\ln \Omega_{i,t}$ in Equation 1.11.

1.6.2 Endogenous switching (sample selection)

The economic model suggests that the effect of divorce on the labour supply of married individuals should be estimated using Equation 1.11, and the effect of divorce on divorced individuals should be estimated using Equation 1.12. However, at any point in time, the decision of spouses to stay together depends on their joint gains from marriage. Therefore, the samples used to estimate Equation 1.11 and Equation 1.12 may not be randomly selected; that is, there is endogenous switching between the regime described by Equation 1.11 and that described by Equation 1.12. To see this, consider first the error term in the equation for the labour supply of married women

$$E(u^M_{1,i,t}|M_{1,i,t} = 1) = E(\eta_{l,1,i,t} + \eta_{L,1,i,t}|M_{1,i,t} = 1) + E(\sum_{j=1}^{t} \varepsilon_{i,j}|M_{1,i,t} = 1)$$

I assume that period-specific shocks to taste for leisure are independent from marital status. For instance, this implies that a person's taste for activities that can be performed as a group, such as hiking, is not impacted by their marital status.[18] This leaves the second term, $E(\sum_{j=1}^{t} \varepsilon_{i,j}|M_{1,i,t} = 1)$. Marital status at time t is a function of assets at time t, $A_{i,t}$, and expectations over future profiles. Therefore,

[18]This does not mean that marital status does not impact choices over leisure. Recall that the quality of time parameter, $g_{\Lambda,i,t}$, enters in the determination of married women's labour supply, Equation 1.11. For instance, this allows for the case where spouses who enjoy spending time together may spend more time hiking than if they were single, because the public leisure is relatively cheaper as a couple.

there is reason to be concerned that the series of shocks that have led to period t, $\sum_{j=1}^{t} \varepsilon_{i,j}$, may be correlated to marital status through the asset level. In turn, this would introduce bias through the correlation between $A_{i,t}$ and elements of $\ln v_{i,t}$ that are part of the right-hand-side variables included in Equation 1.11.

In this context, the probability of divorce enters in the determination of labour force participation through two channels: first, as an endogenous regressor in the equation for hours worked, Equation 1.11, and second, through the necessary selection correction. However, the two are dealt with empirically in very similar ways. To address the first problem, we would estimate the probability of divorce, including a regressor which we have grounds to believe does not belong in the labour supply equation, then use that estimated probability as the right-hand-side variable in Equation 1.11. To address the second problem, we would estimate a bivariate probit for labour force participation and marital status, again including an excluded regressor. If we include both corrections in the labour supply equation, the identification of β_{Ω}^{M} relies on the assumption that divorce probability enters in Equation 1.11 linearly, and that any non-linear effect only operates through selection. To avoid making this assumption, I use a subsample of married women for whom selection is less likely to drive differences in labour force participation; namely, I use the observations for married women, up to three years before the end of the observation period. In other words, I take advantage of the panel nature of my data, and use variation over time within a sample where composition does not change. Although the resulting sample is not randomly selected, any change in hours that is observed cannot come from selection. Indeed, we are concerned that $E(\sum_{j=1}^{t} \varepsilon_{i,j} | M_{1,i,t} = 1) \neq 0$ because of the dependency of $M_{1,i,t}$ on $A_{i,t}$. By restricting the sample as described, the requirement is that $E(\sum_{j=1}^{t} \varepsilon_{i,j} | M_{1,i,t+3} = 1) = 0$. Obviously, this is not a random sample, but the three-year buffer period helps loosen the link between the series of shocks and the asset level.

Next, consider the error in the equation for the labour supply of divorced women

$$E(u_{1,i,t}^{D} | M_{1,i,t} = 0) = E(\eta_{l,1,i,t} + \eta_{L,1,i,t} | M_{1,i,t} = 0) + E(\sum_{j=1}^{t^d-1} \varepsilon_{i,j} | M_{1,i,t} = 0) + E(\sum_{j=t^d}^{t} \varepsilon_{1,i,j} | M_{1,i,t} = 0)$$

As before, I assume that $E(\eta_{l,1,i,t} + \eta_{L,1,i,t} | M_{1,i,t} = 1) = 0$. For now, I treat divorce as an absorbing state, so I ignore the third term. The concern with respect to divorced women is that in the early years of divorce, the second term is non-zero; for instance, that a series of negative shocks may lead to divorce. Although I could use a similar approach as I do for the married sample, and ignore the earlier years of divorce, these are important to understand the effect at the time of dissolution. Therefore, I explicitly correct for selection for the divorced group. Given that both the outcome and the selection term are binary, I address selection through a bivariate probit model instead of specifying the typical Heckman selection model for continuous dependent variables. Going back to Equation 1.12, recall that tax data does not provide information on hours worked. For this reason, I focus on labour market participation, LFP, as the outcome of interest.

$$LFP^D_{1,i,t} = \begin{cases} 1 & \text{if } N^D_{1,i,t} > 0 \ \ (\text{or } \ln(1-N^D_{1,i,t}) < \ln(1)) \\ 0 & \text{if } N^D_{1,i,t} \leq 0 \ \ (\text{or } \ln(1-N^D_{1,i,t}) \geq \ln(1)) \end{cases}$$

Assuming that the errors on the labour supply of divorced individuals are normally distributed, the economic model and the discussion so far suggests that there are three types of contributions to be considered in the estimation of Equation 1.12

$$\begin{aligned} &Pr(M_{1,i,t}=1) \\ &Pr(M_{1,i,t}=0, LFP^D_{1,i,t}=0) \\ &Pr(M_{1,i,t}=0, LFP^D_{1,i,t}=1) \end{aligned}$$

1.6.3 Endogenous regressors: divorce probability

Divorce probability appears in Equation 1.11 through $\ln\Omega_{i,t}$; that is, $\ln\Omega_{i,t} = \ln[f(Pr\{M_{i,t}=0\})]$. As such, it may be endogenous for two reasons: there may be unobserved determinants common to both labour supply and marital status, and labour supply itself may be a relevant determinant of divorce probability. I consider only the first case here, as the second one constitutes a fairly natural extension. The relevant estimation equations are

$$\begin{aligned} LFP^M_{1,i,t} = 1\{&Z_{i,0}\phi_0 + \alpha^A_0\hat{A}_{i,0} + \hat{\Pi}'_{0,i}D + b^M t + \beta^M_\Omega \ln[f(Pr\{M_{i,t}=0\})] + \beta^M_N \ln\bar{N}_{1,i,t} \\ &+ \beta_\lambda \ln\lambda_{i,t} + \beta_g \ln g_{\Lambda,i,t} + \beta^M_1 \ln \mathrm{w}_{1,i,t} + \beta_2 \ln \mathrm{w}_{2,i,t} + u^M_{1,i,t} + a^*_{i,0} < 0\} \end{aligned} \tag{1.41}$$

and

$$M_{i,t} = 1\{\theta_0 + \theta^A A_{i,t} + \hat{\Pi}^{M'}_{0,i}\Theta_0 + \hat{\Pi}^{M'}_{1,i}\Theta_1 + \hat{\Pi}^{M'}_{2,i}\Theta_2 + \iota_{i,t} \geq 0\} \tag{1.42}$$

Recall that $u^M_{1,i,t} = \eta_{l,1,i,t} + \eta_{L,1,i,t} + \sum_{j=0}^{t} \varepsilon_{i,j}$, and $a^*_{i,0}$ is a time-invariant preference shifter. Therefore, the two most likely causes of endogeneity come through the role of $A_{i,t}$ in determining $M_{i,t}$, and through shocks to contemporary variables, such as wages, that also impact the probability of divorce. Indeed, household assets at time t, $A_{i,t}$, are in part determined by $\sum_{j=1}^{t-1} \varepsilon_{i,j}$, the sum of unexpected one-period shocks to wages, prices, etc. Therefore, the error term may be correlated with divorce probability, through the role of assets in determining the latter. Alternatively, it may be that unobserved shocks to labour supply (eg. health shocks) also constitute shocks to the probability of remaining married.

I only observe whether women are married, $M_{i,t}$, not their probability of divorce, $Pr\{M_{i,t}=0\}$. Therefore, I first estimate Equation 1.42 by probit. This yields the predicted probability of divorce, $(1-\hat{\omega}_{i,t}) = (1-Pr\widehat{\{M_{i,t}=1\}})$, which I use to construct $\hat{\Omega}_{i,t}$. This is akin to a standard 2SLS procedure, albeit with two partially observed variable. The method used corresponds closely to the procedure outlined in Maddala (1983, Chapter 8), for models with mixed qualitative, truncated and censored structures.[19] In this context, the effect of divorce risk is identified by the presence of regressors in Equation 1.42 that do not appear in Equation 1.41. From Section 1.6.1, there are two sources of exogenous variation in divorce risk: first, the relative prices of public and private goods; and second, the time-varying impact of the slope and curvature parameters that characterize the various profiles that affect labour supply.[20] Recall that the former are meant to capture economies of scale in consumption, while the latter comes from the spouses' re-evaluation of their (remaining) marital surplus every period.

1.6.4 Endogenous regressors: own and spousal wages

As is typical in models of labour supply, there is concern that

$$E(u^M_{1,i,t}|\ln \mathrm{w}_{1,i,t}) \neq 0$$
$$E(u^M_{1,i,t}|\ln \mathrm{w}_{2,i,t}) \neq 0$$

and

$$E(u^D_{1,i,t}|\ln \mathrm{w}_{1,i,t}) \neq 0$$

I instrument for the natural logarithm of own and spousal wages with education. Specifically, I use the following reduced-form expressions for own and spousal wages

[19] Specifically, this corresponds to Model 6 in Section 8.8. The model outlined by Maddala (1983) is

$$y_1^* = \gamma_1 y_2^* + \beta_1 X_1 + u_1$$
$$y_2^* = \gamma_2 y_1^* + \beta_2 X_2 + u_2$$

where only $y_1 = 1\{y_1^* > 0\}$ and $y_2 1\{y_2^* > 0$ are observed. The estimation procedure first estimates probits of each reduced form equation, and uses the resulting $\hat{y}_1^*$ and $\hat{y}_2^*$ in the structural equations by probit.

In the present case, the corresponding equations would be

$$y_1^* = \gamma_1 P(y_2^* > 0) + \beta_1 X_1 + u_1$$
$$y_2^* = \beta_2 X_2 + u_2$$

which comes down to the same estimation strategy. That is, there is a (scaled) correspondence between net gains from marriage and the probability of divorce. In other words, $\hat{M}^*_{i,t} = Pr(\widehat{M^*_{i,t} = 1})$.

[20] Equation 1.41 includes the following regressors: the estimated value of assets at time of marriage; the estimated intercept, slope and curvature parameters that characterize the own wage, spousal wage, quality of time proxy, and divorce risk profiles; time since marriage; the logarithms of own wage, spousal wage and the quality of time proxy; and the logarithm of estimated divorce risk, and corresponding interactions with the intra-household bargaining power proxy and the couple-type dummies meant to control for whether or not spouses would be worse off after divorce. As for Equation 1.42, the regressors include: the estimated value of assets at time t; the estimated intercept, slope and curvature parameters that characterize the profiles of own wage, spousal wage, the quality of time proxy, the economies of scale in consumption proxy and the intra-household bargaining power proxy; interactions between the slope and curvature parameters and time since marriage; and interactions between the curvature parameters and time since marriage squared.

$$\ln \mathrm{w}_{1,i,t} = Z'_{1,i,t}\zeta_1 + \mu_{1,i,t} \tag{1.43}$$

$$\ln \mathrm{w}_{2,i,t} = Z'_{2,i,t}\zeta_2 + \mu_{2,i,t} \tag{1.44}$$

where $Z_{1,i,t}$ is a vector of age, age squared, education and interactions between education and the quadratic in age, and $Z_{2,i,t}$ is a similar vector, where spouse's age is used instead of the respondent's age. For education to be an appropriate instrument in the context that I'm interested in, it must affect labour supply only through wages; i.e., only by affecting spouses' productivity on the labour market. This is violated, for instance, if education also alters output in home production. A potentially important violation arises therefore with respect to child care. Although the literature on child development is still disentangling the various forces at play, there is evidence of higher returns to the time spent by educated parents on child care (eg. Harding et al., 2015; Kalil et al., 2012). I return to this possibility later on.

1.6.5 Final procedure

For simplicity, re-write

$$\begin{aligned} Pr(M_{i,t}=1) &= Pr\{\iota_{i,t} > -(\theta_t^0 + \theta_t^A A_{i,t} + \hat{\Pi}'_{0,i}\Theta_{i,0} + \hat{\Pi}'_{1,i}\Theta_{i,1} + \hat{\Pi}'_{2,i}\Theta_{i,2})\} \\ &= Pr(\iota_{i,t} > -\Psi_{M,1,i,t}) \end{aligned}$$

and

$$\begin{aligned} Pr(LFP^M_{1,i,t}=1) &= Pr\{u^M_{1,i,t} < \ln(1) - (Z_i\phi_0 + \alpha_0^A\hat{A}_{i,0} + \hat{\Pi}'_{0,i}D_{i,0} + b^M t + \beta_\Omega^M \ln\Omega_{i,t} \\ &\qquad + \beta_N^M \ln\bar{N}_{1,i,t} + \beta_\lambda \ln\lambda_{i,t} + \beta_g \ln g_{\Lambda,i,t} + \beta_1^M \ln \mathrm{w}_{1,i,t} + \beta_2 \ln \mathrm{w}_{2,i,t})\} \\ &= Pr(u^M_{1,i,t} < -\Psi_{LFPm,1,i,t}) \\ Pr(LFP^D_{1,i,t}=1) &= Pr\{u^D_{1,i,t} < \ln(1) - (Z_i\phi_0 + \alpha_0^A\hat{A}_{i,0} + \hat{\Pi}'_{0,i}D_{i,0} + b^M(t^d-1) + \sum_{t=1}^{t^d-1}\beta_\Omega^M \ln\Omega_{i,t} \\ &\qquad + b^d(t-t^d) + \beta_1^D \ln \mathrm{w}_{1,i,t})\} \\ &= Pr(u^D_{1,i,t} < -\Psi_{LFPd,1,i,t}) \end{aligned}$$

Denote $\prod_{M,LFP}$, the product over people belong to the group with marital status $M \in \{0,1\}$ and labour force status $LFP \in \{0,1\}$. Then, the likelihood function for the married sample is

$$
\begin{aligned}
L^M &= \prod_{1,0} Pr(LFP^M_{1,i,t} = 0) \times \prod_{1,1} Pr(LFP^M_{1,i,t} = 1) \\
&= \prod_{1,0} \Phi(\Psi_{LFPm,1,i,t}) \times \prod_{1,1} \Phi(-\Psi_{LFPm,1,i,t})
\end{aligned} \tag{1.45}
$$

and for the divorced sample,

$$
\begin{aligned}
L^D &= \prod_{1,.} Pr(M_{1,i,t} = 1) \times \prod_{0,0} Pr(M_{1,i,t} = 0, LFP^D_{1,i,t} = 0) \times \prod_{0,1} Pr(M_{1,i,t} = 0, LFP^D_{1,i,t} = 1) \\
&= \prod_{1,.} \Phi(\Psi_{M,1,i,t}) \times \prod_{0,0} \Phi_2(-\Psi_{M,1,i,t}, \Psi_{LFPd,1,i,t}; \rho_d) \times \prod_{0,1} \Phi_2(-\Psi_{M,1,i,t}, -\Psi_{LFPd,1,i,t}; \rho_d)
\end{aligned} \tag{1.46}
$$

Structural probit

Both $LFP^M_{1,i,t}$ and $LFP^D_{1,i,t}$ depend on $\ln \mathrm{w}_{1,i,t}$ and $\ln \mathrm{w}_{2,i,t}$, which are only observed for people who are employed. Therefore, wages need to be constructed for individuals who are not participating in the labour market, in a way that accounts for selection; i.e., it must acknowledge that people who are not participating may not have had the same wages as people who are participating and are observably similar. Furthermore, this selection correction must also allow for selection into the labour force to operate differently for married and divorced individuals. To do so, I use a structural probit approach, described as follows:

1. Reduced-form estimation of the probit representing married women's labour force participation, and the bivariate probit representing the marital status and labour force participation choices of divorced women.

   Estimate Equation 1.45 and Equation 1.46, replacing $\ln \mathrm{w}_{s,i,t}$, $s = 1,2$, in the labour force participation regressors by their reduced-form expressions. This yields estimates of combinations of the structural parameters of interest and the reduced-form parameters of the wage equations. These are then used to construct estimates of the sample selection corrections to be applied to the wages of married and divorced individuals, respectively $\hat{\xi}^M_{s,i,t}$ and $\hat{\xi}^D_{s,i,t}$.

2. Selectivity-corrected reduced-form estimation of the wage equations.

   Estimate the reduced-form equations for wages, including the labour force selection correction,

$$
\begin{aligned}
\ln \mathrm{w}^M_{s,i,t} &= Z'_{1,i,t}\zeta_1 + \hat{\xi}^M_{s,i,t} + \mu^M_{1,i,t} \\
\ln \mathrm{w}^D_{s,i,t} &= Z'_{1,i,t}\zeta_1 + \hat{\xi}^D_{s,i,t} + \mu^D_{1,i,t}
\end{aligned}
$$

Using the resulting estimates, compute the selectivity-corrected wages for all individuals,including those that are not in the labour force

$$\ln \tilde{w}^{M}_{s,i,t} = Z'_{s,i,t}\hat{\zeta}_{s,i,t}$$
$$\ln \tilde{w}^{D}_{s,i,t} = Z'_{s,i,t}\hat{\zeta}_{s,i,t}$$

where wages have been rid of the effect of labour force participation selection.

3. Estimation of the structural probit and bivariate probit, using the corrected wages, Equation 1.45 and Equation 1.46, which yields the structural coefficients of interest.

1.7 Results

1.7.1 Effect on married women

Table 1.1 presents the estimated results for married women, corresponding to Equation 1.45. Panel A shows the effect on labour supply of the different profile parameters that contribute to $v_{i,0}$. Panel B shows the baseline updating to the marginal utility of wealth, b^M, and the effect of the contemporary divorce risk through the various components of $\Omega_{i,t}$. Finally, Panel C shows the effect of contemporary values of own and spousal wages, intra-household bargaining power, and economies of scale in leisure (quality of time). The life-cycle model of labour supply informs three elasticities that characterize the response of labour supply to changes in own and spousal wages, economies of scale in leisure, and divorce probability: the intertemporal substitution elasticity, the own-uncompensated elasticity and the cross-uncompensated elasticity.[21] These effects are summarized in Table 1.2 for own wages, spousal wages and quality of time, and in Table 1.3 for divorce probability.

Before discussing the response of married women's labour supply to changes in divorce risk, consider their response to changes in own and spousal wages. This constitutes a useful benchmark to think about the results for divorce risk. The intertemporal substitution elasticity with respect to own wage reflects the labour supply response of individual women as they move along their lifetime wage profile. A 10% wage increase along one's profile results in a 6.94% increase in hours worked; as expected, women shift work hours to periods where these are better remunerated. Similarly, as a woman's husband moves along his wage profile, a 10% increase in his wage leads to a 1.76% increase in her hours worked. This is also in line with theory: an increase in the husband's wage increases the cost of public leisure, which increases labour supply. Next, the own-uncompensated elasticity measures the response of hours worked at time t to a change in wages at time t. Such a wage change affects labour supply through a corresponding change in the lifetime wage profile, and through intertemporal substitution.

[21] In this paper, recall that the parameters discussed for married women are identified subject to the normalization that $Var(u^M_{1,i,t}) = 1$, because I do not observe hours but only labour force participation.

Table 1.1: Effect of divorce risk on married women

		Coefficient	(p-value)
PANEL A			
Assets (t=0)	π_0^y	-0.000	(0.016)
Own wage	$\pi_0^{w_1}$	-0.148	(0.005)
	$\pi_1^{w_1}$	-3.007	(0.001)
	$\pi_2^{w_1}$	-74.500	(0.000)
Spousal wage	$\pi_0^{w_2}$	-0.009	(0.717)
	$\pi_1^{w_2}$	-0.603	(0.222)
	$\pi_2^{w_2}$	-16.002	(0.150)
Prob. divorce	π_0^{ω}	0.646	(0.001)
	π_1^{ω}	20.737	(0.001)
	π_2^{ω}	732.215	(0.000)
Quality of time	$\pi_0^{g_\Lambda}$	-0.291	(0.000)
	$\pi_1^{g_\Lambda}$	-1.943	(0.006)
PANEL B			
Time since marriage		-0.012	(0.000)
ϕ_2 (h. better; w. worse)		0.009	(0.893)
ϕ_3 (h. better; w. better)		0.259	(0.004)
ϕ_4 (h. worse; w. better)		0.186	(0.014)
$\phi_2 \times \ln\lambda$		1.799	(0.332)
$\phi_3 \times \ln\lambda$		0.636	(0.808)
$\phi_4 \times \ln\lambda$		-0.830	(0.672)
$\ln(1-\omega)$		-0.212	(0.000)
$\phi_2 \times \ln(1-\omega)$		0.061	(0.083)
$\phi_3 \times \ln(1-\omega)$		0.154	(0.001)
$\phi_4 \times \ln(1-\omega)$		0.072	(0.046)
$\ln(1-\omega) \times \ln\lambda$		-0.003	(0.997)
$\phi_2 \times \ln(1-\omega) \times \ln\lambda$		0.992	(0.842)
$\phi_3 \times \ln(1-\omega) \times \ln\lambda$		0.268	(0.842)
$\phi_4 \times \ln(1-\omega) \times \ln\lambda$		-1.523	(0.144)
PANEL C			
$\ln w_1$		0.694	(0.000)
$\ln w_2$		0.176	(0.000)
$\ln\lambda$		0.467	(0.758)
$\ln g_\Lambda$		0.507	(0.000)
Constant		-6.001	(0.000)
N		4,800	

Table 1.2: Summary of effects on married women: own wage, spousal wage, and quality of time

	Own wage	Spousal wage	Quality of time
Effect of 10% ↑ at time t			
Intertemp.	6.94%	1.76%	5.07%
Own-unc.	6.90%	1.76%	5.00%
Cross-unc.	−0.04%	0.00%	−0.07%
Effect of 10% ↑ in profile parameters			
$\pi_{0,i}$	5.46%	1.76%	2.16%
$\pi_{1,i}$	$(-3.007+0.694t)\times 10\%$ No effect at age 29	$(0.176t)\times 10\%$	$(-1.943+0.507t)\times 10\%$ No effect at age 29
$\pi_{2,i}$	$(-74.500+0.694t^2)\times 10\%$ No effect at age 35	$(0.176t^2)\times 10\%$	$(0.507t^2)\times 10$

For own wages, it is equal to $\hat{\beta}_1^M + \hat{\alpha}^{w_1}(t)$, approximately 6.90%.[22] The change is mostly driven by intertemporal substitution. Correspondingly, the cross-uncompensated elasticity with respect to own wages is very small. A 10% increase in own wages in period t reduces hours worked in other periods by only 0.04%. As for spousal wage, Table 1.1 shows that shifts in the lifetime profile do not affect the labour supply of women significantly. Therefore, the own-uncompensated elasticity with respect to spousal wage is 1.76%, same as the intertemporal substitution elasticity, and the cross-effect is nil.

The estimates from Table 1.1 can also be used to construct estimates of the effect of anticipated shifts in the lifetime profiles of own and spousal wages. Results are presented in the bottom half of Table 1.2. A 10% increase in $\hat{\pi}_{0,i}^{w_1}$ represents an upward parallel shift of the own wage profile, and results in a change in hours worked of $(\hat{\alpha}_0^{w_1} + \hat{\beta}_1^M)\%$, or 5.46%. A similar shift in the spousal wage profile increases the wife's hours worked by 1.76%. As for changes in the slope and curvature parameters, their effect depends on the point in time that is considered. A 10% increase in the slope of the own wage profile decreases hours worked until age 29, while a similar change in the husband's wage profile increases hours worked at all ages. Finally, a 10% increase in the curvature of the own wage profile decreases hours worked until age 35, and the corresponding change in the husband's wage profile increases women's labour supply at all ages.

The effect of changes in the quality of time spouses spend together is presented in the last column of Table 1.2. The data shows that the profile is fairly linear, so I omit the curvature parameter from the estimation equation. As expected, the coefficients on the profile parameters are negative. Quality of time reduces the cost of public leisure, which has a positive effect on lifetime income, reducing labour supply in all periods. Furthermore, the intertemporal substitution elasticity is positive. As women move along their profile to periods where time spent with their spouse is relatively more enjoyable, they can buy less public leisure and enjoy the same private-leisure equivalent. Thus, movements to higher-

[22] In practice, $\hat{\alpha}^{w_1}(t)$ is not observed, but the empirical estimation yields $\hat{\alpha}_0^{w_1} = \sum_{t=0}^{T} \hat{\alpha}_0^{w_1}(t)$. I follow MaCurdy (1981) and use the average of $\hat{\alpha}_0^{w_1}$ over 40 years, the assumed length of their career.

quality-of-time periods lead to increases in own labour supply. Specifically, a 10% increase along their quality of time profile leads to a 5.07% increase in hours worked. Own- and cross-uncompensated effects for a 10% increase in quality of time are respectively 5.00% and −0.07%. Notably, the cross-effect is much more important for quality of time than for own or spousal wages. Finally, a 10% upward shift in the profile increases labour supply by 2.16% in all periods, and an increase in the slope parameter decreases hours worked until age 29.

Unlike changes to the spousal wage profile, shifts in the quality of time profile have economically and statistically significant effects on the labour supply of married women. The evidence I present here suggests that it is not the husband's lifetime wages that shift women's lifetime income, but the quality of the relationship with said husband. Indeed, women intertemporally coordinate their work with their husband's, but they don't work less over their lifetime because their husband has a higher wage profile. Conversely, they do increase their hours worked every period if their husband contributes positively to the quality of their leisure time. If higher earning husbands are also better company, then failing to control for quality of time would mistakenly lead us to conclude that the husband's wage profile has a meaningful effect on women's lifetime income.

In Table 1.3, I present estimates of the effect of changes in the divorce risk at time t on women's labour supply, and of the effect of changes in the divorce risk profile parameters.[23] Recall that the effect of divorce risk on married women operates entirely through the couple's marginal utility of wealth. Expected lifetime divorce risk decreases lifetime income, and therefore increases women's labour supply in all periods. This is evident from the positive estimated coefficients on $\hat{\pi}^{\omega}_{0,i}$, $\hat{\pi}^{\omega}_{1,i}$ and $\hat{\pi}^{\omega}_{2,i}$ in Panel A of Table 1.1. As women progress through their marriage, they move along their divorce risk profile. Hence, any period in which they remain married is also a period where there was a non-zero risk of divorce. In this context, theory predicts that women will readjust their labour supply as they move along their divorce risk profile and remain married, to reflect the fact that the event against which they were insuring did not occur. Note that higher portions of the divorce risk profile contribute more to the initial increase in labour supply; i.e., to the effect captured by $\hat{\pi}^{\omega}_{0,i}$, $\hat{\pi}^{\omega}_{1,i}$ and $\hat{\pi}^{\omega}_{2,i}$. Therefore, as women move to higher portions of their divorce risk profile, they are expected to decrease their labour supply.

The theoretical model predicts that contemporary divorce risk has a different effect on couples where both spouses would face a tighter budget constraint if divorced, where only the wife would face a tighter budget constraint, where both spouses would face a looser budget constraint if divorced, and where only the husband would face a tighter budget constraint. I label each couple type respectively as $\phi = 1$, $\phi = 2$, $\phi = 3$ and $\phi = 4$. In the empirical specification, this is accounted for by interactions between contemporary divorce risk, $\ln(1 - \omega_{i,t})$, and the type indicator. The other terms in Panel B of Table 1.1 correspond to additional terms that contribute to $\ln \Omega_{i,t}$.

[23] I abuse nomenclature slightly in Table 1.3, and refer to these effects as intertemporal substitution, own-uncompensated and cross-uncompensated elasticities, for the sake of comparison with the effect of changes in own wages, spousal wages and economies of scale in leisure at time t.

Table 1.3: Summary of effects on married women: divorce probability

	Divorce probability, $\phi = 1$	Divorce probability, $\phi = 2$
Effect of 10% ↑ in divorce risk at time t		
Intertemp.	-2.12%	-1.51%
Own-unc.	-1.96%	-1.35%
Cross-unc.	0.16%	0.16%
Effect of 10% ↑ in profile parameters		
$\pi_{0,i}$	4.34%	4.95%
$\pi_{1,i}$	$(20.737 - 0.212t) \times 10\%$	$(20.737 - 0.151t) \times 10\%$
$\pi_{2,i}$	$(732.215 - 0.212t^2) \times 10\%$	$(732.215 - 0.151t^2) \times 10\%$
	Divorce probability, $\phi = 3$	Divorce probability, $\phi = 4$
Effect of 10% ↑ in divorce risk at time t		
Intertemp.	-0.58%	-1.40%
Own-unc.	-0.42%	-1.24%
Cross-unc.	0.16%	0.16%
Effect of 10% ↑ in profile parameters		
$\pi_{0,i}$	5.88%	5.06%
$\pi_{1,i}$	$(20.737 - 0.058t) \times 10\%$	$(20.737 - 0.140t) \times 10\%$
$\pi_{2,i}$	$(732.215 - 0.058t^2) \times 10\%$	$(732.215 - 0.140t^2) \times 10\%$

Consider first the effect of changes in the parameters that characterize the lifetime divorce risk of women. In Table 1.3, I present estimates for each couple type separately. For women in the most vulnerable couples ($\phi = 1$), an upward shift in the divorce risk profile by 10% increases labour supply by 4.34% in all periods. For couples who would be financially better off if divorced, labour supply increases by 5.88%. The difference stems from the smaller re-adjustment during the married years. Recall from Section 1.3 that the size of the re-adjustment follows from the difference between the tightness of the budget constraints spouses would respectively face if divorced, compared to their situation when married. Hence, this result suggests there is a larger gap for more vulnerable couples, which makes sense. Interestingly, couples where either the husband or the wife would face a looser budget constraint if divorced are similar to each other. A 10% parallel shift in their divorce risk profile increases labour supply for women by 4.95% and 5.06% respectively. In all cases, the effect is similar in magnitude to the effect of a 10% downward shift in the own wage profile, but the comparison hides differences in the underlying drivers of the effect. For own wages, movements along the profile play a larger role than shifts in the profile. For divorce risk, shifts in the profile outweigh movements along the profile. As a result, a 10% upward shift in the divorce risk profile results in a similar increase in women's labour supply as a 10% increase their wage profile.

As for increases in the slope and curvature parameters of divorce risk profiles, they unambiguously increase hours worked throughout woman's marriages. This contrasts with the effect of corresponding changes in the own wage profile, which decrease labour supply in the early years of women's careers. Again, this stems from the fact that the effect of parametric shifts outweigh subsequent re-adjustments, for all couple types considered. In other words, differences across women in their divorce probability profiles have less ambiguous effects than differences in their wage profiles.

Next, consider the effect of a 10% increase in divorce risk at time t. Women are expected to re-adjust their hours worked downwards, as they move to higher points of their divorce risk profile. In vulnerable couples, I find that a 10% increase along a woman's divorce risk profile decreases hours worked by 2.12%. The corresponding effect is only approximately one quarter as large for women in couples where spouses are expected to face looser budget constraints if divorced. Furthermore, the own- and cross-uncompensated elasticities are respectively −1.96% and 0.16% for the first couple type, and −0.42% and 0.16% for the better-off group. Consistently with the greater importance of lifetime profile for divorce risk, the cross effect is substantially larger than for own wages.

Finally, a 10% increase in assets at the time of marriage decreases labour supply, but the effect is not economically significant. Moreover, the estimate for b^M, the coefficient on time since marriage, is negative. This term captures two things: first, the comparison between the rate of time preferences and the real rate of interest; and second, expectations over one-period shocks to the marginal utility of wealth. If at time $t-1$ spouses expect shocks in the following period that will increase their marginal utility of wealth, they make consumption, leisure and saving decisions at time $t-1$ that pre-emptively decrease their marginal utility of wealth; vice versa. The negative estimated coefficient on time since marriage is consistent with that scenario, as well as with a rate of time preference that is smaller on average than the real rate of interest.

1.7.2 Effect on divorced women

Table 1.4 shows the results for divorced women, corresponding to Equation 1.46. Again, I construct the effects of interest and present them in Table 1.5 and Table 1.6. Among divorced women, couple-specific variables such as spousal wage and divorce probability operate fully through the corresponding lifetime profile parameters. Before turning to those, consider the labour supply effect of changes in own wages. I find that only the slope and curvature parameters of the own wage profile affect the labour supply of divorced women. Therefore, the cross-uncompensated elasticity is non-existent, and the own-effect is equal to the intertemporal substitution elasticity. A movement along a woman's wage profile corresponding to a 10% increase leads to a 3.13% increase in hours worked, approximately half as much as among married women. Similarly, a 10% increase in $\hat{\pi}_{0,i}^{w_1}$ increases labour supply by 3.13% in every period. Finally, increases of the slope and curvature parameters decrease labour supply for a longer period of time than among married women.

Next, consider the effect of changes in the profile parameters of spousal wage and quality of time. Following the theoretical model, contemporary values were not included in the empirical specification for the labour supply of divorced women. Therefore, intertemporal substitution and own-uncompensated elasticity are irrelevant constructs for spousal wage and quality of time after divorce. Furthermore, note that the cross-uncompensated elasticity necessarily refers more specifically to the effect of a change in spousal wage or quality of time before dissolution on the labour supply of women

Table 1.4: Effect of divorce on divorced women

		Coefficient	(p-value)
PANEL A			
Assets (t=0)	π_0^y	-0.000	(0.000)
Own wage	$\pi_0^{w_1}$	-0.176	(0.151)
	$\pi_1^{w_1}$	-4.646	(0.028)
	$\pi_2^{w_1}$	-114.262	(0.008)
Spousal wage	$\pi_0^{w_2}$	-0.209	(0.002)
	$\pi_1^{w_2}$	-3.764	(0.004)
	$\pi_2^{w_2}$	-72.114	(0.010)
Prob. divorce	π_0^{ω}	1.480	(0.002)
	π_1^{ω}	56.441	(0.001)
	π_2^{ω}	1771.678	(0.000)
Quality of time	$\pi_0^{g_\Lambda}$	-0.479	(0.014)
	$\pi_1^{g_\Lambda}$	-5.938	(0.015)
Time of dissolution		-0.011	(0.000)
PANEL B			
Time since dissolution		-0.018	(0.000)
$\ln w_1$		0.313	(0.002)
Constant		5.460	(0.035)
N		1,500	

Table 1.5: Summary of effects on divorced women: own wage, spousal wage, and quality of time

	Own wage	Spousal wage	Quality of time
Effect of 10% ↑ at time t			
Intertemp.	3.13%	–	–
Own-unc.	3.13%	–	–
Cross-unc.	0.00%	−0.05%	−0.12%
Effect of 10% ↑ in profile parameters			
$\pi_{0,i}$	3.13%	−2.09%	−4.79%
$\pi_{1,i}$	$(-4.646+0.313t)\times 10\%$ No effect at age 40	−37.64%	−59.38%
$\pi_{2,i}$	$(-114.262+0.313t^2)\times 10\%$ No effect at age 44	−721.14%	–

Table 1.6: Summary of effects on divorced women: divorce probability

	Divorce probability
Effect of 10% ↑ in divorce risk at time t	
Intertemp.	–
Own-unc.	–
Cross-unc.	0.37%
Effect of 10% ↑ in profile parameters	
$\pi_{0,i}$	14.80%
$\pi_{1,i}$	564.41%
$\pi_{2,i}$	17,716.78%

after dissolution. I estimate that a 10% increase in spousal wage decreases women's hours worked by 0.05% after dissolution. In comparison, a similar change was found to have no effect before divorce. In other words, an increase in spousal wage while spouses are married has no cross-effect until after divorce. Similarly, the cross-uncompensated effect of a 10% increase in the quality of time spent with one's spouse is almost twice as large among divorced women as it is among married women. Therefore, in both cases divorce exacerbates pre-existing differences between women. That is, women who were in different marriages exhibit larger differences in labour supply after divorce than during their married years.[24] Furthermore, in the absence of intertemporal substitution, changes to the spousal wage and quality of time profiles all have unambiguously negative effects on labour supply after marital dissolution. For instance, divorced women whose spousal wage and quality of time profiles are uniformly higher by 10% work 2.09% and 4.79% fewer hours respectively.

Like spousal wage and quality of time, divorce risk only affects divorced women through its lifetime profile. The effect of an increase in divorce probability before dissolution on labour supply after dissolution is twice as large as the effect on labour supply before dissolution: a 10% increase leads to an increase in hours worked of 0.37%, compared to 0.16% before dissolution. Again, divorce exacerbates differences that characterized marriages before dissolution. Shifts in the profile parameters unambiguously increase divorced women's labour supply, because there is no re-adjustment after divorce. Recall that the re-adjustment among married women comes from the fact that they reach points of their divorce risk profile where there was a non-zero possibility of divorce, yet they remain married.

Finally, assets at the time of marriage have a statistically significant but economically small negative effect on the labour supply of divorced women, as they did for married women. Time of dissolution also reduces labour supply, as expected. This variable enters the specification for divorced women through its effect on the married couple's marginal utility of wealth prior to dissolution. Therefore, we expect the estimated coefficient on time of dissolution in the divorced equation to be similar to the estimated

[24] The larger effect of spousal wage and quality of time changes on the labour supply of women after divorce may also arise from complementarities in leisure that are not fully accounted for by the quality of time proxy. In that case, after divorce the complementarities disappear and women are left only with the higher lifetime income. I leave the exploration of that channel for further research.

coefficient on time since marriage in the married equation. The results presented in Table 1.4 and Table 1.1 are in line with this, showing estimated coefficients equal to −0.011 and −0.012 respectively. As for time since dissolution, it reflects the updating to the marginal utility of wealth in the divorced household. It is negative, suggesting a decrease in the marginal utility of wealth, i.e. an increase in wealth. Furthermore, it is one and a half times the magnitude of the updating prior to dissolution. One possibility is that divorced women have a greater expectation of one-period shocks that would increase their marginal utility of wealth than spouses in married couples. Alternatively, the difference in magnitude is also consistent with divorced women having a lower rate of time preferences than married spouses.

1.7.3 Results and the reduced-form literature

The theoretical model and empirical analysis I develop in this paper can be used to understand the conflicting results from the reduced-form literature. To illustrate this, I consider four papers from the literature and set them in the life-cycle model of labour supply with divorce: Peters (1986), Johnson and Skinner (1986), Stevenson (2007) and Mueller (2005). Peters (1986) is one of the first contributions to the economics literature on divorce, and its use of divorce laws to drive variation in divorce risk set the tone for much of what came after in the reduced-form literature. Johnson and Skinner (1986) is a useful paper to consider, because it is one of few who attempted to control for individual divorce risk. As for Stevenson (2007), it is particularly interesting to consider in the context of this paper, since it explicitly accounts for the time since marriage, which I argue is central to our understanding of divorce. Finally, I include Mueller (2005) both because it is part of a relatively scarce literature using Canadian data, and it uses a fixed effects estimator which is of interest in light of the present study.

To simplify exposition, I start by abstracting from the effect of intra-household bargaining power and economies of scale. I do this in part because the reduced-form literature does not account for this couple-specific factors, and so there is no comparison to draw here that is not heavily based on assumptions regarding where these factors come in in their empirical models. The labour supply for married and divorced women is determined respectively by the following two estimating equations

$$
\begin{aligned}
\ln N^M_{1,i,t} &= \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{1,i,j} + \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{2,i,j} + \sum_{j=0}^{T} E_0 \ln \omega_{i,j} + \sum_{j=1}^{t} \ln \Omega_{i,j} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^M_{1,i,t} \\
&= \sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + \sum_{k=0}^{2} \pi^{\omega}_{k,i} + \sum_{j=1}^{t} \ln \Omega_{i,j} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^M_{1,i,t}
\end{aligned}
\tag{1.47}
$$

$$
\begin{aligned}
\ln N^D_{1,i,t} &= \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{1,i,j} + \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{2,i,j} + \sum_{j=0}^{T} E_0 \ln \omega_{i,j} - \sum_{j=1}^{t} \ln \Omega_{i,j} + b^D t + \ln \mathrm{w}_{1,i,t} + u^D_{1,i,t} \\
&= \sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + \sum_{k=0}^{2} \pi^{\omega}_{k,i} - \sum_{j=1}^{t} \ln \Omega_{i,j} + b^D t + \ln \mathrm{w}_{1,i,t} + u^D_{1,i,t}
\end{aligned}
\tag{1.48}
$$

Life-cycle labour supply and changes to divorce law

The reduced-form literature has largely relied on the variation afforded by legislative changes regarding grounds for divorce and marital property distribution in the US. From the late 1960s and for several decades, states have progressively transitioned from mutual consent and fault-based divorce rules, to unilateral and/or no-fault divorce. That is, they have moved from a legal setting in which spouses could divorce if both agreed, or if one of them had committed abuse or adultery, to a setting where a spouse could choose unilaterally to leave the other one, without having to prove that that person had committed an offence.[25] Moreover, the US also experienced several state-wide shifts from title-based property distribution at divorce to equitable or communal distribution systems.

I assume that divorce is impossible under mutual consent laws, or that married individuals perceive it as such. The same analysis would apply if I allowed for a non-zero, but lower, divorce probability in mutual consent states. For now, I treat the labour supply of married women in divorced states as determined by

$$\begin{aligned}\ln N^{M,mut}_{1,i,t} &= \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{1,i,j} + \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{2,i,j} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^{M,mut}_{1,i,t} \\ &= \sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^{M,mut}_{1,i,t}\end{aligned} \tag{1.49}$$

In this context, consider what happens to women who live in states that transition from mutual consent to unilateral grounds for divorce. This potentially affects women differently depending on whether they were already married at when the law changed. Let t^*_i denote the year of marriage couple i was in when the new legislation was introduced. The labour supply of women who were married before the change is

$$\begin{aligned}\ln N^{M,uni-b}_{1,i,t} &= \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{1,i,j} + \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{2,i,j} + \sum_{j=t^*}^{T} E_{t^*} \ln \omega_{i,j} + \sum_{j=t^*_i}^{t} \ln \Omega_{i,j} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^{M,uni-b}_{1,i,t} \\ &= \sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + \underbrace{\sum_{k=0}^{2} \pi^{\omega}_{t^*,k,i}}_{1b} + \underbrace{\sum_{j=t^*_i}^{t} \ln \Omega_{i,j}}_{2b} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u^{M,uni-b}_{1,i,t}\end{aligned} \tag{1.50}$$

[25] Unilateral and no-fault divorce are not technically the same thing, and some states have adopted one but not the other (Jacob, 1989). In practice however the admission of no-fault grounds has resulted in unilateral divorce, so that both terms are frequently used interchangeably in the literature.

and the labour supply of women who were married only after the change is

$$\begin{aligned}\ln N_{1,i,t}^{M,uni-a} &= \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{1,i,j} + \sum_{j=0}^{T} E_0 \ln \mathrm{w}_{2,i,j} + \sum_{j=0}^{T} E_0 \ln \omega_{i,j} + \sum_{j=1}^{t} \ln \Omega_{i,j} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u_{1,i,t}^{M,uni-a} \\ &= \sum_{k=0}^{2} \pi_{k,i}^{\mathrm{w}_1} + \sum_{k=0}^{2} \pi_{k,i}^{\mathrm{w}_2} + \underbrace{\sum_{k=0}^{2} \pi_{k,i}^{\omega}}_{1a} + \underbrace{\sum_{j=1}^{t} \ln \Omega_{i,j}}_{2a} + b^M t + \ln \mathrm{w}_{1,i,t} + \ln \mathrm{w}_{2,i,t} + u_{1,i,t}^{M,uni-a}\end{aligned} \tag{1.51}$$

Equation 1.50 and Equation 1.51 differ in two ways. First, expectations over the divorce risk profile are formulated at time t_i^* for couples who were married before the legislative change. This is illustrated in Equation 1.50, where the $\pi_{k,i}^{\omega}$'s are indexed by t^* (term $1b$). The $\pi_{k,i}^{\omega}$'s in Equation 1.51 are the same ones we've been referring to in the rest of the paper (term $1a$). Second, terms $2b$ and $2a$ differ because the re-adjustment period starts at the time of legislative change in Equation 1.50.

Peters (1986)

The paper uses the 1979 cross-section of the CPS to compare the labour force participation of married women across states with different grounds for divorce. It estimates the following equation

$$LFP_{i,1979}^{M} = \beta_0 + \beta_1 UNI_{i,1979} + X_i' \beta_2 + u_i^M \tag{1.52}$$

where $LFP_{i,1979}^{M}$ denotes the labour force participation of married woman i in 1979, $UNI_{i,1979}$ is equal to 1 if woman i lived in a unilateral divorce state in 1979, X_i is a vector of individual characteristics, including own age, age of children, education, an indicator for white sample members, husband earnings, and region. Peters (1986) finds that labour force participation is higher by 2.22% in unilateral divorce states.

Before turning to the implication of this model for labour supply in the presence of divorce risk, it is worth recalling what is being estimated as a cross-sectional model of labour supply. It is particularly hard to pin point exactly what it is that the controls included in X_i capture. On the one hand, Peters (1986) controls for education and race, which are expected to capture differences in wage profiles and initial assets levels across the women included in the sample. Therefore, any wage variation used in the specification would come from changes within individuals over time, and the coefficient on own wages would be interpreted as the intertemporal substitution elasticity. However, the vector of covariates X_i does not include own wages. Part of the effect of women's movements along their wage profile is captured by the coefficient on age, but any non-linearities are omitted and end up in the error term. Broadly speaking, if $\ln \hat{\mathrm{w}}_{1,i,t_1} = \pi_{0,i}^{\mathrm{w}_1} + \pi_{1,i}^{\mathrm{w}_1} t_1 + \pi_{2,i}^{\mathrm{w}_1} t_1^2$, where t_1 denotes the wife's age, the $\pi_{j,i}^{\mathrm{w}_1}$, $j = 0, 1, 2$ are functions of time-invariant characteristics, then $\pi_{2,i}^{\mathrm{w}_1} t^2$ is in the error term. Since theory predicts that $\pi_{2,i}^{\mathrm{w}_1} < 0$, then this should introduce downward bias in the estimate of the coefficient on age. Furthermore, this bias is likely to be reinforced by the omission of time since marriage from the empirical specification, for which I estimated a negative effect on labour supply.

As mentioned, education and race are most likely picking up differences in own wage profiles and initial asset levels across women. The coefficient on race is consistent with that: Peters (1986) estimates that white women are 12.52% less likely to participate in the labour force. As for the coefficient on education, it suggests that one more year of education increases labour force participation by 3.46%. Therefore, education is not only capturing the income effect from shifts in wage and asset profiles. Because wages are not controlled for, it may also capture part of the effect of movements along women's wage profiles. Women included in the sample were on 44 years old on average in 1979, or born in 1934. Together, the eras in which women were respectively born and interviewed implies that there may have been later-life variation in education that captures a function of age. This would explain the positive coefficient on education.

Next, to see what is picked up by the inclusion of husband's earnings in the control variables, consider that $\ln \hat{w}_{2,i,t_2} = \pi^{w_2}_{0,i} + \pi^{w_2}_{1,i} t_2 + \pi^{w_2}_{2,i} t_2^2$, where t_2 denotes the husband's age and again the $\pi^{w_2}_{j,i}$, $j = 0, 1, 2$ are functions of time-invariant characteristics. Suppose (as I do) that the wife's characteristics do a good enough job of capturing the husband's characteristics. If husband's age was used as part of the covariates, then variation in the husband's earnings would come from the square term, and so the coefficient on husband's earnings would correspond to the intertemporal substitution elasticity with respect to spousal wages. However, since only the wife's age is used, the variation in husband's earnings used in fact corresponds to $\pi^{w_2}_{1,i} \tau_i + \pi^{w_2}_{2,i} t_2^2$, where τ_i is the age gap between the spouses in couple i. If age gap is negatively correlated with the quality of time spent together, then some of the variation used here comes from differences in the quality of time profiles of couples. The intertemporal substitution elasticity with respect to husband's earnings is positive, but the theory and results from this paper show that higher quality of time profiles lead to decreases in labour supply in all periods.

Finally, I expect her regional controls to capture some of the variation that affects labour supply through intra-household bargaining power and economies of scale in my estimation. Most problematic for the comparison with my estimates is the inclusion of children in the estimation.

Next, consider the specific implications with respect to the effect of divorce risk on the labour supply of married women. In 1979, 31 states had switched over to unilateral divorce, with transitions occurring between 1970 (California, Iowa) and 1977 (Wyoming).[26] Therefore, the sample of women who are coded as having $UNI_{i,1979}$ is a combination of women who've been exposed to the law for two to nine years, including marriages that took place before and after the changes. According to Equation 1.51 and Equation 1.50, the effect of unilateral divorce law for women who were married after and before divorce is respectively

[26] In practice, there is a range of laws that fall somewhere between pure mutual consent and pure unilateral, including combinations of fault and no-fault grounds for divorce and requirements of a minimum separation period. Different papers have used different classifications, and Jacob (1989) has shown that results can be sensitive to the classification used.

$$
\begin{array}{llccccc}
\text{married after:} & \underbrace{\pi^{\omega}_{0,i} \quad + \quad \pi^{\omega}_{1} \quad + \quad \pi^{\omega}_{2}}_{(1a)} & - & \underset{(2a)}{\sum_{j=1}^{t^*-1} \ln \Omega_{i,j}} & - & \underset{(3a)}{\sum_{j=t^*}^{t} \ln \Omega_{i,j}} \\
\text{married before:} & \underbrace{\pi^{\omega}_{t^*,0,i} \quad + \quad \pi^{\omega}_{t^*,1,i} \quad + \quad \pi^{\omega}_{t^*,2,i}}_{(1b)} & - & \underset{(2b)}{\sum_{j=1}^{t^*-1} \ln(1)} & - & \underset{(3b)}{\sum_{j=t^*}^{t} \ln \Omega_{i,j}}
\end{array}
$$

In the context of the model I developed in this paper, the variation used by Peters (1986) is expected to yield a positive effect for $UNI_{i,1979}$. However, the size of the effect changes depending on how we think the effect on women differs if they were married before or after the legislative transition. First, consider terms $1a$ and $1b$ above. We might expect that any adjustment made at time $t > 0$ is greater than the adjustment that would have been made at the time of marriage, because some of the uncertainty about the marriage is already resolved. Second, term $2b$ is unambiguously smaller (larger) than term $2a$ if $\Omega_{i,j} > 1$ ($\Omega_{i,j} < 1$) for $j = 1,\ldots,t^* - 1$. Finally, it appears reasonable to expect that terms $3a$ and $3b$ are equal, i.e., that the re-adjustment in any given period only depends on the time since marriage. This is the case, for instance, if the contemporary probability of divorce and intra-household bargaining power are not dependent on previous values. In this context then, whether it matters that some women were already married at the time of the legislative change depends on the comparison between the sum of terms $1a$ and $2a$ and the sum of terms $1b$ and $2b$. Suppose that $\pi^{\omega}_{0,i} + \pi^{\omega}_{1,i} + \omega^{\omega}_{2,i} - \sum_{j=1}^{t^*-1} \ln \Omega_{i,j} = \pi^{\omega}_{t^*,0,i} + \pi^{\omega}_{t^*,1,i} + \omega^{\omega}_{t^*,2,i}$; that is, if married people go from a regime where divorce is not (less) possible to one where it is, their adjustment is identical to where they would have been had they adjusted fully at time of marriage then re-adjusted. Conversely, one may imagine that $\pi^{\omega}_{0,i} + \pi^{\omega}_{1,i} + \omega^{\omega}_{2,i} - \sum_{j=1}^{t^*-1} \ln \Omega_{i,j} > \pi^{\omega}_{t^*,0,i} + \pi^{\omega}_{t^*,1,i} + \omega^{\omega}_{t^*,2,i}$; i.e., that any net adjustment that happens from the time of marriage is larger, because of the greater uncertainty for instance.

To draw a comparison between Peters (1986)'s result and the effect of a similar legislative change in the life-cycle model of labour supply with divorce (with Canadian data), I proceed as follows:

1. Pick a year in the sample (eg. 1982), and randomly assign a regime to each woman who is observed married in that year: mutual consent, unilateral before own marriage, unilateral after own marriage.

2. Assign the effect of divorce on each woman, based on the randomly-selected regime and the estimated results from Section 1.7:

   - Mutual consent: no effect.
   - Unilateral, married before change: $\sum_{k=0}^{2} \pi^{\omega}_{t^*,k,i} - \sum_{j=t_i^*}^{t} \ln \Omega_{i,j}$.
   - Unilateral, married after change: $\sum_{k=0}^{2} \pi^{\omega}_{k,i} - \sum_{j=1}^{t} \ln \Omega_{i,j}$.

3. Calculate the average labour force participation of women in the mutual consent regime, and subtract from the average labour force participation of women in the unilateral regime.

Depending on the assumptions I make with respect to the way the timing of legislative changes affects women who are already married, using Peters (1986)'s methodology I find that labour force participation is higher by 5 to 19 percentage points in unilateral states.

Johnson and Skinner (1986)

This paper is similar to Peters (1986), in that it uses the 1972 cross-section of the CPS. However, in addition to $UNI_{i,1972}$, it also controls for individual divorce probability, $\omega_{i,1972}$. That is,

$$LFP^M_{i,1972} = \beta_0 + \beta_1 UNI_{i,1972} + X_i' \beta_2 + \beta_3 \omega_{i,1972} + u^M_i \tag{1.53}$$

where X_i includes controls for both spouses' age and education, household income, race, and state unemployment rate. Johnson and Skinner (1986) find that living in a unilateral divorce state reduces labour force participation among married women, and that individual divorce probability has no statistically significant effect.

Johnson and Skinner (1986) estimate a number of models which differ in the variables and interactions of variables included. I limit my discussion to their baseline model, which shares many of the features of Peters (1986) with the exception of the inclusion of individual-level divorce risk. Some differences with respect to the covariates included are worth noting. First, both the wife's and husband's education are controlled for, which is expected to capture differences in wage profiles. As in Peters (1986), the estimated coefficients are positive. The sample members in Johnson and Skinner (1986) are approximately ten years younger than in Peters (1986) on average, so again I expect that the positive coefficients are picking up life-cycle variation in education. Notably, only the wife's education is significant. The wife's race is also included, and again, being white has an statistically significant, negative effect on labour supply, in line with expectations. Second, both the wife's and the husband's age are included, so that if the two spouses' earnings were included, the paper would identify the respective intertemporal substitution elasticities. However, the authors instead control for family income minus the wife's earnings. Again then, life-time non-linear variation in the wife's earnings is in the error term, which is expected to bias the results. Furthermore, family income potentially includes capital income, which conflates the associated estimated coefficient even in the absence of bias. As a result, it is unclear what any of this is estimating. The coefficients on own and spousal wages are respectively positive and negative, but neither is estimated to be statistically different from zero. The coefficient on family income minus wife's earnings is negative and statistically significant, confirming that it is probably picking up some of the income effect of assets through the inclusion of capital income.

Ignoring for a second the baseline issues with the specification for labour supply, I consider next the way the role of divorce risk is evaluated. At the time Johnson and Skinner (1986) measure the effect of grounds for divorce laws on female labour supply, only six States had switched over to unilateral divorce, and for at most two years (California and Iowa). If the effect on married women is the same regardless of when the law is introduced, with respect to their marriage year, then this plays no part in explaining the difference in estimated effect between Peters (1986) and Johnson and Skinner (1986).

On the other hand, if the effect is larger among women who were married after the change, then the estimated effect of unilateral divorce laws should be smaller in Johnson and Skinner (1986), since a smaller proportion of the women they observe were married after the change. However, this doesn't explain the reversal in the estimated effect of unilateral divorce. To understand this, it is useful to look at how Johnson and Skinner (1986) control for divorce probability. They estimate individual divorce probability by first constructing an indicator equal to one if a woman gets divorce in the following six years,

$$\mathcal{D}_{i,1972} = \prod_{t=1973}^{1978} 1\{M_{i,t} = 0\}$$

where $M_{i,t} = 0$ if a couple dissolved in year t. That is, their measure of divorce probability today captures a combination of divorce probability today and divorce probability over the next six years. In terms of the model I develop in this paper, this captures current divorce probability, as well as part of individual divorce probability profiles. Hence, women who live in unilateral divorce states face the effects detailed in Equation 1.50 and Equation 1.51. However, controlling for $\omega_{i,1972}$ also captures part of $\pi^{\omega}_{0,i} + \pi^{\omega}_{1,i} + \pi^{\omega}_{2,i}$ (or $\pi^{\omega}_{t^*,0,i} + \pi^{\omega}_{t^*,1,i} + \pi^{\omega}_{t^*,2,i}$, since most women were married before). Consider the extreme case where $\omega_{i,1972}$ perfectly captures the difference in profiles across women in the sample, then $UNI_{i,1972}$ only picks up the difference from the re-adjustment, $-\sum_{j=1}^{t} \ln \Omega_{i,j}$ (or $-\sum_{j=t_i^*}^{t} \ln \Omega_{i,j}$). Therefore, in the life-cycle framework, we expect the effect of $UNI_{i,1972}$ to be negative, in line with Johnson and Skinner (1986)'s results. Earlier, I hypothesized that if the legislative change matters less for people who've already had some time to learn about their marriage, then $\pi^{\omega}_{0,i} + \pi^{\omega}_{1,i} + \omega^{\omega}_{2,i} - \sum_{j=1}^{t^*-1} \ln \Omega_{i,j} > \pi^{\omega}_{t^*,0,i} + \pi^{\omega}_{t^*,1,i} + \omega^{\omega}_{t^*,2,i}$. This scenario would explain why the estimated effect of $\omega_{i,1972}$ is statistically insignificant. Note that although the effect is not statistically significant, Johnson and Skinner (1986) estimate that individual divorce probability increases labour force participation, also consistent with the expected effect from a shift in divorce probability profile.

Stevenson (2007)

Stevenson (2007) uses Census data from 1970 and 1980 to compare the labour supply of newly married women in states with different divorce legislation. Newly weds refer to couples that have been married for two years or less. Hence, she estimates

$$LFP^{M}_{i,t} = \beta_0 + \beta_1 UNI_{i,t} + X_{i,t}\beta_2 + \gamma Year + \sum_{s} State + u^{M}_{i,t} \tag{1.54}$$

where $UNI_{i,t}$ is equal to one if the state had adopted unilateral divorce before woman i's marriage, and $X_{i,t}$ includes own and spousal age, race and ethnicity, as well as the couple's metropolitan status. Stevenson (2007) also includes a control for the length of marriage. She finds that unilateral divorce law increased the labour supply of newly married women by 1.5%.

The specification identifies none of the elasticities discussed in Section 1.7. If wife's earnings were included as part of the covariates, the corresponding coefficient would provide an estimate of intertemporal substitution elasticity. However, Stevenson (2007) leaves income out to capture the effect of unilateral divorce laws through their effect on income. In practice however, if the age distribution changes in unilateral divorce states differently than in mutual-consent states, than movements along the wage profile will be captured in the estimated effect of $UNI_{i,t}$. Although Stevenson (2007) controls for length of marriage, she does not report the results and so it is difficult to draw comparisons with the present paper.

Because Stevenson (2007) uses only newly married women, residents of states that maintain mutual consent divorce are characterized by Equation 1.49 throughout, while women in states that transition to unilateral divorce are characterized by Equation 1.51 in 1980. Hence, the life-cycle model of labour supply with divorce predicts a positive effect, in line with Stevenson (2007)'s results.

It is worth noting that in alternative specification, Stevenson (2007) controls for marital property division laws. She finds that the labour supply of married women in unilateral divorce states is higher irrespective of how property is to be divided in case of dissolution. In the theoretical model I develop in Section 1.3, what happens at divorce only matters through the re-adjustment term, $-\sum_{j=1}^{t} \ln \Omega_{i,j}$. This result would be relatively small, compared to the original shift from the divorce probability profile, among newly married women.

Mueller (2005)

Finally, I consider Mueller (2005), who uses Canada's longitudinal Labour Market Activity Survey for 1988 to 1990. The paper studies the effect of divorce on work behaviour by comparing individual women's labour supply one year before and one year after divorce. Although Mueller (2005) looks a various measures of labour supply (eg. days worked per week), I focus on his labour force participation results. He estimates

$$\begin{aligned}\Delta LFP_{i,1990} &= LFP^{D}_{i,1990} - LFP^{M}_{i,1988} \\ &= \beta_0 + \beta_1 Div_{i,1990} + X_{i,1990}\beta_2 + \Delta u_{i,1990}\end{aligned}$$

where $Div_{i,1990}$ is equal to one if woman i divorced between the two periods, and $X_{i,1990}$ includes controls for age, presence of children, educational attainment, minority and immigrant status, and language. Mueller (2005) estimates that divorce has a positive has a positive but statistically insignificant effect on female labour force participation. His findings are qualitatively similar for different measures of labour supply.

Mueller (2005)'s use of a first-difference estimator makes it effectively impossible to estimate the effect of parametric shifts in the wage profiles. His inclusion of time-invariant controls as the right-hand-side variables for a regression of the difference in labour force participation controls for some life-cycle differences, but the estimates are not reported.

In the framework I develop, Mueller (2005)'s estimates correspond to the difference between Equation 1.48 and Equation 1.47, respectively the labour force participation of divorced and married women. Here I re-introduce intra-household bargaining power and economies of scale in leisure, since they are part of what changes with marital dissolution. That is,

$$
\begin{aligned}
\Delta LFP_{i,1990} &= LFP^{D}_{i,1990} - LFP^{M}_{i,1988} \\
&= \Bigg[\sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + \sum_{k=0}^{2} \pi^{\omega}_{k,i} + \sum_{k=0}^{2} \pi^{\lambda}_{k,i} + \sum_{k=0}^{2} \pi^{g_\Lambda}_{k,i} - \sum_{j=1}^{t^d-1} \ln \Omega_{i,j} + b^M (t^d - 1) \\
&\qquad + \ln \mathrm{w}_{1,i,t^d-1} + \ln \mathrm{w}_{2,i,t^d-1} + \ln \lambda_{i,t} + \ln g_{\Lambda,i,t} + u^{D}_{1,i,t^d-1} \Bigg] \\
&\quad - \Bigg[\sum_{k=0}^{2} \pi^{\mathrm{w}_1}_{k,i} + \sum_{k=0}^{2} \pi^{\mathrm{w}_2}_{k,i} + \sum_{k=0}^{2} \pi^{\omega}_{k,i} + \sum_{k=0}^{2} \pi^{\lambda}_{k,i} + \sum_{k=0}^{2} \pi^{g_\Lambda}_{k,i} + \sum_{j=1}^{t^d-1} \ln \Omega_{i,j} + b^M (t^d - 1) \\
&\qquad + b^D [(t^d + 1) - t^d] + \ln \mathrm{w}_{1,i,t^d+1} + u^{M}_{1,i,t^d+1} \Bigg] \\
&= b^D + (\ln \mathrm{w}_{1,i,t^d+1} - \ln \mathrm{w}_{1,i,t^d-1}) - \ln \mathrm{w}_{2,i,t^d-1} - \ln \lambda_{i,t} - \ln g_{\Lambda,i,t} + \Delta u_{1,i,1990}
\end{aligned}
$$

where b^D is captured by the intercept in Mueller (2005)'s equation, β_0. Therefore, the estimated effect of divorce is equal to the loss of the husband's wages, net of any growth in own wages between 1988 and 1990, the gain in control over household resources, and the loss of economies of scale in leisure. We can reasonably expect the net change in household earnings to be negative, while the gain of control is necessarily positive. As for economies of scale, they have an ambiguous impact, depending on the couple's relationship before divorce. Consistent with this, Mueller (2005) finds that the estimated effect is not statistically different than zero. The estimate is negative however, suggesting that the biggest effect comes from the net loss in earnings, or from the loss of economies of scale. This estimation can be implemented directly using my data.

Mueller (2005) also estimates a second specification which includes an indicator variable equal to one if the woman's spouse worked full-time in 1988, $FT_{i,1988}$,

$$
\begin{aligned}
\Delta LFP_{i,1990} &= LFP^{D}_{i,1990} - LFP^{M}_{i,1988} \\
&= \beta_0 + \beta_1 Div_{i,1990} + X_{i,1990}\beta_2 + \beta_3 FT_{i,1988} + \beta_4 Div_{i,1990} \times FT_{i,1988} + \Delta u_{i,1990}
\end{aligned}
$$

Once more, he finds no statistically significant effect of divorce, regardless of spousal labour force participation in 1988.

1.8 Model with myopia

The model presented in Section 1.3 suggests that there is no effect of divorce on the divorced, conditional on what they leave their marriage with. However, that result hinges on the assumption that people foresee divorce and that they are capable of smoothing its effect over their married years. In this section, I explore the first of these two premises.

Recall that $Q_{i,t} = \{l_{1,i,t}, l_{2,i,t}, L_{1,i,t}, L_{2,i,t}, c_{1,i,t}, c_{2,i,t}, C_{1,i,t}, C_{2,i,t}, A_{i,t+1}\}$ is the set of choice variables for the household, and suppose that spouses do not foresee divorce. Then, the married problem becomes

$$\begin{aligned} V^M_{1,i}(A_{i,t}) = \max_{Q_{i,t}} \lambda_{i,t} U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) + \lambda_{i,t}\beta E_t\left[V^M_{1,i}(A_{i,t+1})\right] \\ + (1-\lambda_{i,t})U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) + (1-\lambda_{i,t})\beta E_t\left[V^M_{2,i}(A_{i,t+1})\right] \end{aligned}$$

subject to the budget constraint

$$\begin{aligned} \mathrm{w}_{1,i,t} l_{1,i,t} + \mathrm{w}_{2,i,t} l_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t})\kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t\kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1+r_t)A_{i,t} \end{aligned}$$

As shown in Section A.4, the MUWC labour supply of household member 1 when they are married is then

$$N_{1,i,t} = \bar{N}_{1,i,t} - \frac{\lambda_{i,t}\gamma_{l,1,i,t}}{\nu_{i,t}\mathrm{w}_{1,i,t}} - \frac{\lambda_{i,t}\gamma_{L,1,i,t}}{\nu_{i,t} g_{\Lambda,i,t}(\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t})} \tag{1.55}$$

where $\nu_{i,t}$ is the shadow price of relaxing the married household's budget constraint. Note that the labour supply of married women is unchanged, with the exception of $\nu_{i,t}$, which is different under myopia. Specifically,

$$\ln\nu_{i,t} = \ln\nu_{i,0} - \sum_{j=1}^{t} b_{i,j} + \sum_{j=1}^{t}\varepsilon_{i,j} \tag{1.56}$$

The re-adjustment term $\sum_{j=1}^{t}\ln\Omega_{i,j}$ is absent from this equation. Furthermore, $\ln\nu_{i,0}$ is also different. In both the baseline model and the model with myopia, $\nu_{i,0}$ is a function of initial assets, as well as initial expectations over wages, prices and interest rates, leisure and consumption needs, and preferences. Under myopia however, it doesn't depend on the expected lifetime divorce risk profile. I assume that $\ln\nu_{i,0}$ can be approximated by

$$\begin{aligned} \ln\nu_{i,0} = Z_i\phi + \alpha^A A_{i,0} + \sum_{t=0}^{T}\alpha_t^{\mathrm{w}_1} E_0 \ln \mathrm{w}_{1,i,\underline{t}_1+t} + \sum_{t=0}^{T}\alpha_t^{\mathrm{w}_2} E_0 \ln \mathrm{w}_{2,i,\underline{t}_2+t} \\ + \sum_{t=0}^{T}\alpha_t^{g_\Lambda} E_0 \ln g_{\Lambda,i,t} + a_{i,0} \end{aligned} \tag{1.57}$$

which I implement empirically as follows

$$\ln\nu_{i,0} = Z_i\phi + \alpha^A \hat{A}_{i,0} + \hat{\Pi}^{mm'}_{0,i} D^{mm} + a^{mm*}_{i,0} \tag{1.58}$$

where

$$\hat{\Pi}^{mm}_{0,i} = (\hat{\pi}^{\mathrm{w}_1}_{0,i}, \hat{\pi}^{\mathrm{w}_1}_{1,i}, \hat{\pi}^{\mathrm{w}_1}_{2,i}, \hat{\pi}^{\mathrm{w}_2}_{0,i}, \hat{\pi}^{\mathrm{w}_2}_{1,i}, \hat{\pi}^{\mathrm{w}_2}_{2,i}, \hat{\pi}^{g_\Lambda}_{0,i}, \hat{\pi}^{g_\Lambda}_{1,i}, \hat{\pi}^{g_\Lambda}_{2,i})',$$

$$D^{mm} = (\alpha^{\mathrm{w}_1}_0, \alpha^{\mathrm{w}_1}_1, \alpha^{\mathrm{w}_1}_2, \alpha^{\mathrm{w}_2}_0, \alpha^{\mathrm{w}_2}_1, \alpha^{\mathrm{w}_2}_2, \alpha^{g_\Lambda}_0, \alpha^{g_\Lambda}_1, \alpha^{g_\Lambda}_2)',$$

$$\alpha^h_j = \sum_{t=0}^{T} \alpha^h_t t^j, \quad j = 0, 1, 2, \; h \in \{\mathrm{w}_1, \mathrm{w}_2, g_\Lambda\},$$

$$a^{mm*}_{i,0} = a_{i,0} + \textstyle\sum_{t=0}^{T} \left[\alpha^{\mathrm{w}_1}_t \mathrm{v}^{\mathrm{w}_1 *}_{i,t_1} \right] + \sum_{t=0}^{T} \left[\alpha^{\mathrm{w}_2}_t \mathrm{v}^{\mathrm{w}_2 *}_{i,t_2} \right] + \sum_{t=0}^{T} \left[\alpha^{g_\Lambda}_t \mathrm{v}^{g_\Lambda *}_{i,t} \right]$$

Therefore, in the presence of myopia the profile parameters for divorce risk and the contemporary divorce risk do not determine the work choices of married women. The equation for the labour supply of married women is

$$\ln N^M_{1,i,t} = F^M_i + b^M t + \beta^M_N \ln \bar{N}_{1,i,t} + \beta_\lambda \ln \lambda_{i,t} + \beta_g \ln g_{\Lambda,i,t} + \beta^M_1 \ln \mathrm{w}_{1,i,t} + \beta_2 \ln \mathrm{w}_{2,i,t} + u^M_{1,i,t} \tag{1.59}$$

where

$$F^M_i = \ln \gamma_{l,1,i} \ln \gamma_{L,1,i} + \ln \nu_{i,0} \tag{1.60}$$

$$u^M_{1,i,t} = \eta_{l,1,i,t} + \eta_{L,1,i,t} + \sum_{j=1}^{t} \varepsilon_{i,j} \tag{1.61}$$

This model is naturally nested in the baseline model for married labour supply. Looking back to Table 1.1, it is evident that the model with myopia is not supported by the estimation for married women. Both the profile parameters corresponding to divorce probability and the contemporary risk of divorce enter statistically significantly.

To see what happens post-divorce, recall that the contemporary value of the marginal utility of wealth at divorce was derived assuming continuity in the transition from the married to the divorced problem. That is, women's marginal utility of wealth after divorce is

$$\ln \nu_{1,i,t} = \ln \nu_{1,i,t^d} - \sum_{j=t^d+1}^{t} b_{1,i,j} + \sum_{j=t^d+1}^{t} \varepsilon_{1,i,j}, \;\; \forall t > t^d$$

If spouses foresee divorce, then in their married years they set their couple's marginal utility of wealth taking into account the potential transition to divorce. At time $t < t^d$ married spouses make decisions by solving their problem at time T, then working their way backward to their current problem. If their solution to future problems takes into account the possibility of divorce, than at dissolution they don't need to make further adjustments. In that case, $\ln \nu_{1,i,t^d} = \ln \nu_{i,t^d}$, where

$$\begin{aligned} \ln \nu_{i,t^d} &= \ln \nu_{i,t^d-1} - \ln \beta (1 + r_{t^d}) - \ln E_{t^d-1}[\exp(\varepsilon_{i,t^d})] + \varepsilon_{i,t^d} \\ &= \ln \nu_{i,t^d-1} - b_{i,t^d} + \varepsilon_{i,t^d} \end{aligned}$$

That is, the divorced women's marginal utility of wealth at the time of dissolution comes directly from the marginal utility of wealth of the married couple right before divorce. Then,

$$\ln \nu_{1,i,t} = \underbrace{\ln \nu_{i,0} - \sum_{j=1}^{t^d} b^*_{i,j} - \sum_{j=1}^{t^d} \ln \Omega_j + \sum_{j=1}^{t^d} \varepsilon^*_{i,j}}_{\ln \nu_{i,t^d}} - \sum_{j=t^d+1}^{t} b_{1,i,j} + \sum_{j=t^d+1}^{t} \varepsilon_{1,i,j}$$

This is the expression for $\ln \nu_{1,i,t}$ that is used in the baseline model for the labour supply of divorced women. On the contrary, if married spouses don't foresee the possibility of marital dissolution, then the divorced problem is considered for the first time at divorce. Again, the marginal utility of wealth after dissolution is

$$\ln \nu_{1,i,t} = \ln \nu_{1,i,t^d} - \sum_{j=t^d+1}^{t} b_{1,i,j} + \sum_{j=t^d+1}^{t} \varepsilon_{1,i,j}, \ \ \forall t > t^d$$

but $\ln \nu_{1,i,t^d} \neq \ln \nu_{i,t^d-1} - b_{i,t^d} + \varepsilon_{i,t^d}$. To see where $\ln \nu_{1,i,t^d}$ comes from, consider the rest-of-the-lifetime budget constraint associated with the new problem

$$\begin{aligned} & A_{i,t^d} + E_{t^d}\left[\sum_{t=t^d}^{T}\left(\frac{1}{1+r_t}\right)^t \mathrm{w}_{1,i,t}\right] \\ & = E_{t^d}\left[\sum_{t=t^d}^{T}\left(\frac{1}{1+r_t}\right)^t \left(\mathrm{w}_{1,i,t}(\bar{l}_{1,i,t} + \bar{L}_{1,i,t}) + \frac{\gamma_{l,i,1,t} + \gamma_{L,i,1,t}}{\nu_{1,i,t}} + \bar{c}_{1,i,t} + p_t \bar{C}_{1,i,t} + \frac{\gamma_{c,i,1,t} + \gamma_{C,i,1,t}}{\nu_{1,i,t}}\right)\right] \end{aligned} \tag{1.62}$$

I assume that $\ln \nu_{1,i,t^d}$ can be approximated as follows

$$\ln \nu_{1,i,t^d} = Z_i \phi_{t^d} + \alpha^A_{t^d} A_{i,t^d} + \sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,j} E_{t^d} \ln \mathrm{w}_{1,i,t} + a_{i,t^d} \tag{1.63}$$

Similar to what was done in Section 1.5.1 and Section 1.6.1, I also assume that

$$E_{t^d} \ln \mathrm{w}_{1,i,t} = \pi^{\mathrm{w}_1}_{0,i} + \pi^{\mathrm{w}_1}_{1,i} t + \pi^{\mathrm{w}_1}_{2,i} t^2 + \mathrm{v}^{\mathrm{w}_1 *}_{i,t}$$

Therefore,

$$\sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,t} E_{t^d} \ln \mathrm{w}_{1,i,t} = \pi^{\mathrm{w}_1}_{0,i} \alpha^{\mathrm{w}_1}_{t^d,0} + \pi^{\mathrm{w}_1}_{1,i}[t^d \alpha^{\mathrm{w}_1}_{t^d,0} + \alpha^{\mathrm{w}_1}_{t^d,1}] + \pi^{\mathrm{w}_1}_{2,i}[(t^d)^2 \alpha^{\mathrm{w}_1}_{t^d,0} + 2t^d \alpha^{\mathrm{w}_1}_{t^d,1} + \alpha^{\mathrm{w}_1}_{t^d,2}] + \sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,t} \mathrm{v}^{\mathrm{w}_1 *}_{i,t}$$

where $\alpha^{\mathrm{w}_1}_{t^d,0} = \sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,t}$, $\alpha^{\mathrm{w}_1}_{t^d,1} = \sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,t}(t - t^d)$, and $\alpha^{\mathrm{w}_1}_{t^d,2} = \sum_{t=t^d}^{T} \alpha^{\mathrm{w}_1}_{t^d,t}(t - t^d)^2$.
Therefore, I implement $\ln \nu_{1,i,t^d}$ as follows

$$\ln \nu_{1,i,t^d} = Z_i \phi_{t^d} + \alpha^A_{t^d} A_{i,t^d} + \hat{\Pi}^{md'}_{0,i} D^{md}_0 + \hat{\Pi}^{md'}_{1,i} D^{md}_1 + \hat{\Pi}^{md'}_{2,i} D^{md}_2 + a^{md*}_{i,t^d}$$

where

$\hat{\Pi}^{md}_{0,i} = (\hat{\pi}^{w_1}_{0,i}, \hat{\pi}^{w_1}_{1,i}, \hat{\pi}^{w_1}_{2,i})'$,

$\hat{\Pi}^{md}_{1,i} = (\hat{\pi}^{w_1}_{1,i} \times t^d, \hat{\pi}^{w_1}_{2,i} \times t^d)'$,

$\hat{\Pi}^{md}_{2,i} = \hat{\pi}^{w_1}_{2,i} \times (t^d)^2$,

$D_0 = (\alpha^{w_1}_{t^d,0}, \alpha^{w_1}_{t^d,1}, \alpha^{w_1}_{t^d,2})'$,

$D_1 = (\alpha^{w_1}_{t^d,0}, \alpha^{w_1}_{t^d,1})'$,

$D_2 = \alpha^{w_1}_{t^d,0}$,

$a^{md*}_{i,0} = a_{i,0} + \sum_{t=t^d}^{T} \alpha^{w_1}_{t^d,t} v^{w_1*}_{i,t}$

The labour supply equation for divorced women is then

$$\ln N^D_{1,i,t} = F^D_i + b^D(t^d - t) + \beta^M_N \ln \bar{N}_{1,i,t} + \beta^D_1 \ln w_{1,i,t} + u^D_{1,i,t} \tag{1.64}$$

where

$$F^D_i = \ln \gamma_{l,1,i} + \ln \gamma_{L,1,i} + \ln \nu_{1,i,t^d} \tag{1.65}$$

$$u^D_{1,i,t} = \eta_{l,1,i,t} + \eta_{L,1,i,t} + \sum_{j=t^d}^{t} \varepsilon_{1,i,j} \tag{1.66}$$

This differs from the baseline specification for divorced women in three ways. First, only own wage profile parameters are included. The spousal wage, quality of time and divorce risk profiles are omitted. Second, interactions between the slope and curvature parameters and the time of dissolution are added to the set of regressors. Third, assets at the time of dissolution are also included. The corresponding results are shown in Table 1.7. Notably, the spousal wage, quality of time and divorce risk profile parameters are statistically significant. Again, Table 1.8 and Table 1.9 provide summaries of those results.

The parameters characterizing the own wage profile enter as interactions with the time at dissolution. In Table 1.7, two of the coefficients on these added regressors are statistically significant. To gain some intuition for their direction and magnitude, consider where they come in with respect to effects we've already looked at. Table 1.8 shows uncorrected and corrected results for own wage. In the top half of the table, the first column shows summarized results using the estimates from the myopia model, but without taking into account the added terms. The second column shows results with the added terms. Since the parameter interactions only affect the slope and curvature, there is no difference for any effect other than the impact of changes in those parameters. The estimated coefficient on $(\pi^{w_1}_{1,i} \times t^d)$ is statistically significant, but it is two orders of magnitude smaller than the estimated coefficient on $\pi^{w_1}_{1,i}$. In practice, its effect is not large. A 10% change in $\pi^{w_1}_{1,i}$ decreases labour supply up to age 41 without the correction, and up to ages 40 and 38 with the correction, respectively for marriages

Table 1.7: Effect of divorce on divorced women, with myopia

		Coefficient	(p-value)
PANEL A			
Assets (t=0)	π_0^y	-0.000	(0.000)
Own wage	$\pi_0^{w_1}$	-0.176	(0.158)
	$\pi_1^{w_1}$	-5.052	(0.022)
	$\pi_2^{w_1}$	-118.562	(0.008)
Spousal wage	$\pi_0^{w_2}$	-0.192	(0.004)
	$\pi_1^{w_2}$	-3.715	(0.005)
	$\pi_2^{w_2}$	-71.630	(0.011)
Prob. divorce	π_0^{ω}	1.739	(0.001)
	π_1^{ω}	63.856	(0.000)
	π_2^{ω}	1929.943	(0.000)
Quality of time	$\pi_0^{g_\Lambda}$	-0.460	(0.003)
	$\pi_1^{g_\Lambda}$	-5.843	(0.002)
Time of dissolution		-0.003	(0.639)
PANEL B			
Time since dissolution		-0.018	(0.000)
$\ln w_1$		0.321	(0.002)
PANEL C			
Assets (t=t^d)		0.000	(0.002)
Own wage, interac.	$\pi_1^{w_1} \times t^d$	0.030	(0.098)
	$\pi_2^{w_1} \times t^d$	-0.746	(0.163)
	$\pi_2^{w_1} \times (t^d)^2$	0.066	(0.000)
Constant		5.076	(0.052)
N		1,500	

dissolved after five and 25 years. Alternatively, consider the difference in the labour supply of two 50-year-old women, one of which divorced five years into her marriage, and the other, 25 years (assuming they both got married at age 25). In the first case, a 10% change in the slope of her wage profile results in a 31% increase in hours worked at age 50; in the second case, that same shift leads to a 37% increase in hours worked. The difference stems from the fact that the increase in slope has a smaller negative income effect (in absolute value), the later a woman is divorced. Next, consider the effect of a change in the curvature of a woman's own wage profile. A 10% change in $\pi_{2,i}^{w_1}$ decreases labour supply up to age 44 without the correction, and up to ages 44 and 41 with the correction, respectively for marriages dissolved after five and 25 years.

In brief, there is only limited evidence that the model with myopia is a better representation for what happens after divorce. The profile parameters for spousal wage, quality of time and divorce probability are statistically significant determinants of labour supply among divorced women. Furthermore, the interactions between the own profile parameters and time of dissolution are statistically significant, but their effect is economically small. Similarly, assets at the time of divorce have an economically

Table 1.8: Summary of effects on divorced women, with myopia: own wage, spousal wage, and quality of time

	Own wage	Own wage, corrected
Effect of 10% ↑ at time t		
Intertemp.	3.21%	3.21%
Own-unc.	3.21%	3.21%
Cross-unc.	0.00%	0.00%
Effect of 10% ↑ in profile parameters		
$\pi_{0,i}$	3.21%	
$\pi_{1,i}$	$(-5.052+0.321t)\times 10\%$ No effect at age 41	$(-5.052+0.030t^d+0.321t)\times 10\%$ No effect at age 40 ($t^d=5$) / 38 ($t^d=25$)
$\pi_{2,i}$	$(-118.562+0.321t^2)\times 10\%$ No effect at age 44	$(-118.562+0.066(t^d)^2+0.321t^2)\times 10\%$ No effect at age 44 ($t^d=5$) / 41 ($t^d=25$)

	Spousal wage	Quality of time
Effect of 10% ↑ at time t		
Intertemp.	–	–
Own-unc.	–	–
Cross-unc.	−0.05%	−0.12%
Effect of 10% ↑ in profile parameters		
$\pi_{0,i}$	−1.92%	−4.60%
$\pi_{1,i}$	−37.15%	−58.43%
$\pi_{2,i}$	−716.30%	–

Table 1.9: Summary of effects on divorced women, with myopia: divorce probability

	Divorce probability
Effect of 10% ↑ in divorce risk at time t	
Intertemp.	–
Own-unc.	–
Cross-unc.	0.43%
Effect of 10% ↑ in profile parameters	
$\pi_{0,i}$	17.39%
$\pi_{1,i}$	638.56%
$\pi_{2,i}$	19,299.43%

insignificant impact of labour supply among divorced women. Finally, the estimates of intertemporal substitution elasticity and own- and cross-uncompensated elasticities are essentially unchanged by the inclusion of the myopia-specific terms, for own wages, spousal wages, and quality of time.

1.9 Conclusion

In this paper, I show that divorce and divorce risk impact the labour supply of women, from the time they get married to the period following marital dissolution. A 10% increase in the lifetime divorce risk faced by women when they get married leads to a 4.34 to 5.88% increase in their labour supply in every period throughout their marriage. This is similar in magnitude to the effect of a 10% increase in their lifetime wages. At the time of marriage, the expectation that they may get divorced in the future leads women to increase their labour supply. As they progress through their marriage and divorce is not realized, women then re-adjust their labour supply downwards. That re-adjustment is greatest among women in couples where both spouses would face tighter budget constraints after divorce: it is approximately four times as large as in couples where spouses would face looser budget constraints after dissolution.

Among women who do experience divorce, the event exacerbates pre-existing differences across marriages. In other words, the effect of divorce on women's labour supply is conditional on the marriage they were in. Differences in spousal wage, quality of time spent with one's spouse and divorce risk while married have larger impacts on women's labour supply after than before divorce. For instance, an increase in spousal wage at time t while women are married increases their labour supply in that period by 1.76%. However, it has no effect in other periods, until after divorce. In other words, changes in spousal wage have no cross-effect until after marital dissolution. Similarly, the cross-effect of an increase in the quality of time spent with one's spouse is one and a half time as large after divorce; and the cross-effect of an increase in divorce risk is twice as large.

The reduced-form literature on divorce has extensively documented the fact that marital dissolution should be viewed as a process, starting well before the actual event. However, it has produced largely contradictory results on the direction of the impact on women's labour supply. In this paper, I develop a model of life-cycle labour supply that takes into account the effect of marital dissolution, and provides the necessary structure to understand how this effect unfolds over women's lifetime. In every period, decisions are taken collectively within the married household, accounting for the role of intra-household bargaining power and economies of scale. I show that the theoretical model and its empirical results are consistent with findings from the reduced-form literature. Taking into account how previous work has empirically modelled labour supply and controlled for marital dissolution, a life-cycle perspective on the effect of divorce and divorce risk explains disparities in findings. Furthermore, I also contribute to the structural literature on marital dissolution, which has focused on divorce risk as a determinant of bargaining power in marriage. I show that even after controlling for intra-household resource allocation, women's insurance response to divorce risk has a significant impact on their labour supply.

Understanding how marital dissolution affects women at different points in their lifetime provides a stronger basis for policy making. Being divorced is correlated with a wide range of negative outcomes, including poverty and morbidity (eg. Curtis and Rybczynski, 2014; Joung et al., 1997). In this paper, I show that divorce exacerbates pre-existing differences across women's marriages. As such, it constitutes a significant shock for those who experience it. Although economic and other differences across women may not be driven by marital dissolution, their effect is more severe after divorce than during marriage. Furthermore, the fact that divorce impacts women by exacerbating pre-existing factors means that corresponding policies should at least in part be targeted earlier in life, and at factors that directly affect women's economic security. Indeed, there are vulnerabilities that may not be realized during marriage but matter if and when women experience divorce.

The role played by factors that already exist during marriage in the impact of divorce points to important directions in which to extend the theoretical model and empirical analysis I present in this paper. Further research should integrate human capital accumulation and fertility choices made by women during marriage. Existing evidence from the structural literature (eg. Mazzocco et al., 2013) suggests that these may interact with the income effect of lifetime divorce risk and with the re-adjustment process throughout marriage, and may result in differences across women that are magnified by divorce.

Chapter 2

An exploration of intergenerational income mobility with the Longitudinal and International Study of Adults for Canada

2.1 Introduction

Existing work on the intergenerational transmission of (dis)advantage has found Canada to be fairly mobile: a 10% increase in a man's income is associated with a 2.3 to 3.8% increase in their son's adult income, and with a 1.9 to 2.9% increase in their daughter's income (Chen et al., 2017). This is half as much as in the United States, but 30% higher than in Finland, Norway or Denmark (Corak, 2013). Everywhere, institutions and policies that affect investments in and returns to human capital play a significant part in determining the extent to which different circumstances early on lead to gaps in later-life outcomes. For instance, Corak (2013) discusses important distinctions between Canada and the United States, which pave the path of children from conception to labour market participation. In the United States, children's physical and mental health outcomes, school readiness, and postsecondary attendance are all more tightly associated with parental outcomes than in Canada. In turn, the college premium is substantially higher in the United States, and so is income inequality within a given generation. In that context, research has found cognitive skills and education to account for over half of the intergeneration correlation in incomes in the United States (Bowles and Gintis, 2002). In this paper, we evaluate the extent to which Canadians' experience of intergenerational mobility is associated with their education.

The literature on intergenerational mobility has broadly followed two courses. On the one hand, it has sought to refine the measurement of mobility, exploring the effect of data limitations (eg. Atkinson et al., 1983; Chen et al., 2017; Solon, 1992; Zimmerman, 1992), but also seeking to draw comparisons over time (eg. Connolly et al., 2019) and across places (eg. Chetty, Hendren, Kline, and Saez, 2014;

Chetty, Hendren, Kline, Saez, and Turner, 2014; Connolly et al., 2019; Corak, 2017). On the other hand, the development of increasingly rich data sets has given impetus to a growing literature on the mechanisms that underlie some of the broader numbers. Bridging these two strands are a small number of papers that have tried to decompose measures of intergenerational mobility; i.e., that have tried to map these numbers back to some of the investigated drivers. In the United States and the United Kingdom, these have found that approximately one half of the observed intergenerational correlation in income is associated with education (respectively Bowles and Gintis, 2002, and Blanden et al., 2007). This relationship hinges on two links: first, parental income must be associated with education, and second, education must be associated with traits and skills that are valued on the labour market. The institutional distinctions Corak (2013) draws between Canada and the United States suggest that both of these links are weaker here than south of the border. Therefore, we not only expect that the correlation would be smaller in Canada, but also that education would explain a smaller portion of the correlation that we do see.

To investigate this, we take advantage of the Longitudinal and International Study of Adults (LISA), which provides detailed survey data on respondents, as well as a panel of administrative data covering 1982 to 2013 for both respondents and their parents. The link between survey and administrative data allows us to measure the part played by education in the intergenerational income correlation of Canadians. To do so, we use regressions of child income rank on parental income rank. We decompose the intergenerational income correlation coefficient and find that educational attainment accounts for 38.2 to 50.1% of the association, depending on the measures of child and parental income used. Furthermore, we look separately at individuals whose highest attainment is a Bachelor's degree, and at those who have completed graduate studies. In most cases, we find evidence that field of study strengthens income correlation. Although more research is required to understand the mechanisms behind the observed relationships, this result is particularly interesting given the massification of tertiary education.

Unsurprisingly, more educated respondents exhibit higher incomes: children born at the very bottom of the parental income distribution, and who don't complete high school, end up at the 20th percentile of their generation's income distribution; conversely, those who graduate university make it past the median. However, we find that differences in education do not translate to differences in the size of the intergenerational income correlation. This contrasts with results from the United States, where mobility is highest among university graduates (eg. Hout, 1988; Torche, 2011; see also Wanner and Hayes, 1996, for Canada). It implies that education plays less of an equalizing role in Canada, or that selection into higher education is not as strongly tied to mobility as in the United States. In comparison, we find evidence suggesting that intergenerational income correlation is higher for men than for women, and much lower for members of visible minority groups.

As mentioned, approximately half of intergenerational mobility is associated with education-related attributes and skills; i.e., characteristics that are leveraged or acquired in education, or that are themselves correlated with such characteristics. To better understand the nature of these skills, we further decompose intergenerational mobility to account for the role played by a number of skills respondents report using at work (reading, writing, communication, mathematics, physical strength and dexterity) and specific indicators of job quality (unionization status, permanence of contract and authority over other employees). We find that almost half of the portion of intergenerational mobility that is associated with education is itself correlated with job skills. However, very little of the role played by education has to do with selection into higher-quality jobs, net of job skills. Altogether, job characteristics are also associated with 16.5 to 27.5% of the portion of intergenerational mobility that is not correlated with education. Again, job skills play the biggest part.

In this context, our contribution is two-fold. First, we show that despite greater mobility in Canada, the portion of the intergenerational correlation in incomes that is explained by education is similar to what is reported for the United States and United Kingdom (respectively Bowles and Gintis, 2002, and Blanden et al., 2007), if not higher. This suggests a role for educational policy to further its potential in leveling the playing field. Second, to the best of our knowledge, our paper provides the first evidence on the part played by job skills in intergenerational mobility. In particular, we find that half of the portion of income correlation that is associated with education is linked to job skills, including reading, writing, communication, mathematics, physical strength and dexterity. This is useful to understand the role of education and related institutions in Canada, but it also contributes to the broader literature on the relative role of cognitive skills. Research has generally focused on ability measured in childhood and adolescence (eg. Blanden et al., 2007). Our analysis speaks directly to the skills people use at work, thus more closely to the skills people are remunerated for.

2.2 Literature review

Apart from the role of wealth and inheritance, children born to high- (low-) income parents themselves have high (low) income in adulthood if being born in a high-income family is associated with the development of characteristics that are rewarded on the labour market; for instance, abilities, human capital, social networks, and health. As such, the effect of parental income on children's income results from the cumulative and complementary interactions of pre- and post-natal environments (Björklund et al., 2006; Corak, 2013; Cunha et al., 2006). In this context, education is an intermediary on the path from parental to child income. The former impacts child education directly as well as indirectly: higher-income parents may have more opportunities to invest in their children's education; and may possess characteristics that contribute to higher educational attainment. In turn, these same investments and characteristics can affect income through education, or by helping children develop skills, attitudes and habits that contribute directly to employment.

Becker and Tomes (1979, 1986) provide an early theorization of the impact of parental investment in human capital development on the intergenerational transmission of income. In the presence of credit constraints, lower-income parents are less likely to be able to invest in their children's skills. As expected, the empirical literature finds gaps between low-income and high-income families in their level of parental expenditures in child education (Kaushal et al., 2011; Kornrich and Furstenberg, 2013;

Schneider et al., 2018). These include financial resources to cover school costs, including tertiary education. In the Canadian context, Frenette et al. (2017) find that high-income/high-wealth families are more likely to invest in RESPs for their children. The presence of RESP by age 15 is associated with a greater probability of attending postsecondary education (Frenette et al., 2017). Expenditures also cover "shadow curriculum" activities, such as test preps (Bray, 1999), and other resources, like books and computers (Duncan and Murnane, 2011; see also Corak, 2013). Higher income and more educated parents are also likely to spend more time with their children (Ramey and Ramey, 2010; Schneider et al., 2018).

There is evidence that differences in parental investments do lead to disparities in children's skills and abilities, starting early on. Several studies highlight the relationship between family income when children are young and educational attainment (Bailey and Dynarski, 2011; Belley and Lochner, 2007) as well as cognitive and non-cognitive skills (Cunha and Heckman, 2007; Knudsen et al., 2006; Waldfogel and Washbrook, 2011). In turn, early cognitive and non-cognitive skill development matters for later skill acquisition and educational success (Cunha et al., 2006). Evidence of the limited impact of credit constraints on teenagers further hint to the importance of the timing of investment opportunities (eg. Carneiro and Heckman, 2002). Finally, with the expansion of postsecondary education differences in parental income may be expressed not only in terms of attainment, but also with respect to graduates' choice of field of study (Zarifa, 2012a).

As mentioned before, the effect of parents on their children's education can operate specifically through the parents' own human capital. A large literature has investigated the intergenerational transmission of education, with an increasing focus on producing causal estimates (see Black and Devereux, 2011, for a review). Results vary with the identification strategy, and have generally been found smaller than non-causal estimates. However, the literature broadly finds that parental education is important among post-natal environment factors (eg. Björklund et al., 2006, and Holmlund et al., 2011, in Scandinavia; De Haan, 2011, and Oreopoulos et al., 2006, in the United States); that its effect is lasting (Dickson et al., 2016); and that mother's and father's education don't necessarily have the same effect (Björklund et al., 2006; Black et al., 2005). Although they could not identify causation, Aydemir et al. (2013) find positive correlation between parents and children's education among Canadians, especially among those whose parents were also born in Canada. Other forms of human capital transmission have also been studied; for instance, Currie and Moretti (2003, 2007) have found a positive impact of mothers' education and health on their children's health.

In the United States and the United Kingdom respectively, Bowles and Gintis (2002) and Blanden et al. (2007) have found that approximately half of the relationship between parental and child incomes is explained by the combination of children's education and cognitive and non-cognitive skills. Child education also plays an important role in the transmission of socioeconomic status, beyond income (Blau and Duncan, 1967; Passeron and Bourdieu, 1964; Erikson and Goldthorpe, 1992).[1] Although

[1]Research in sociology predominantly relies on measures of status attainment based on social class categories (see Erikson and Goldthorpe, 1992). Research following this tradition construct their measures of social class using broad occupational

descriptive, this evidence is a crucial step in understanding the link between national (and subnational) estimates of intergenerational mobility, and the abundant literature that has zeroed in on the various possible mechanisms at play. It contrasts the role of education-related skills, skills that either contribute to or are developed through education, and non-education-related skills. The latter include a broad range of abilities and resources, encompassing such things as skills acquired through on-the-job training as well as social networks that are likely to be at least partially shared across generations. Our paper contributes to this literature, by documenting the part played by education in the Canadian context. Differences in the social and institutional setting in Canada may imply a different role for education than estimates from Bowles and Gintis (2002) and Blanden et al. (2007). For instance, Corak et al. (2011) show that along many dimensions Canadian children face less unequal environments than children in the United States. This may reduce both the overall intergenerational transmission of (dis)advantage and the role of education.

Importantly, earlier work using occupations as a proxy for socioeconomic status has found a significant role for education, even in the Canadian context. Investigating a cohort of Ontario males, Ornstein (1981) estimated that educational attainment accounted for a large portion of the association between father occupation and child household income (the measure of socioeconomic status used in the paper is an index largely based on occupational levels of income and education). Likewise, Wanner and Hayes (1996) finds that the association between father's and son's occupational status in Canada decreases by more than 40% when controlling for education. Furthermore, Wanner and Hayes (1996) found that university graduates in the General Social Survey (GSS) exhibited a weaker association with their parents' socioeconomic status than less-educated respondents. Similar observations were made for US college graduates, compared to workers with a lower level of education or those with a graduate degree (Hout, 1988; Torche, 2011). Although valuable, these earlier Canadian contributions focus on the transmission of occupations, limiting the comparability with more recent research in economics. In addition, the studies focus on older birth cohorts, predominantly born before the 1960s.[2] Our study allows to consider younger birth cohorts for which postsecondary education is more common. We also investigate the role of field of study and degree type in the intergenerational transmission of income among Bachelor degree graduates.

Understanding the role of education in the intergenerational transmission of income also requires us to investigate the nature of the education-related skills that are rewarded on the labour market. Bowles and Gintis (2002) and Blanden et al. (2007) find that the role of cognitive and non-cognitive skills

groupings, often supplemented by distinctions based on the employment relationship of workers (self-employed/owner versus employee, for example). Occupation prestige scales ranking occupations based on their average wage level and educational attainment of incumbents are also common (Blau and Duncan, 1967; Blishen, 1967). The use of these measures is often justified by the fact that they provide permanent measures of socioeconomic status that vary little over one's life course after a certain age, and that capture more accurately the social rank of individuals than income alone.

[2]For example, the study by Wanner and Hayes (1996) uses those aged 25 to 64 in the 1986 General Social Survey, corresponding to the 1922-1961 birth cohorts. Only a small proportion of 15.9% of the respondents had a Bachelor's degree or more, and 51.7% had no post-secondary education (Wanner and Hayes, 1996, p.65). It is possible that Bachelors degree graduates formed a relatively homogenous group.

in intergenerational mobility largely overlaps with that of education. Using the LISA, we can speak specifically to the relationship between education and skills used at work with respect to intergenerational mobility; and between education and selection into jobs with desirable qualities (unionization, permanence and authority over other employees). Previous work has found that sons born from higher income fathers are more likely to be employed at the same firm as their father (Corak and Piraino, 2011; Kramarz and Skans, 2014; Lindquist et al., 2015). In addition, studies of hiring for "elite" positions in investment banking, management consulting and law show a desire to consider individuals with certain cultural, status, and network markers. These signal an upper middle-class background as being a "better fit" (Rivera, 2016), in addition to selection on credentials and skills. Finally, Ornstein (1981) also found that part of the intergenerational association between a father's occupation and the household income of their son operated through the first and current occupations of the child. Occupations are often conceptualized as a bundle of tasks, and job tasks have received a large amount of attention from labour economics over the last two decades (e.g. Autor et al., 2003). However, relatively little research focus on job characteristics related to the job task concept as a pathway for the intergenerational transmission of socioeconomic status.

2.3 Data

For this paper, we use the second wave of the Longitudinal and International Study of Adults (LISA). LISA is a biennial panel survey, first conducted in 2012 on a sample of Canadian households encompassing 23,926 respondents. Wave 2 was collected from January 2014 to June 2014. It counts with 19,178 respondents,[3] and covers topics including education and training, employment and job characteristics, and family background and life events. For every survey wave, respondents were linked to their T1 Family File (T1FF) records, from 1982 to the fiscal year preceding the survey collection year (see Hemeon, 2016 for details). The administrative linkage was also extended to include the T1FF records of all past and present family members, including spouses and common-law partners, parents and children, and siblings. Hence, we have detailed income data for both respondents and their family members, from 1982 to 2013. In this paper, we focus on the pairs formed by LISA respondents and their parents; that is, respondents correspond to the children in the intergenerational relationships.

The analysis we present in Section 2.6 focuses on the sample of LISA respondents aged 31 to 51 in 2014 (born between 1963 and 1983). The upper bound on age is imposed by the linkage methodology, and we chose the lower bound to minimize income measurement issues. Individuals were successfully linked with a parent only if both the respondent and their parent filed a T1 in the same fiscal year, in at least one year between 1982 and 2013; and the respondent and their parent resided at the same address (as reported in the T1 file) in a year they both filed a T1. As a result, the majority of LISA respondents were first linked to a parent between ages 15 and 21, when they entered the labour market, and the

[3]Of Wave 2 respondents, 16,895 were also respondents in Wave 1. The rest are either temporary sample members or original sample members who were non-respondents in Wave 1.

greatest number of parents were observed when respondents were 19 years old.[4] [5]

Research has identified two income measurement issues that bias intergenerational mobility estimates upwards (Atkinson et al., 1983; Solon, 1992; Zimmerman, 1992). First, errors-in-variables bias arises from using annual income or the average of a small number of annual income values instead of true permanent income. Second, life-cycle bias is introduced by measuring income for parents and children at different points in their life-cycles. Restricting the sample to those aged 31 to 51 in 2014 ensures that respondents were at most 19 years old in 1982,[6] constraining their parents' age, and at least 30 years old in 2013. To preserve a workable sample size, we do not directly limit parents' age. However, we use income ranks rather than levels for the analysis, which are less sensitive to both errors-in-variables and life-cycle biases (Chetty, Hendren, Kline, and Saez, 2014; Connolly et al., 2019; Corak, 2017). In this context, we proxy for parent permanent income by averaging a parent's income when their child is 15 to 19 years old, corresponding approximately to the time of first parental link. For children, we use average income between 2009 and 2013, to reflect income at the time respondent characteristics were measured.[7] In both cases, individuals are treated as having zero income in any given year their tax records are missing. However, they are dropped from the sample if their average over the five years used for the income measurement is under \$500.[8] In parts of the analysis where family income is used instead of individual income, exclusions are applied at the individual level before calculating family income, following Corak (2017). All values are CPI adjusted in 2013 constant Canadian dollars.

Approximately 78% of LISA respondents born between 1963 and 1983 were linked to a parent and satisfied the parental and own income restrictions. Overall, this yields a sample of 3,785 observations.[9] In results not shown here, we observed no difference in the total pre-tax parental income distributions

[4] Children may be linked to more than one parents. Parents in two child-parent pairs first observed on the same year form a couple (married or common-law), and are both identified as parents (for two child-parent pairs to be observed on the same year, both parents must reside at the same address). When only one parent is observed on the year the first child-parent pair is observed, we consider this parent to be a single parent. We assume that the first observed child-parent pair or pairs are the "real" parents of the children in our sample. Consistent with previous research, any parent subsequently paired with a child are discarded.

[5] The analysis is also restricted to respondents who were born in Canada or who arrived in Canada before age 16. Immigrants who arrived later are less likely to be observed filing their taxes with their parents, but also less likely to have immigrated with their parents. For the purpose of this study, it is also unclear what the intergenerational mobility of later-life immigrants would capture, given that their education and their parents' career would have taken place in large part outside of Canada.

[6] For respondents born between 1963 and 1966, who were over 15 years old in 1982, parental income is measured using the first five years of data. Their parental income is therefore measured when they were slightly older.

[7] Income is residualized on age to account for the fact that respondents from different birth cohorts have different ages in 2009-2013. This correction does not account for life cycle biases associated with changes in the ranking of children from families with different income levels (a lower intercepts and steeper age-income profiles of children of high-income parents, for example), which introduces a small downward bias in the estimate of intergenerational mobility.

[8] We do this to reflect the fact that individuals that are observed as having less than \$500 yearly income on average over a period of five years are likely to be a very unique group, one for which tax records may not draw a proper picture of economic situation.

[9] The sample size varies slightly with the measure of income used: it is lower when using individual income and/or employment income. The figure provided refers to our largest sample, where total family income is used for both parents and children.

across respondents included in the intergenerational sample, compared to those who weren't. This finding is in line with results from Corak and Heisz (1999) and Oreopoulos (2003), who worked with purely administrative data. Nonetheless, Oreopoulos et al. (2008) note that low-income adults were less likely to be included as children in the Intergenerational Income Database (IID), Canada's main administrative data set of parent-child pairs. In our own work, we also found that respondents who were successfully linked to a parent were more likely to have better-educated parents and to have grown up in a two-parent household. Our approach to the parent-child linkage is very similar to that used for the IID, which suggests that any selection associated with the intergenerational linkage in our sample also affects other studies of mobility in Canada. The impact of this selection on estimates of intergenerational income correlation is beyond the scope of this paper, and is left for further research. However, the differences in the family background of respondents included in our analysis may have implications for the interpretation of our results regarding education. We discuss those more carefully in Section 2.6.

Finally, Table 2.1 further shows the number of respondents observed for each birth year, as well as the corresponding linkage rate and average age at which parental income is measured. There is little variation in mean parent age by respondent birth year. For birth years 1966 to 1980, the linkage rate hovers around 80%. It is at its lowest at 66.7% for the 1963 birth year, and highest at 90.2% for the 1981 birth year. Despite their lower linkage rates, we include respondents born between 1963 and 1965 to reflect the cohorts used in existing work on intergenerational mobility in Canada. Furthermore, the difference in linkage rates is consistent with the fact that these respondents are over 16 years old in 1982. As such, we expect that selection for these groups is along the lines discussed above.

2.4 Child income and characteristics across the parental income distribution

2.4.1 Transition matrices

While Canada exhibits substantially more mobility than the United States, Corak (2013) shows that most of the difference comes from the bottom and top of the parental income distribution. These non-linearities are best seen in transition matrices, which show the probability for a respondent to be in a given income quintile conditional on their parents' income quintile. If there was no association between parental and child income, each cell would take a value of 20%. In Table 2.2, we present results for LISA respondents and their parents, using couple total income; that is, individual income, plus spousal income when present. Results using couple market and employment income follow similar patterns.

Our results show that 32.9% of children born from top-quintile parents also end up in the top income quintile. This is three times larger than the percentage of children born from bottom-quintile parents who reach the top, 10.9%. We find an almost mirror image of these conditional probabilities among children in the bottom income quintile, with 10.9% of children born in top-quintile families experiencing downward mobility to the bottom quintile, and 32.1% of children born in bottom-quintile

Table 2.1: Sample size, linkage rate and mean parent age by respondent birth year

Birth year	Sample size	Linkage rate (%)	Mean parent age
1963	236	66.7	49.0
1964	234	69.9	48.3
1965	227	67.4	47.4
1966	206	75.7	45.7
1967	206	83.7	44.9
1968	222	79.6	45.7
1969	185	79.4	44.2
1970	179	78.5	44.9
1971	199	80.6	45.6
1972	185	75.2	45.0
1973	156	78.0	44.9
1974	180	83.7	44.7
1975	157	77.7	46.1
1976	145	80.1	45.8
1977	132	81.0	44.1
1978	155	82.0	44.7
1979	162	81.4	46.8
1980	158	82.7	46.0
1981	157	90.2	44.6
1982	163	87.2	45.5
1983	141	87.6	44.4
Total	3,785	78.2	45.7

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013. Sample size refers to the number of respondents used for the analysis. Linkage rate refers to the percentage of respondents born in each year that are included in the analysis.

families remaining in the bottom quintile. These results are consistent with Corak (2017), who finds that 32.3% of children born in the top quintile remain there, 30.1% of children born in the bottom quintile remain there, and 11.4% of children born in the bottom quintile reach the top quintile in adulthood (downward mobility from top to bottom is not reported).

Overall, 46.9% of children born from parents in the bottom quintile experienced an upward mobility of more than one quintile. Coincidently, 43.3% of children born in the top quintile experienced a downward mobility of more than one quintile. These results suggest that social mobility out of bottom and top parental income quintiles is unlikely, but not exceptional. Children born in the middle of the parental total family income distribution are much more likely to experience mobility. Those born in the middle quintile experience almost perfect mobility (around 20%) per cell. Those born in the second to last quintile have low probabilities of reaching the top two quintiles, and those born in the fourth quintile have low probabilities of falling into the bottom quintile. The conditional probability of

Table 2.2: Child family income quintile conditional on parental family income quintile

Parental family income quintile	Child family income quintile Bottom	Q2	Q3	Q4	Top	Total
	Total income					
Bottom	32.1	21.0	19.1	17.0	10.9	100.1
Q2	22.4	22.7	23.3	16.7	14.9	100.0
Q3	20.7	19.2	20.6	20.8	18.7	100.0
Q4	14.6	19.4	21.6	21.7	22.7	100.0
Top	10.9	17.0	15.4	23.8	32.9	100.0
	Market income					
Bottom	30.1	21.8	20.8	15.5	11.8	100.0
Q2	23.1	20.5	23.8	17.4	15.2	100.0
Q3	18.9	22.0	19.3	21.8	18.0	100.0
Q4	18.0	18.6	19.5	21.1	22.8	100.0
Top	10.7	16.9	16.4	23.9	32.1	100.0
	Employment income					
Bottom	29.8	19.9	21.1	16.0	13.3	100.1
Q2	24.2	19.6	23.1	18.7	14.4	100.0
Q3	17.7	23.4	19.8	19.8	19.3	100.0
Q4	16.4	21.4	20.5	20.6	21.2	100.1
Top	12.0	16.0	15.3	25.1	31.6	100.0

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

reaching the other quintiles is relatively even.

2.4.2 Characteristics of respondents by their parental income quintile

Table 2.2 confirms that much of the intergenerational correlation in incomes stems from the bottom and the top of the parental income distribution. Correspondingly, it is worth looking at the characteristics of children born to those segments of the distribution, to get a sense of where the mechanisms may lie. Table 2.3 presents cross-tabulations of parental income quintile with the characteristics of their children, including education, marital status, employment, and job characteristics. This is a first step towards understanding which factors may contribute to the intergenerational transmission of income.

There is a clear relationship between parental income quintile and child education. Education increases with parental income quintile, with a particularly large portion of respondents in the top quintile reporting a bachelor's degree or more: only 22.6% of respondents born to bottom-quintile parents obtained a university degree, compared to 27.8% for those born to the middle quintile, and 59.3% for those from the top of the distribution, with 22.8% having obtained a graduate degree. Note that the probability of completing graduate studies, conditional on obtaining a Bachelor degree, is higher in the top three quintiles. This is reported in the column of Table 2.3 showing the percentage

Table 2.3: Child education by parental family income quintile

Parent family income quintile	Lt HS	High school	Trade	Some postsec	Bachelor	Graduate	% Bachelor w/ Graduate
Bottom	13.1	23.4	13.9	27.0	15.3	7.3	32.3
Q2	5.4	23.0	15.8	28.0	21.1	6.7	24.1
Q3	8.2	19.0	15.3	29.7	17.2	10.6	38.1
Q4	3.6	17.1	10.1	33.4	23.2	12.7	35.4
Top	2.1	9.5	5.0	24.0	36.5	22.8	38.4
Total	6.5	18.4	12.0	28.4	22.6	12.0	34. 7

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

with a graduate degree. Meanwhile, children born in the top quintile are much less likely to have a trade or vocational certificate, or some postsecondary education below Bachelor.[10]

We now turn our analysis to the variation in job characteristics among employed respondents. Job characteristics variables are divided into two categories, job skills used at work and indicators of job quality. First, we used the LISA questions on the use of skills at work to create six skill use indices for reading, writing, math, communication, dexterity and physical strength. The LISA collects detailed information about skills used at work. Respondents are asked how important the use of certain skills is for their work, on a scale from one (not important) to five (extremely important). Those who gave an answer above one are also asked about the complexity of that skill ("What level of [skill] is needed to perform your [current] job?"), on a scale from one to seven. We develop a skill intensity index combining those two measures by recording the first question on a scale from zero to four and by multiplying the values of the two variables together.[11]

In Figure 2.1, we plot mean job skill intensity index values against parental family total income deciles (results are similar for employment income). There is a good level of variation in job skill intensity across the parental income distribution. The skills with the highest intensity on average are reading and communication, with a similar distribution of mean values by parental income deciles. The average level of writing skill intensity is lower, but it also follows a similar distribution as reading and communication. For those three skills, we find that respondents in the bottom decile have lower values, skill intensity is relatively constant across the second to eighth decile, with only a slight monotonic increase between them, and that skill intensity is much higher for those in the top two deciles. The

[10]The probabilities of being employed and of being married or cohabiting conditional on parental income quintile are also reported, in Table B.1. Those born in higher parental income quintiles are more likely to be employed (a difference of 13.3 percentage points with those born in the bottom quintile) and to be married or cohabiting (a difference of 13.4 percentage points with those born in the bottom quintile). Overall, Table 2.3 and Table B.1 show a great level of variation in sociodemographic characteristics among children born in different parental income quintiles.

[11]This yields a higher skill intensity score for those who either use skills frequently or use complex skills. The highest skill intensity scores are attributed to those who use skills frequently and with a high level of complexity. This mirrors the approach used in the Dictionary of Occupational Titles (DOT) and the O*Net, two occupation-level US datasets that rate the importance or intensity in the use of certain occupational skills (see Autor et al., 2003; Autor and Handel, 2013). For information on the structure of the index, see Table B.2 reporting all possible values.

Figure 2.1: Job skill intensity by parental total family income decile

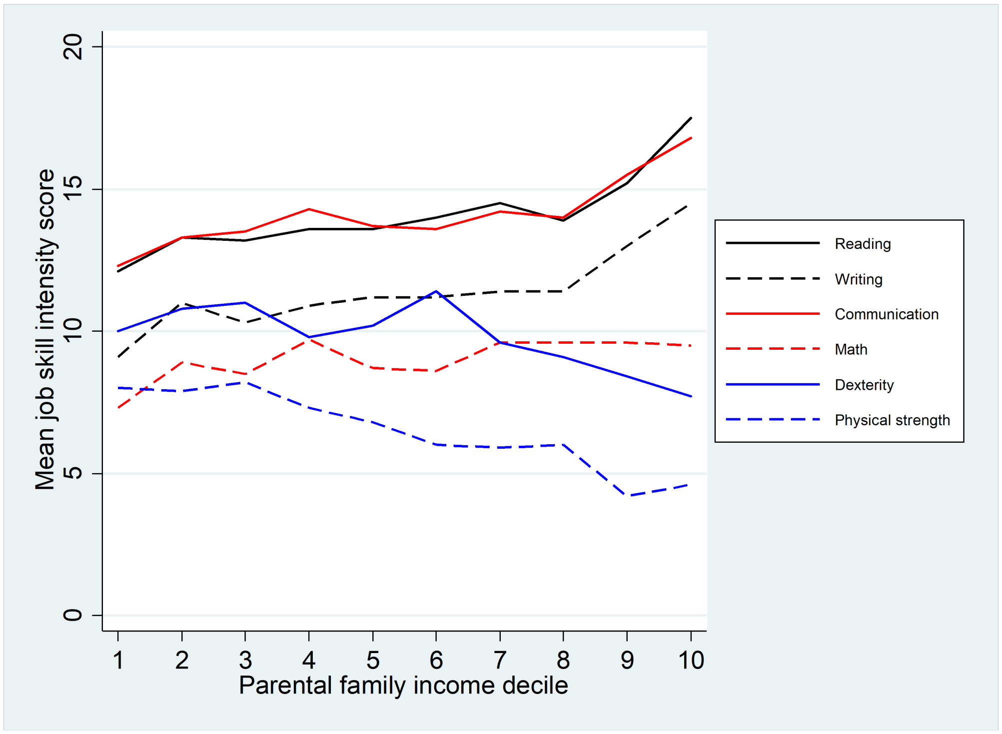

distribution of mean skill intensity values by parental income decile is similar to other skills at the bottom of the parental income distribution (but at a lower level). However, mean values are similar for those born across the rest of the parental income distribution. Finally, dexterity skill intensity and physical strength intensity are both much higher for those in the bottom parental income deciles than other deciles, with a greater difference between the bottom and the top for physical strength. Overall, physical strength is the least intensely used skill at work in our sample, especially among those born from high income parents.

Second, we use three measures of job quality: unionization (dummy), contract type (permanent, temporary, or self-employed), and supervision of employees, either as a manager/supervision or as an employer (dummy). Table 2.4 reports the conditional probability of job quality indicators. While there appears to be a relationship between parental total income or employment income and whether respondents are employed under a permanent contract (as opposed to having a non-permanent contract or being self-employed), the probability of being a union member or supervising employees varies little across parental income quintiles.

2.5 Methodology

We use rank-rank regressions to formally investigate the role of children's education in the intergenerational transmission of income. Let i denote individual respondents, such that $i = 1, \ldots, N$; $y_{i,c}$, the

Table 2.4: Job quality by parental family income quintile, employed respondents

Parent family income quintile	Permanent contract	Union member	Supervision
		Total income	
Bottom	65.9	31.6	32.2
Q2	66.0	30.1	32.9
Q3	70.3	30.0	38.5
Q4	72.9	35.7	31.9
Top	72.2	31.8	36.1
Total	69.5	31.9	34.4
		Employment income	
Bottom	66.4	31.8	36.3
Q2	65.2	30.8	31.3
Q3	73.0	30.7	37.9
Q4	72.0	35.8	34.8
Top	74.9	33.0	35.1
Total	70.4	32.5	35.1

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

respondent's ranking in the income distribution; and $y_{i,p}$, their parent's income rank. The rank-rank specification is

$$y_{i,c} = \alpha + \rho y_{i,p} + \varepsilon_{i,c} \tag{2.1}$$

where α and ρ respectively denote absolute and relative mobility. Specifically, α corresponds to the average percentile rank of children born to parents from the very bottom of the parental income distribution. As for ρ, it documents the correlation between parent and child income percentile ranks. This specification presumes nothing with respect to the direction of causality, but merely captures the relationship between the respective income rankings of children and their parents.

2.5.1 Estimating the role of education in the intergenerational transmission of income

To better understand the role education plays in the intergeneration transmission of income, we follow Blanden et al. (2007) and decompose ρ into a part that is associated with education, and a part that isn't. As outlined in Section 2.2, we expect the association between parental and child income to arise in part in relation to children's education, either because higher-income parents can invest directly in their children's education, or because they have abilities that are transmitted to children and contribute to their educational attainment. Denote $H_{i,c}$, child educational attainment, and consider two reduced-form equations, one linking parental income to child education and the second one, child education to child income:

$$H_{i,c} = \pi_0 + \pi_1 y_{i,p} + \upsilon_{i,c} \tag{2.2}$$

$$y_{i,c} = \beta_0 + \beta_1 H_{i,c} + \nu_{i,c} \tag{2.3}$$

Because Equation 2.1 is a univariate regression, the relative measure of mobility, ρ, can easily be decomposed into a portion that is associated with education through the two-link channel embodied by Equation 2.2 and Equation 2.3, and a portion that is orthogonal to educational attainment:

$$\begin{aligned} \rho &= \frac{Cov(y_{i,c}, y_{i,p})}{Var(y_{i,p})} \\ &= \frac{Cov(\beta_0 + \beta_1 H_{i,c} + \nu_{i,c}, y_{i,p})}{Var(y_{i,p})} \\ &= \frac{\beta_1 Cov(H_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\nu_{i,c}, y_{i,p})}{Var(y_{i,p})} \\ &= \beta_1 \pi_1 + \mu \end{aligned} \tag{2.4}$$

Correspondingly, we can estimate the product $\beta_1 \pi_1$ and divide it by the estimate of ρ to obtain the proportion of the relative measure of mobility that is attributable to education, through the combination of the association between parental income and child educational attainment and the association between child education and income.

In Section 2.6, we present decomposition results for a continuous and a binary measure of education, respectively years of education and a dummy variable equal to one for respondents who've obtained a Bachelor's degree. We also investigate a categorical measure, highest degree completed. A categorical variable for educational attainment might better capture meaningful differences in human capital and credentials. In that case, the linear reduced-form equation for educational attainment, Equation 2.2, is hard to justify. For this part of the analysis, we adopt an alternative approach to the decomposition. It corresponds to the method used by Blanden et al. (2013) to understand differences in the results obtained by economists and sociologists looking at intergenerational mobility in the UK using income and social class respectively (see also Rothstein, 2019). Let $\hat{y}^{Ed}_{i,c}$, the value of child income predicted by child education; that is,

$$\begin{aligned} y_{i,c} &= \hat{\beta}_0 + \hat{\beta}_1 H_{i,c} + \hat{\nu}_{i,c} \\ &= \hat{y}^{Ed}_{i,c} + \hat{\nu}_{i,c} \end{aligned}$$

Then,

$$\begin{aligned}
\rho &= \frac{Cov(y_{i,c}, y_{i,p})}{Var(y_{i,p})} \\
&= \frac{Cov(\hat{y}^{Ed}_{i,c} + \hat{\nu}_{i,c}, y_{i,p})}{Var(y_{i,p})} \\
&= \frac{Cov(\hat{y}^{Ed}_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\nu}_{i,c}, y_{i,p})}{Var(y_{i,p})} \\
&= \lambda_1 + \eta
\end{aligned}$$

where λ_1 is estimated by regressing $\hat{y}^{Ed}_{i,c}$ on $y_{i,p}$:

$$\hat{y}^{Ed}_{i,c} = \lambda_0 + \lambda_1 y_{i,p} + \varepsilon_{i,c}$$

When using the linear reduced-form specification for educational attainment, Equation 2.2, the two approaches yield the same result. We also extend the approach to understand the relationship between mobility, education, and job characteristics. Denote S_i, a vector of indices characterizing the intensity of different skills respondents report using at work, and Q_i, a vector of job quality indicators (unionization, permanence, authority). Consider a reduced-form equation linking child income to these job characteristics:

$$y_{i,c} = \zeta + S_i\delta + Q_i\theta + \vartheta_{i,c} \tag{2.5}$$

and denote the partial predicted values of income, $\hat{y}^S_{i,c} = S_i\hat{\delta}$ and $\hat{y}_{i,c} = Q_i\hat{\theta}$. Then, the relative measure of mobility can be expressed as follows:

$$\begin{aligned}
\rho &= \frac{Cov(y_{i,c}, y_{i,p})}{Var(y_{i,p})} \\
&= \frac{Cov(\hat{y}^S_{i,c} + \hat{y}^Q_{i,c} + \hat{\vartheta}_{i,c}, y_{i,p})}{Var(y_{i,p})} \\
&= \frac{Cov(\hat{y}^S_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{y}^Q_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\vartheta}^S_{i,c}, y_{i,p})}{Var(y_{i,p})}
\end{aligned}$$

This provides a similar decomposition as discussed earlier for education: it indicates the portion of ρ that comes through the transmission of skills from parent to child that are associated to job skills, job quality, and to neither type of job characteristics. In turn, it is possible to relate this back to education-related skills. To do so, we regress each income part on education:

$$\begin{aligned}
\hat{y}^S_{i,c} &= \gamma^S_0 + \gamma^S_1 H_i + \psi^S_{i,c} \\
\hat{y}^Q_{i,c} &= \gamma^Q_0 + \gamma^Q_1 H_i + \psi^Q_{i,c} \\
\hat{y}^\vartheta_{i,c} &= \gamma^\vartheta_0 + \gamma^\vartheta_1 H_i + \psi^\vartheta_{i,c}
\end{aligned}$$

Again, denote $\widehat{y^S_{i,c}} = \hat{\gamma}^S_0 + \hat{\gamma}^S_1 H_i$, $\widehat{y^Q_{i,c}} = \hat{\gamma}^Q_0 + \hat{\gamma}^Q_1 H_i$ and $\widehat{y^\vartheta_{i,c}} = \hat{\gamma}^\vartheta_0 + \hat{\gamma}^\vartheta_1 H_i$. This yields the final decomposition:

$$\rho = \frac{Cov(\widehat{y^S_{i,c}}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\psi}^S_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\widehat{y^Q_{i,c}}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\psi}^Q_{i,c}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\widehat{y^\vartheta_{i,c}}, y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\psi}^\vartheta_{i,c}, y_{i,p})}{Var(y_{i,p})} \tag{2.6}$$

Hence, although we cannot characterize the direction of the causality, the six components of Equation 2.6 allow us to speak broadly to six mechanisms that are involved in the intergenerational transmission of income:

1. $\frac{Cov(\widehat{y^S_{i,c}}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income and with both education and job skill intensity (net of job quality);

2. $\frac{Cov(\hat{\psi}^S_{i,c}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income and with job skill intensity (net of job quality), but not education;

3. $\frac{Cov(\widehat{y^Q_{i,c}}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income and with both education and job quality (net of job skill intensity);

4. $\frac{Cov(\hat{\psi}^Q_{i,c}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income and with job quality (net of job skill intensity), but not education;

5. $\frac{Cov(\widehat{y^\vartheta_{i,c}}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income and with education, but not job characteristics;

6. $\frac{Cov(\hat{\psi}^\vartheta_{i,c}, y_{i,p})}{Var(y_{i,p})}$: the part of child income correlated with parental income, independently of education and job characteristics.

2.6 Results

2.6.1 The role of education in the intergenerational transmission of income

In Table 2.5, we present the rank-rank estimates of absolute and relative mobility, α and ρ, and the results of the intermediate regressions corresponding to Equation 2.2 and Equation 2.3, π_1 and β_1. We include results using total income, market income and employment income; and for each income source, using parental and respondent couple income, parental couple income and respondent individual income, and parental and respondent individual income. The top panel of Table 2.5 uses years of education, and the bottom panel, a dummy variable equal to one if the respondent has a Bachelor degree.

Table 2.5: Decomposition of relative intergenerational mobility

Income definition Parent	Child	α	ρ	π	β	$\pi\beta$	$\pi\beta/\rho$ (%)
				Years of education			
Total income							
couple	couple	36.98	0.268	0.034	2.688	0.091	34.1
couple	indiv	39.27	0.223	0.034	2.817	0.096	42.9
indiv	indiv	40.26	0.203	0.032	2.817	0.090	44.2
Market income							
couple	couple	37.94	0.249	0.032	2.550	0.080	32.4
couple	indiv	40.41	0.200	0.032	2.584	0.081	40.7
indiv	indiv	41.60	0.176	0.030	2.584	0.078	44.4
Employment income							
couple	couple	38.76	0.232	0.031	2.406	0.075	32.1
couple	indiv	40.83	0.192	0.031	2.392	0.074	38.7
indiv	indiv	42.45	0.160	0.030	2.392	0.071	44.6
				Bachelor Degree			
Total income							
couple	couple	36.76	0.273	0.004	19.76	0.081	29.7
couple	indiv	39.04	0.229	0.004	20.08	0.082	36.1
indiv	indiv	40.19	0.206	0.004	20.08	0.076	37.1
Market income							
couple	couple	37.87	0.249	0.004	19.18	0.076	30.7
couple	indiv	40.32	0.201	0.004	18.79	0.075	37.2
indiv	indiv	41.52	0.177	0.004	18.79	0.070	39.5
Employment income							
couple	couple	38.66	0.233	0.004	18.31	0.071	30.4
couple	indiv	40.71	0.193	0.004	17.59	0.068	35.3
indiv	indiv	42.34	0.161	0.004	17.59	0.063	39.3

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

Before turning to the role of education, it is worth noting two things with respect to the baseline mobility estimates. First, the association is stronger for total income than for market income, and weakest for employment income, in line with previous findings from the Canadian literature using definitions of income similar to ours (eg. Corak and Heisz (1999)). On the one hand, using employment income rather than total income excludes respondents from the bottom of the parental income distribution, where intergenerational correlation is relatively high (Corak, 2013). On the other hand, it also excludes a relatively larger portion of the income of higher-income respondents and their parents, who also exhibit greater intergenerational correlation. Second, irrespective of income source, the specifications using couple income for both parents and children show the strongest association: from 0.232 to 0.268 when using years of education. Conversely, using individual income for (highest-earning) parents and children produces the weakest association, from 0.160 to 0.203. Estimates using the couple income of parents and the individual income of children have a value of 0.192 to 0.223. This suggests an important role for marital status and assortative mating in the intergenerational transmission of income.[12]

The last column of Table 2.5 shows that education accounts for at least 30% of relative intergenerational mobility. Using graduation from a Bachelor degree program instead of years of education allows to explain only slightly less of the association, suggesting that this education credential plays a large role in the intergenerational transmission of income. The return to a Bachelor degree is indeed substantial, at around 20 percentiles compared to the baseline category of those who obtained some postsecondary education or less.[13]

For both variables, education accounts for a larger percentage of the association between parental and child income in specifications using individual child income. This suggests one or both of two things. First, that there is positive correlation between a respondent's parental income and their spouse's education, even net of the assortative matching on education between spouses. Second, that the part of spousal income that is orthogonal to their education is itself positively correlated with the respondent's parental income. The relative importance of education is even greater in specifications using the individual income of the highest-earning parent. Interestingly however, the proportion of relative mobility associated with education is similar across specifications using different income sources (total, market, and employment income).

In Table 2.6, we report results from the projection decomposition method, using a categorical variable that records the highest level of schooling attained by children. We also include results using years of education and the Bachelor degree dummy, for comparison. Education is associated with a larger proportion of relative intergenerational mobility when measured by children's highest completed degree or certificate, confirming that it picks up relevant differences between individuals. In this case, education accounts for 40 to 50% of the relationship between parental and child income ranks. The differences by source of income and income definition (individual or couple income) are similar to those we observed in Table 2.5. We rely on the projection approach for the decomposition results presented in the rest of the paper.

[12] In a decomposition exercise reported in Table B.3, we estimate the share of the association explained by marital status, net of education and employment status. This is done using a three-step approach as described in the methods section, substituting marital and employment status for job characteristics. The mediating role of marital status is larger in specifications using a measure of child income including spouse income (when present), than in specifications using child individual income. Marital status accounts for up to 16.8% of the association between parental and child income in specifications that use family income (sum of spouses income, if a spouse is present) as a measure of child income. The effect of marital status shows little correlation with education. This suggests differences in marriage decisions and outcomes between children born from parents with different income levels. These differences appear to be uncorrelated with child educational attainment, for the most part.

[13] The small differences in the estimates of absolute and relative mobility presented in the top and bottom panels of Table 2.5 stem from minor sample differences.

Table 2.6: Decomposition of relative intergenerational mobility by projection decomposition method

Income definition			Years of education		Bachelor Degree		Highest degree	
Parent	Child	ρ	λ	$\lambda/\rho(\%)$	λ	$\lambda/\rho(\%)$	λ	$\lambda/\rho(\%)$
Total income								
couple	couple	0.273	0.092	33.6	0.081	29.7	0.111	40.5
couple	indiv	0.229	0.096	42.1	0.082	36.1	0.110	48.1
indiv	indiv	0.206	0.090	43.7	0.076	37.1	0.103	50.1
Market income								
couple	couple	0.249	0.080	32.3	0.076	30.7	0.095	38.2
couple	indiv	0.201	0.081	40.7	0.075	37.2	0.090	44.8
indiv	indiv	0.177	0.078	44.3	0.070	39.5	0.085	48.1
Employment income								
couple	couple	0.233	0.075	32.0	0.071	30.4	0.090	38.6
couple	indiv	0.193	0.074	38.6	0.068	35.3	0.082	42.6
indiv	indiv	0.161	0.071	44.6	0.063	39.3	0.078	48.2

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

2.6.2 Horizontal inequality and differentiation by parental background among higher education graduates

The LISA data set allows us to further evaluate the role of education by considering horizontal educational stratification among university graduates. While vertical stratification corresponds to variations in the level of education by parental background, horizontal stratification is the sorting of children with the same level of education into fields of study associated with different labour market rewards depending on their parental background (Lucas, 2001; Torche, 2011). This has been identified as a possible mechanism of socioeconomic status reproduction in an era of mass participation to university education in Canada and the US. For instance, Zarifa (2012a) finds a strong association (direct and indirect) between parental background (income and education) and the choice of field of study in Canada. Meanwhile, the sorting into graduate school among Bachelor degree graduates is another way individuals with more privileged socioeconomic backgrounds differentiate themselves from other Bachelor degree graduates (Mullen et al., 2003; Zarifa, 2012b). These empirical studies focus exclusively on the relationship between parental socioeconomic status and child educational decisions, not the impact of these decisions on their socioeconomic status attainment.

We use the projection decomposition approach to investigate horizontal educational stratification among LISA respondents. For the sample of university graduates, we compute predicted income given degree type, field of study, or both. We then regress predicted income on parental income to obtain an estimate of λ_1. Degree type includes Bachelors degree (reference category), professional degrees, Master degrees, law degrees (JD), medical degrees, and PhDs. The Wave 2 LISA data does not include information about the institution respondents attended. In Table 2.7, we report the estimate of λ_1 for the subgroup of LISA respondents who hold a Bachelor degree, and the percentage of ρ that is associated with degree type and/or field of study, $(\lambda_1/\rho \times 100)$.

Table 2.7: Decomposition of relative intergenerational mobility among university graduates

Income definition		Bachelor and above				Bachelor degree			Graduate degree			
			% explained				% explained			% explained		
Parent	Child	Total effect	Degree type	Degree type and FoS	n	Total effect	FoS	n	Total effect	Degree type	Degree type and FoS	n
Total income												
couple	couple	0.230	4.8	7.7		0.196	5.4		0.275	2.5	14.5	
couple	indiv	0.142	7.7	11.2		0.130	4.3		0.145	5.6	27.6	
indiv	indiv	0.132	7.8	8.5	2,174,100	0.110	-3.2	1,422,700	0.154	6.3	29.8	751,300
Market income												
couple	couple	0.221	6.7	9.4		0.199	6.6		0.237	3.5	8.6	
couple	indiv	0.144	9.2	11.4		0.135	3.5		0.132	4.9	17.5	
indiv	indiv	0.129	9.6	8.9	2,140,800	0.112	-3.2	1,409,100	0.138	6.2	19.7	731,800
Employment income												
couple	couple	0.199	6.3	9.8		0.202	7.5		0.165	2.2	13.7	
couple	indiv	0.130	6.8	12.1		0.147	8.9		0.072	-2.1	31.7	
indiv	indiv	0.101	7.9	8.5	2,074,500	0.113	2.9	1,365,000	0.058	1.3	37.4	709,500

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

Among all university graduates, degree type alone accounts for 5 to 10% of relative intergenerational mobility, depending on the specification. Together, degree type and field of study account for 8 to 12% of the mobility estimate. We further split the sample into two groups: those whose highest educational attainment is the Bachelors degree, and those who also completed graduate studies. In the first group, field of study accounts for 4 to 9% of relative intergenerational mobility in specifications that use couple parental income.[14] Field of study also appears to play a substantial role among respondents with graduate degrees. Degree type alone accounts at most for approximately 6% of relative intergenerational mobility in that group, whereas degree type and field of study together account for 8 to 37% of mobility. This may be explained by the higher proportion of children from higher income parents graduating with MBAs or other expensive business degrees, compared to children from lower income parents. While we can't speak to the mechanisms that underlie these estimates, they do point to the value of further research into the role played by field of study and degree type, and how it relates with family structure.

These results are significant in the Canadian context for two reasons. First, they suggest that other factors such as selection into institutions with different reputations (elite universities, for example) and labour market outcomes may play a greater role in accounting for the association between parental and child income than selection into different fields of study or into graduate school among Bachelor degree graduates. Second, although the role played by field of study among university graduates is economically small, it nonetheless appears to be a channel through which mobility may be achieved.

2.6.3 Intergenerational mobility across groups

Intergenerational mobility by highest educational attainment

One third to one half of relative mobility is associated with differences in the educational attainment of children from different parental backgrounds. However, educational attainment may interact with parental income in a way that is not captured in Section 2.6.1 and Section 2.6.2. For example, results for Bachelor degree graduates in Table 2.7 reports smaller relative mobility coefficients than those for the overall population in Table 2.6. Another way of investigating the role of education in the intergenerational transmission of income is therefore to obtain group-specific estimates of absolute and relative mobility. As mentioned in Section 2.2, previous work in the United States and in Canada has shown less correlation in socio-economic status among university graduates, suggesting either that education acts as an equalizing force or that more mobile children select into higher education, or both (Hout, 1988; Torche, 2011; for Canada, see Wanner and Hayes, 1996).

To explore this, we regress respondent income rank on parental income rank, highest educational

[14] Interestingly, the estimate of λ_1 is negative in specifications that use individual parental income, where income is defined as total or market income. Hence, while couple parental income is associated with selection into higher-return fields of study, the opposite is true for individual parental income: among respondents who obtain a Bachelor's degree, those from the bottom of the parental income distribution are more likely to have studied in higher-paying fields. However, note that this effect is economically extremely small (-3.2% of the 0.11 relative mobility coefficient).

attainment, and the corresponding interaction. Respondents who haven't graduated from high school constitute the reference group. A negative interaction term between parental income and a specific level of child education indicates that the association between parental and child income is weaker and intergenerational mobility is higher among those with that given level of education.

We report results of specifications using different income sources and income definitions in Table 2.8. Importantly, none of the interaction terms are statistically significant at the 95% level. This suggests no difference in intergenerational mobility across groups. Parental income is just as strongly correlated with child income among high school drop-outs as it is among university graduates. This implies that university education does not play an equalizing role in Canada, or that selection into education is not linked with mobility, or a combination of both. As expected however, the absolute size of the interaction coefficients for some postsecondary education is large and borders significance in most specifications.

In Section 2.3, we noted that samples of parent-child pairs constructed using personal income tax records may under-represent children of low-educated parents, and children who didn't grow up with both of their parents. Given the wide use of administrative data in studies of intergenerational mobility, this is unlikely to affect the comparison between Canada and other countries. However, the selection may affect any given set of estimates, depending on how it operates. Suppose for instance that all high-income parents are captured, but that we only see better-educated low-income parents. This would be the case if less-educated parents are less responsive to tax-filing incentives for low-income people. In this case, we expect the baseline estimate of relative intergenerational mobility to be biased downward. Furthermore, if parental and child education levels are positively correlated, this form of selection would over-state the effect of education for children of low-income parents, thus under-stating the difference in the effect of education along the parental income distribution. Further work is needed to correct for selection and confirm that respondents with differing levels of education share a similar experience of intergenerational mobility.

Intergenerational mobility by sex, visible minority status and immigration background

As a comparison, we investigate whether absolute and relative intergenerational mobility differ across groups defined by sex, visible minority status and immigration background. We present the respective results in the first, second, and third panels of Table 2.9, using total income. A respondent is considered to have an immigration background if they were born outside of Canada or if at least one of their parents was. Note that in our sample, selected immigrants are likely to have a specific set of characteristics because only those who arrived in Canada before age 15 are included. They are therefore more likely to have completed at least part of their education in Canada and to have Canadian education credentials. This is also likely to affect the composition of members of visible minority groups in our sample to the extent that both statuses are sometimes related.

Using couple income for both respondents and their parents, we find no statistically significant difference in the intergenerational transmission of income between men and women. Conversely, absolute

Table 2.8: Varying relative intergenerational mobility level by highest degree obtained

	Parent and child couple income	Parent couple and child indiv. income	Parent and child individual income
Total income			
Parental income (ParInc)	0.252 *	0.235 *	0.147
Less than HS (reference)			
High school or equivalent	6.356	8.123	7.385
Vocational/Apprenticeship	15.496 *	16.342 *	13.589 *
Some postsecondary	23.399 *	24.174 *	21.132 *
Bachelor Degree	27.538 *	31.722 *	29.967 *
Graduate Degree	27.991 *	36.283 *	32.929 *
ParInc × High school or equivalent	-0.037	-0.049	-0.016
ParInc × Vocational/Apprenticeship	-0.097	-0.070	0.012
ParInc × Some postsecondary	-0.164	-0.179	-0.090
ParInc × Bachelor Degree	-0.057	-0.105	-0.038
ParInc × Graduate Degree	0.020	-0.092	0.005
Constant	22.317 *	22.254 *	25.236 *
Employment income			
Parental income (ParInc)	0.187 *	0.167	0.091
Less than HS (reference)			
High school or equivalent	0.669	0.923	0.966
Vocational/Apprenticeship	4.779	6.832	4.418
Some postsecondary	14.741 *	13.987 *	11.451 *
Bachelor Degree	17.837 *	20.284 *	19.291 *
Graduate degree	26.846 *	30.632 *	28.610 *
ParInc × High school or equivalent	-0.007	0.014	0.019
ParInc × Vocational/Apprenticeship	0.001	0.028	0.088
ParInc × Some postsecondary	-0.110	-0.114	-0.047
ParInc × Bachelor Degree	0.014	-0.021	0.022
ParInc × Graduate Degree	-0.025	-0.097	-0.034
Constant	30.082 *	30.763 *	33.753 *

* Significant at the 95 % level. Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

and relative mobility are both lower for women when their individual income is used. Women earn less than men on average (4.98 percentile points), and the gender income gap is larger among individuals born from higher income parents: it grows by 0.11 percentile point for each one-percentile-point increase in parental income. In other words, relative intergenerational mobility is greater among women, but around a lower average income ranking. Alternatively, these results suggest that women achieve the same level of intergenerational transmission of income than men, but that marriage and couple formation plays a greater role in this process than personal income level in comparison with men.

Visible minority status is also associated with much larger relative mobility, whether couple or individual child income is used: members of visible minorities experience a relative mobility of 0.05,

Table 2.9: Group differences in relative and absolute intergenerational mobility, total income

	Parent and child couple income		Parent couple and child individual income	
	(1)	(2)	(1)	(2)
Sex				
Parental income	0.28 *	0.19 *	0.28 *	0.17 *
Women	2.96	-0.95	-4.98 *	-9.60 *
Interaction	-0.04	-0.01	-0.11 *	-0.08 *
Education		Yes		Yes
Constant	35.68 *	24.89 *	41.15 *	27.07 *
Visible minority status				
Parental income	0.29 *	0.19 *	0.24 *	0.14 *
Visible minority	9.36	2.86	13.38 *	6.70
Interaction	-0.24 *	-0.19 *	-0.22 *	-0.17
Education		Yes		Yes
Constant	36.20 *	24.46 *	37.95 *	25.24 *
Immigration background				
Parental income	0.28 *	0.17 *	0.23 *	0.12 *
Immigrant parents/self	2.07	-3.58	4.37	-1.33
Interaction	-0.04	0.02	-0.04	0.03
Education		Yes		Yes
Constant	36.34 *	25.59 *	37.91 *	26.13 *

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

compared to 0.29 among the rest of our sample (0.02 vs. 0.24, using individual child income). There is also a large and statistically significant difference in absolute mobility when individual child income is used: members of visible minorities are 13.38 percentile points above non-members on average. Hence, members of visible minority groups have a greater level of mobility around a higher average income. As a result, low-parental-income members of visible minority groups do better than comparable non-members, while high-parental-income members fare worse than non-members. This suggests greater barriers in the reproduction of a high socioeconomic status across generations for members of visible minority groups than for the rest of the population, but that visible minority children of lower income parents can expect opportunities for upward mobility. Interestingly, the third panel of Table 2.9 suggests that these results are not driven by immigration background.

Finally, in the second and fourth columns of Table 2.9, we show the same results, with controls for highest educational attainment. As expected, this decreases the association between parental and child incomes for all groups and all specifications.

2.6.4 The mediating role of job characteristics among employed individuals

Understanding the role education plays in the intergenerational transmission of income requires understanding the skills that are leveraged or acquired through education (and rewarded on the labour market), and the skills that aren't. To do so, we use a sample of LISA respondents who were employed during the Wave 2 reference period, and for whom we have information on a number of job characteristics. For these respondents, we decompose relative intergenerational mobility into components that are associated with education and/or job characteristics, following Equation 2.6. Recall that we use two sets of job characteristics: skills used at work (reading, writing, communication, mathematics, physical strength, and dexterity), and indicators of job quality (unionization status, permanence, and authority over other employees). We restrict our analysis to employment income, although the results show little difference with other income sources.[15] We report results in Table 2.10. Note that relative mobility for the employed sample is slightly smaller than for the full sample. In line with previous results, using couple income yields larger estimates of ρ, or lower mobility. Again, this is consistent with the matching of higher-earning respondents with high-earning spouses.[16]

[15]In a three-step decomposition of the relative intergenerational mobility coefficient for the full sample reported in Table B.3, employment status accounts for 7.4 to 19.2% of the total association between parental and child income, net of marital status. In the case of total income, this mediating effect is evenly split between the component correlated with education and the component mediating the direct or residual association. The later component increases in relative importance for specifications using market or employment income (but the size of the partial correlation remains the same), while the former decreases in relative and absolute size.

Overall, these results suggest that employment status is an important pathway that accounts for a significant part of the association between parental and child income. More specifically, it accounts for part of the association between parental and child income that is explained by differences in educational attainment levels for children across the parental income distribution. An interpretation for this finding is that children of higher income parents earn a higher level of income because they have a higher probability of employment, which is associated with higher income.

The probability of employment is partly correlated with educational attainment. Table B.3 shows that the mediating role of education operates largely net of employment status (and of marital status). The reason may be that there is a large variation in the income of employed individuals based on their level of education. The decomposition of the pathways that account for intergenerational income transmission among employed respondents conducted in this section will contribute to a better understanding of the role of education in the sorting of respondents into jobs with characteristics associated with different levels of earnings.

[16]In a separate decomposition analysis including marital status in Table B.4, we find almost no role for marital status in specifications using child individual income, and a small role in the specification using child and spouse income, when present (explaining less than 10% of the association). Marital status has a marginal mediating role in the intergenerational transmission of employment income among employed individuals.

Table 2.10: Decomposition of relative mobility in employment income, employed respondents

		Child couple income		Child individual income		Child and parent indiv. income	
		Effect	%	Effect	%	Effect	%
Total effect	$\frac{Cov(y_{i,c},y_{i,p})}{Var(y_{i,p})}$	0.211	100.0	0.168	100.0	0.141	100.0
Explained	$\frac{Cov(y_{i,c},y_{i,p})}{Var(y_{i,p})} - \frac{Cov(\widehat{\hat{y}^{\vartheta}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.105	49.8	0.096	56.8	0.091	64.7
Education, total	$\frac{Cov(\widehat{\hat{y}^{Ed}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.084	39.9	0.077	45.6	0.072	51.4
Education, through job char.	$\frac{Cov(\widehat{\hat{y}^{S}_{i,c}},y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\widehat{\hat{y}^{Q}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.040	18.9	0.037	21.7	0.035	24.9
Skill use	$\frac{Cov(\widehat{\hat{y}^{S}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.038	18.1	0.034	20.2	0.032	23.0
Job quality	$\frac{Cov(\widehat{\hat{y}^{Q}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.002	0.8	0.003	1.5	0.003	1.9
Education, net	$\frac{Cov(\widehat{\hat{y}^{\vartheta}_{i,c}},y_{i,p})}{Var(y_{i,p})}$	0.044	21.0	0.040	23.9	0.037	26.5
Job char., uncorr. w/ ed., total	$\frac{Cov(\hat{\psi}^{S}_{i,c},y_{i,p})}{Var(y_{i,p})} + \frac{Cov(\hat{\psi}^{Q}_{i,c},y_{i,p})}{Var(y_{i,p})}$	0.021	10.0	0.019	11.3	0.019	13.3
Skill use	$\frac{Cov(\hat{\psi}^{S}_{i,c},y_{i,p})}{Var(y_{i,p})}$	0.015	7.0	0.012	7.0	0.011	7.9
Job quality	$\frac{Cov(\hat{\psi}^{Q}_{i,c},y_{i,p})}{Var(y_{i,p})}$	0.006	2.9	0.007	4.3	0.008	5.4
Unexplained	$\frac{Cov(\hat{\psi}^{\vartheta}_{i,c},y_{i,p})}{Var(y_{i,p})}$	0.106	50.2	0.073	43.2	0.050	35.3
N		5,236,200		5,236,200		5,236,200	

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

Education accounts for 39.9 to 51.4% of the association between parental and child income. Approximately half of this relationship is linked to job characteristics, almost exclusively skills used at work. Indeed, the part of respondent income that is correlated with job quality only is itself largely unrelated to education. This pattern might arise, for instance, if the skills and networks that facilitate access to high-quality jobs are not otherwise leveraged nor acquired in education.

Overall, job skills account for 25.1 to 30.9% of the association between parental and child employment income ranks. Approximately 75% of that effect comes through their association with educational attainment. In other words, a quarter of the effect operates through income transmission mechanisms that do not manifest themselves until the labour market. As for job characteristics, they account for 3.7 to 7.3% of the overall association between parental and child employment income ranks. Opposite to what we observe for job skills, a quarter of that effect is associated with education-related skills, and three quarters with non-education-related skills.

Of the portion of relative intergenerational mobility that is not associated with education, 16.5 to 27.5% is linked to job characteristics, mostly job skills. This means that a large portion of the association between parental and child income that is not correlated with education is accounted for by job matches with characteristics associated with higher rewards among children born from higher income parents. This may reflect unobserved skills that are rewarded in the labour market, but not in the education system. Alternatively, it may be the result of network or cultural capital effects, as discussed in Section 2.2.

2.7 Conclusion

Estimates of the intergenerational transmission of income in Canada suggest that it is a fairly mobile society. The literature on the topic has demonstrated that the link between parents' and children's income is complex, arising from the cumulative and complementary interactions of pre- and post-natal environments (Björklund et al., 2006; Corak, 2013; Cunha et al., 2006). Higher-income parents have more opportunities to invest in their children, and disparities that arise early on in life have long-lasting effects. On the other hand, there is evidence that skills retain some degree of malleability throughout life (Cunha et al., 2006), and that parental characteristics may affect children's income through channels revealed to researchers at later stages of life (eg. Corak and Piraino, 2011).

In the United States, Bowles and Gintis (2002) have found that education and cognitive and non-cognitive skills account for approximately half of intergenerational income correlation. In this paper, we investigate the extent to which this result also applies in Canada. Along many dimensions Canadian children face less unequal environments than children in the United States (Corak et al., 2011). In this context, it would be reasonable to expect not only that intergenerational income correlation would be smaller in Canada, but also that education would account for a smaller portion of the correlation that does exist. In fact, we find that a similar proportion of Canadian relative mobility is associated with education. Furthermore, our estimates show comparable experiences of intergenerational mobility across individuals with differing levels of education. Both findings suggest untapped opportunities in the education system to improve people's mobility.

The rich LISA survey data allows us to push the investigation further, and consider whether children of high-income parents further differentiate themselves by pursuing graduate studies or by choosing more lucrative fields of study. Interestingly, we find suggestive evidence that respondents with a Bachelor's degree may have successfully used their field of study to increase mobility; or that more mobile individuals choose better-paying fields. Field of study also plays a meaningful role among individuals who have completed graduate studies, and is associated with reduced mobility. As tertiary education has become increasingly widespread, uncovering the role of more specific decisions, like degree type and field of study, is also increasingly central to our understanding of the part played by education in people's experience of mobility.

Finally, we further decompose intergenerational mobility to account for the skills respondents report using at work (reading, writing, communication, mathematics, physical strength and dexterity) and specific indicators of job quality (unionization status, permanence of contract and authority over other employees). Almost half of the portion of intergenerational mobility that is associated with education is itself correlated with job skills. However, very little of the role played by education-related skills has to do with selection into higher-quality jobs, net of job skills. Altogether, job characteristics are also associated with 16.5 to 27.5% of the portion of intergenerational mobility that is not correlated with education. Our findings suggest that it is despite, not because of the education system that Canadians experience greater mobility than their southern neighbours. As such, education policy offers much untapped potential for improving mobility. However, several institutions pave the way from early childhood learning to post-secondary education. Further work should seek to determine at which point along the way the system limits rather than bolsters mobility.

Chapter 3

The evolution of wealth inequality in Canada between 1982 and 2011: Evidence from personal income tax records

3.1 Introduction

Research on the evolution of wealth inequality has produced often broadly varying results, both quantitatively and qualitatively, depending on the data source used. One of the most visible debates about the measurement of wealth inequality has come from the United States. On the one hand, researchers using household surveys or estate tax records have estimated that wealth held by the top 1% and the top 0.1% experienced very little variation starting in the 1980s, after a period of decline initiated around the Great Depression (Kopczuk, 2015; Kopczuk and Saez, 2004). Conversely, Saez and Zucman (2016) exploited income tax record to produce estimates for the period between 1913 and 2012, and found that the share of wealth held by the top 0.1% climbed from 7 to 22 % between 1978 and 2012.

Since then, a number of papers have discussed the factors underlying these differences, and proposed methods that play to the strengths of various data sources (eg. Bricker et al., 2016; Kopczuk, 2015). Increasingly, researchers combine data sources to try and circumvent data limitations (Bricker et al., 2016; Davies and Di Matteo, 2018; Garbinti et al., 2017; Martínez-Toledano, 2017). In general, five data sources have been used to measure wealth inequality: surveys of household assets and liabilities, estate tax records, personal income tax records, lists of wealthy individuals, and wealth tax data. Kopczuk (2015) provides a thorough review of the first four data sources, along with an in-depth discussion of the sources of differences in the estimates they have produced. Although there are variations across countries, surveys generally offer detailed data, with direct reporting of wealth, but may suffer from poor coverage at the top and irregular coverage over time, under-reporting of assets and liabilities, and in some cases, small sample size. Both the estate tax and income capitalization methods typically offer more frequent observations. The estate tax method also provides direct observations of wealth, but it relies on assumptions about differences (or lack thereof) in mortality rates across the wealth distribution, and about the evolution of these differences over time. Similarly, the income capitalization method relies on assumptions about differences in rates of return within asset classes across the wealth distribution. As for lists of wealthy individuals, while they are often considered a useful tool to complement other data sources, especially to address under-coverage at the top in survey data, they are not necessarily compiled systematically over time. Finally, like household surveys, wealth tax data provides more direct observations on wealth; and like estate tax and personal income tax records, it offers fairly broad coverage. Authors that have used wealth tax data have found conflicting evidence on the validity of some of the assumptions adopted by other methods. For instance, Roine and Waldenström (2009) in Sweden and Martínez-Toledano (2017) in Spain respectively found that the estate tax and income capitalization methods produce similar results as the wealth tax data. Conversely, Fagereng et al. (2016) and Lundberg and Waldenström (2018) both reported violations of the income capitalization method's constant returns assumption, using data for Norway and for Sweden. Across countries and across data sources, variations in results also vary with the unit of analysis. For instance, wealth surveys such as the Survey of Consumer Finances (SCF) in the United States and the Survey of Financial Security (SFS) in Canada are conducted at the household level. Conversely, personal income tax records report information at the tax unit level, which is defined differently across countries.

In Canada like elsewhere, existing research has encountered a substantial challenge in coming to a conclusion on the state and evolution of wealth inequality, especially in more recent periods.[1] Relying on estate tax records, Davies and Di Matteo (2018) find that the share of wealth held by the top 1% decreased between 1946 and 1970, from 29.6 to 19.6 %. For the period between 1970 and 2012, Davies and Di Matteo (2018) combine household survey data and lists of wealthy individuals, and find that wealth concentration has remained approximately constant over the period, with the top 1% holding approximately 23 to 24%. Earlier work using the same (unadjusted) household survey found somewhat different results. Oja (1983) concludes that the period between 1970 and 1977 was characterized by a moderate decline in wealth inequality, as measured by the percentage of wealth held by different quintiles. That change was driven in part by a decrease in the concentration of financial wealth. For that period, Morissette and Zhang (2006) found similar results. Between 1977 and 1984, their results are in line with Davies and Di Matteo (2018). According to them, wealth inequality later increased, between 1984 and 1999, due not only to an increase in net worth at the top of the distribution, but also to a decrease at the bottom of the distribution; that is, certain groups became worse off both in relative and in absolute terms (Morissette et al., 2006; Morissette and Zhang, 2006). Those who experienced the largest declines include couples with children and lone-parent families, recent immigrants, and individuals aged 25-34 years old. In comparison, the net worth of people aged 55 and over increased between 1984 and 1999. This last finding is consistent with Brzozowski et al. (2010), who observe that

[1] The greater consensus in earlier periods may stem from the smaller variety of data sets that have been used to address them.

top wealth shares have remained constant among the working-age population.

Of particular interest is the fact that wealth inequality has experienced a very modest increase, if it has changed at all. This contrasts with what has happened in other countries, and most notably with what has happened in the United States. For the period following the Great Recession, Wolff (2017) and Kuhn et al. (2018) have shown that the slumping of housing prices and increase of stock prices has affected differently individuals at different parts of the wealth distribution, thereby contributing to the exacerbation of wealth concentration. In both cases, Canada had a much more mitigated experience, which would explain the modest or non-existent change in wealth inequality over the last decade. However, as Davies et al. (2017) point out, this does not explain the lack of movement since the 1970s. Using the SFS, they demonstrate that growth in household wealth inequality was stifled by demographic changes. If it wasn't for an ageing population and for the rise of couples without children, they estimate wealth inequality would have grown substantially more between 1999 and 2012.

In this context, we contribute to the overall portrait of wealth inequality in Canada, by bringing to bear a new source of data. More specifically, we present new evidence on wealth inequality by applying the income capitalization method to the personal income tax records of Canadians. As mentioned, most of the existing research on wealth inequality has relied mostly on the SFS. Although an invaluable source of information on the wealth held by Canadian households, it is limited in its coverage of the wealthier segments of the population, in its sample size, and in the frequency of its data collection.[2] According to Statistics Canada, there is also evidence that survey respondents under-report their assets and liabilities, in particular their financial assets. To the extent that we expect people at different points of the distribution to differ in the composition of their portfolio, this implies that any analysis that relies on the Survey of Financial Security is likely to under-estimate wealth inequality at any given point in time. Whether it would affect trends is harder to determine. In addition to contributing to the overall understanding of what has happened with wealth inequality in Canada over the past three decades, we add to existing work by providing the first set of high-frequency (yearly) estimates for that period. This allows us to consider more specifically what happened around key moments, like changes in the redistributive role of the tax and transfer system in the 1990s, which is known to have contributed to the increase in after-tax income inequality Green et al. (2016), and the Great Recession. Indeed, the Survey of Financial Security's sporadic coverage makes it ill-suited to investigate these events. Furthermore, our paper presents the first estimates of trends in individual-level inequality in Canada. Among other things, this makes for more straightforward comparisons with the literature on income inequality.

Consistent with earlier results from Davies and Di Matteo (2018), we find that the top 10% share has been fairly stable over the period considered, with a very modest decrease from 1988 to 1999 and a modest increase thereafter. Our higher-frequency data allows us to determine that the latter arises from an increase in the top 10% share between 1999 and 2007, followed by an almost equal decrease between 2007 and 2011. Furthermore, whereas Davies and Di Matteo (2018) found that the top 1% share was fairly stable between 1984 and 2012, our results suggest that very wealthy people have

[2]The SFS also uses top-coding as a strategy to protect respondents' privacy.

in fact become relatively wealthier between 1995 and 2007. Their gains were subsequently halved between 2007 and 2011. Although we do not present the results here, we observe a similar trend for the top 0.1%. Importantly, using individuals as the unit of analysis yields higher estimates of inequality than using households, as the SFS does. We estimate that the top 10% share peaked at 76% in 2006 and 2007, compared to approximately 56% in the SFS Davies and Di Matteo (2018). Corresponding numbers for the top 1% share are 37% compared to 23%.

The large sample size of the LAD also allows us to look at wealth inequality within groups, namely by sex, cohort and region. We find that the increase in the top 1% share initiated in the mid-1990s is most pronounced among men; and that the share held by the bottom 50% of men decreased over the period covered, while the share held by the bottom 50% of women increased. We also find evidence that inequality decreases over the life-cycle, and that it does so fairly steadily between ages 30 and 65. Lastly, we observe that wealth concentration increased most in the Prairies, between 1988 and 2005, and that the Atlantic provinces exhibited not only the highest top 10% share over the period covered, but also the most stable.

The rest of the paper proceeds as follows. In Section 3.2, we clarify the definition of wealth used in this paper. In Section 3.3, we present the two data sources that form the basis of our analysis, along with our application of the income capitalization method. In Section 3.4, we discuss the trends in wealth inequality between 1982 and 2011. One of the main limitations of the income capitalization method is that it cannot account for wealth that does not generate taxable income. In Section 3.5, we present a possible way of addressing this issue, as well as preliminary evidence on the impact of that correction on our results. Finally, we conclude in Section 3.6.

3.2 Concept of wealth

In addition to the data sources used for the analysis, the concept of wealth used has an incidence on the results and on the interpretation of these results. As mentioned in Section 3.1, we use wealth measured at the individual level, to the extent that personal income tax records reflect individual income. Our main analysis is limited to wealth that generates taxable income. However, we also present preliminary work that better accounts for debts and assets that do not generate taxable income in Section 3.5. For the moment, we exclude social security wealth, although Milligan (2005) has shown that it plays a substantial role, especially in the later stages of life. In the future, longitudinal personal income tax records lend themselves very well to the imputation of social security wealth as done in Baker et al. (2003) and Milligan (2005). Finally, we exclude human capital from our definition of wealth. Notably, Abbott and Gallipoli (2019) have found that it mitigates inequality, although to a lesser degree in more recent years.

3.3 Data and methodology

The starting point of our analysis is the data on national wealth recorded in Canada's National Balance Sheet Accounts (NBSAs). Following Saez and Zucman (2016), we use the income capitalization method to allocate assets proportionally to the value of capital income Canadians report in their taxes. This data is documented in the Longitudinal Administrative Databank (LAD). The intuition behind the income capitalization method is that several assets generate capital income flows. Therefore, if both rates of return and income at the individual level were observed, it would be possible to calculate wealth. In lieu of rates of return, capitalization factors are used. Let W_a, the total wealth for a given asset class a, and Y_a, the total income stemming from that asset class. The capitalization factor is then $c_a = \frac{W_a}{Y_a}$. Observed individual capital income flows are then multiplied by the corresponding capitalization factors to obtain estimates of individual wealth. In turn, these can be used in calculating wealth inequality. For this paper, we focus on top and bottom shares; that is, we consider trends in top 10% and top 1%, and in the bottom 50% share and the middle 40% share (P50-P90).

3.3.1 National Balance Sheet Accounts

The National Balance Sheet Accounts (NBSAs) inform non-financial assets, financial assets, liabilities and net worth for the six sectors of the Canadian economy: financial and non-financial corporations, the government, households (or persons and unincorporated businesses) and non-profit institutions serving households, and the non-resident sector. They were first produced in 1983-1984, and are part of the Canadian System of Macroeconomic Accounts, which itself goes back to the 1940s.

For the period from 1982 to 2011, the NBSA data for the household sector is drawn from CANSIM table 36-10-0510-01 (Statistics Canada, n.d.-e). We classify assets that produce taxable income into five categories: non-corporate business equity, non-owner-occupied residential and non-residential real estate (henceforth, real estate), public equity, fixed-income claims, and life insurance and pensions. Figure 3.1 shows trends in the five categories of wealth held by the household sector, from 1970 to 2011, and corresponding changes in aggregate wealth composition. Wealth held by the household sector has increased considerably over the period considered, driven in great part by growth in non-corporate and public equity, and life insurance and pensions. Non-corporate business equity is the largest wealth component throughout the period covered by our analysis.

Residential structures: owner-occupied and otherwise

The NBSAs document the stock of residential structures held by Canadian households, but do not distinguish between rental and owner-occupied structures. Since the latter do not generate taxable capital income, the distinction is important to the income capitalization method. To address this issue, we follow previous work that has relied on the ratio of imputed to paid rents to approximate the value of owner-occupied and rental residential structures (eg. Baldwin et al., 2007; Diewert and Yu, 2012). More specifically, we use household expenditures on rental fees, imputed and paid, recorded in CANSIM Table 36-10-0107-01 (Statistics Canada, n.d.-d). Between 1982 and 2011, 70 to 75% of residential structures are thus considered to have been owner occupied.

Figure 3.1: Wealth composition from 1970 to 2011

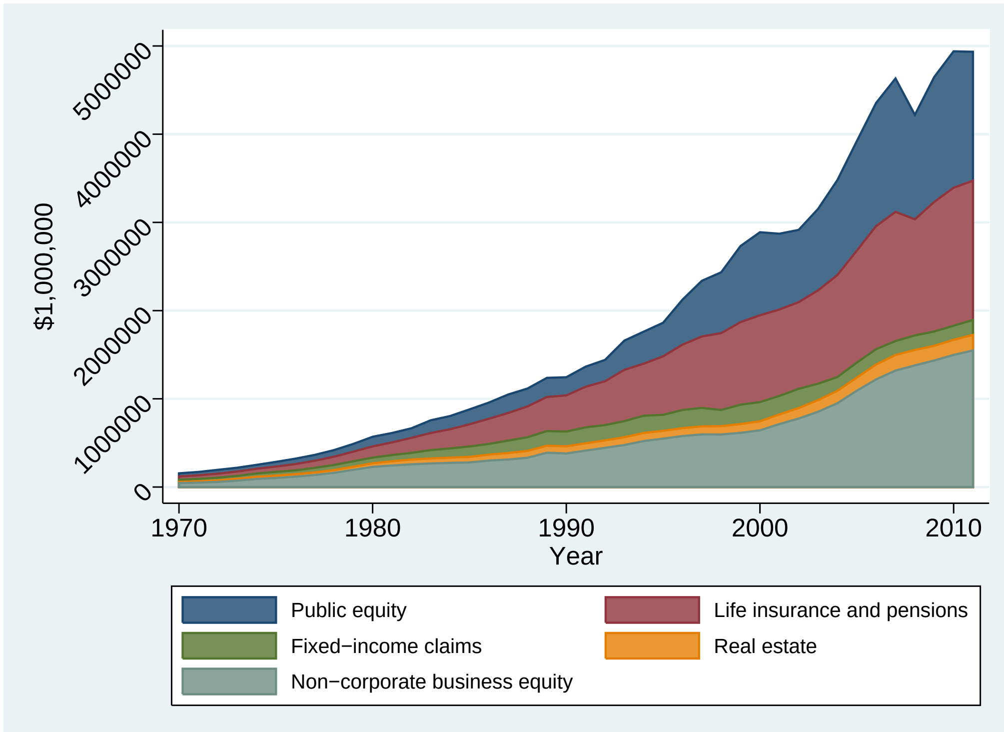

3.3.2 Longitudinal Administrative Databank

The Longitudinal Administrative Databank (LAD) regroups tax information for approximately 20% of Canadian tax filers, starting in 1982. Every year, Statistics Canada uses personal income tax (T1) records collected by the Canada Revenue Agency (CRA) to construct the T1 Family Files (T1FF): the T1s of census family members are linked, and a number of family-specific variables are constructed to document these families. The LAD takes cross-sections of the T1FF and links them across years through peoples Social Insurance Number (SIN). Sample members are followed so long as they can be found in the T1FF, and new members are added over the years to maintain the sample size and representativeness of the data. Although the LAD samples individual tax filers, not families, it includes a large number of variables that inform family-level characteristics, including family size and composition, partner characteristics and family income. These are crucial to create the link between the LAD and the SFS that makes possible the imputation of assets that do not generate taxable income.

Application of the income capitalization method to Canadian tax data

The Canadian tax system records several sources of capital income, which form the basis of the capital income aggregates used for the income capitalization method. Table 3.1 shows the equivalence between the five asset classes discussed in Section 3.3.1 and the corresponding capital income sources from the LAD. It is worth noting that life insurance is included with pensions in the NBSAs, and we only observe contributions and income pertaining to the latter in the LAD. By capitalizing pensions

Table 3.1: Equivalence between asset classes and capital income flows

Asset class (NBSA)	Capital income flow (LAD)
Non-corporate business equity	Net business income
	Net commission income
	Net farming income
	Net fishing income
	Net limited partnership income
Real estate	Net rental income
Public equity	Dividends
	Net capital gains
Fixed-income claims	Interest and investment income
Life insurance and pensions	RRSP contributions
	RRSP spousal contributions
	Registered pension plan contribution
	RRSP income
	Annuity income from RRSP
	Pension and superannuation income

based on aggregate wealth held in life insurance and pensions, we assume that the two are distributed similarly in the population.

Some of the variables available on the LAD have changed over time, reflecting both changes to the way the information is recorded and changes to the tax system. There are two sets of variables that are affected by the first kind of change. First, net limited partnership income has only been recorded independently since 1988. Prior to that date, it was included with net business income, net rental income, or other income. Although this is not ideal, net limited partnership income is very small compared to net business income and other components of self-employment income, as well as compared to net rental income. Second, variables that describe pension contributions and pension income have also changed over the years. On the one hand, Registered Pension Plan contributions are only recorded on the LAD going back to 1986. Similarly, Registered Retirement Savings Plan income was included with other income until 1987. Given the important role that pensions play in the portfolio of middle- and upper-middle-income families, this implies that changes in wealth inequality in the second half of the 1980s that are driven by pension wealth should be taken with a grain of salt. We discuss these issues in Section 3.4 when we present the results.

Next, sources of capital income that are associated with public equity have experienced changes in taxation between 1982 and 2011. Taxation of dividends has decreased over the period, with particularly low levels from 1988 to 2005. Taxation of capital gains has also changed, with an increase from 1988 to 1999, and a return to pre-1988 levels starting in 2001. Furthermore, there was a change in legislation in 1994, which restricted the size of businesses for which deductions could be claimed.

3.4 Results

Figure 3.2 shows the share of total wealth held by four segments of the wealth distribution. The left panel shows the share of the top 10% (P90-P100) and the top 1% (P99-P100); and the right panel shows the share of the bottom 50% (P0-P50) and the middle 40% (P50-P90). Note that each panel has two axes, the left one corresponding to the first series and the right one corresponding to the second series. We do not present results for the top 0.1% (P99.9-P100) here, but they are qualitatively similar to those for the top 1%. Before looking at the results, it is worth noting one thing that reflects an important limitation of our data. We include in all our graphs a line at 1988, to highlight the fact that there were important changes in the way some variables were recorded between 1982 and 1987. Of particular relevance for our results is the fact that Registered Pension Plan contributions and Registered Retirement Savings Plan income became their own variables only in 1986 and 1988 respectively. As such, the dramatic decrease in the share of wealth held by the top 10%, and the corresponding increase in the bottom 90% share are artifacts of the data. Although we cannot speak directly to the pre-1988 trends, we include those years nonetheless since they highlight the importance of pensions in the portfolio of the bottom 90%.

Bearing in mind this limitation, three interesting results come out of Figure 3.2. First, the top wealth shares we estimate using the income capitalization method are considerably higher than previous estimates obtained with the SFS. To illustrate the comparison, Figure 3.3 shows trends for the top 10% and the top 1%, using the LAD and the adjusted SFS estimates from Davies and Di Matteo (2018). At its highest in 2006 and 2007, we calculate that the wealth share held by the top 10% (top 1%) is 76% (37%), while it is approximately 56% (23%) in the SFS. This is in line with what we expect, given that the LAD allows us to measure inequality at the individual level, while the SFS provides estimates at the household level. Second, our results seem to confirm the findings of Davies and Di Matteo (2018) for the top 10%, but we find more evidence suggesting the growth of the top 1% share. Among the top 10%, there is a very modest decline in the share of wealth between 1988 and 1999, from 74 to 73%, and a modest increase from 1999 on. There is a three percentage point increase in the share held by the top 10% between 1999 and 2007, but most of that gain is lost thereafter. For the top 1%, we find that there is a fairly steady increase in wealth share starting in the mid-1990s and culminating with the 2006-2007 peak. Despite the drop in their share following 2007, the top 1% still experienced an overall two percentage point increase between 1995 and 2011, a 7% increase from their 1995 share. Finally, patterns among the middle 40% mirror those observed for the top 10%: the share of wealth undergoes a light increase from 1988 to 1999, then decreases until 2007 and recovers afterwards.

3.4.1 Changes in wealth inequality in different groups

Next, we consider how inequality changed within different groups. More specifically, we look at inequality by sex, birth cohort, and region. Figure 3.4 shows trends by sex; that is, we calculate the share of wealth held by women at different point of the women's wealth distribution, and similarly for men. First, wealth concentration among women has changed respective to men over the period

Figure 3.2: Overall trends in wealth concentration

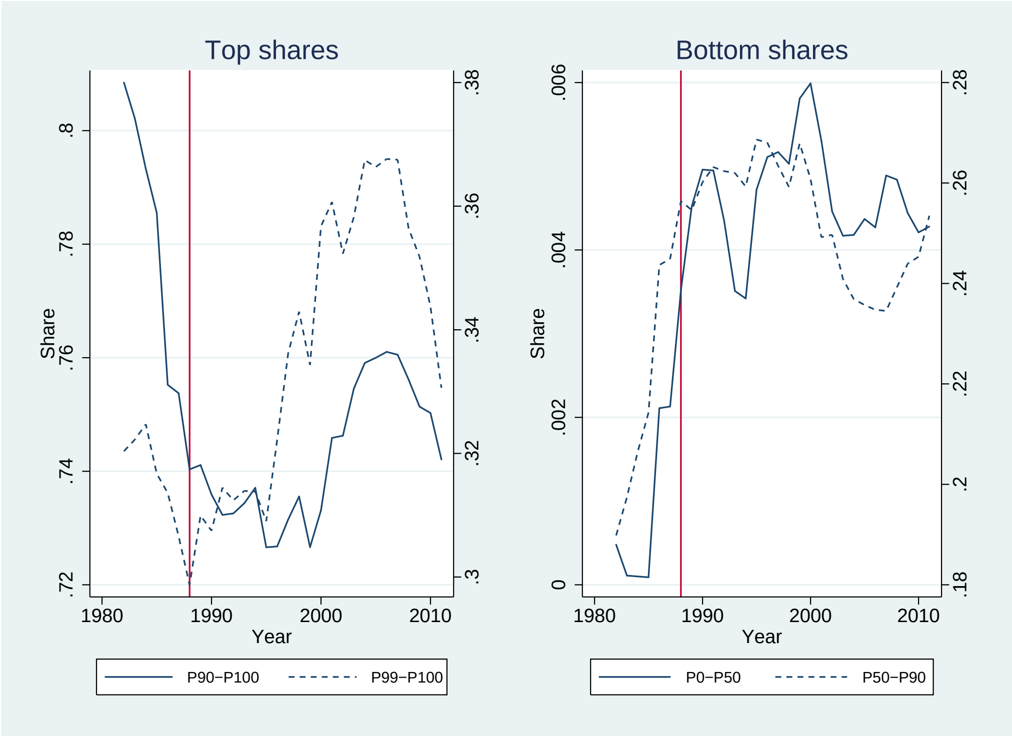

Figure 3.3: Overall trends in wealth concentration, comparison of LAD and SFS estimates

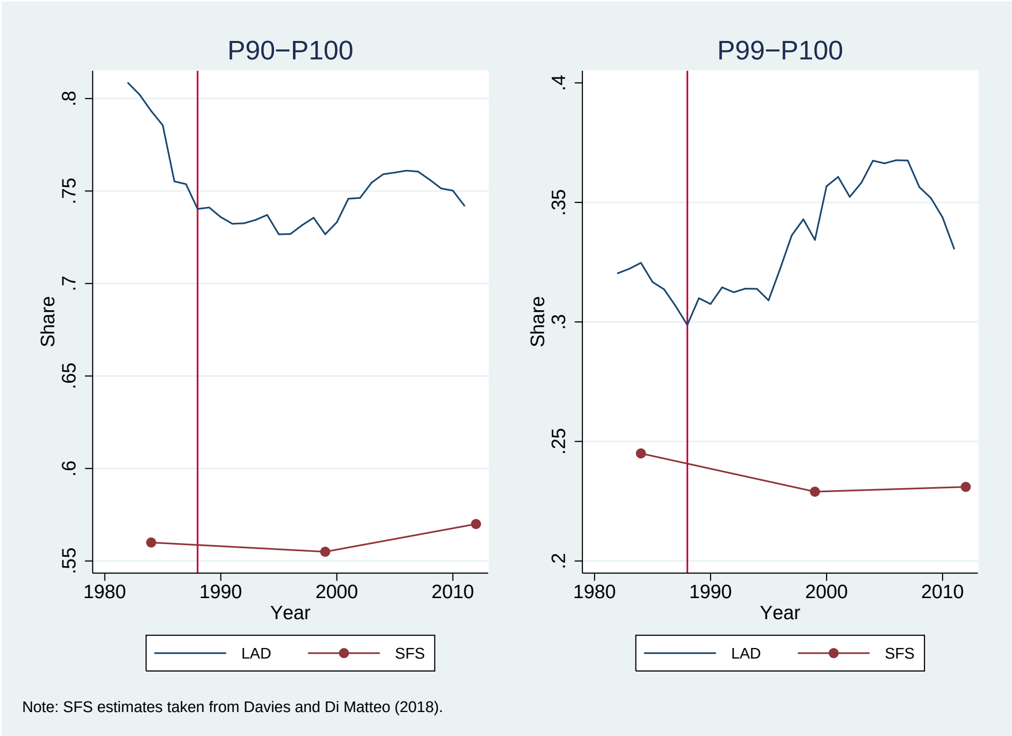

Figure 3.4: Trends in wealth concentration, within sex

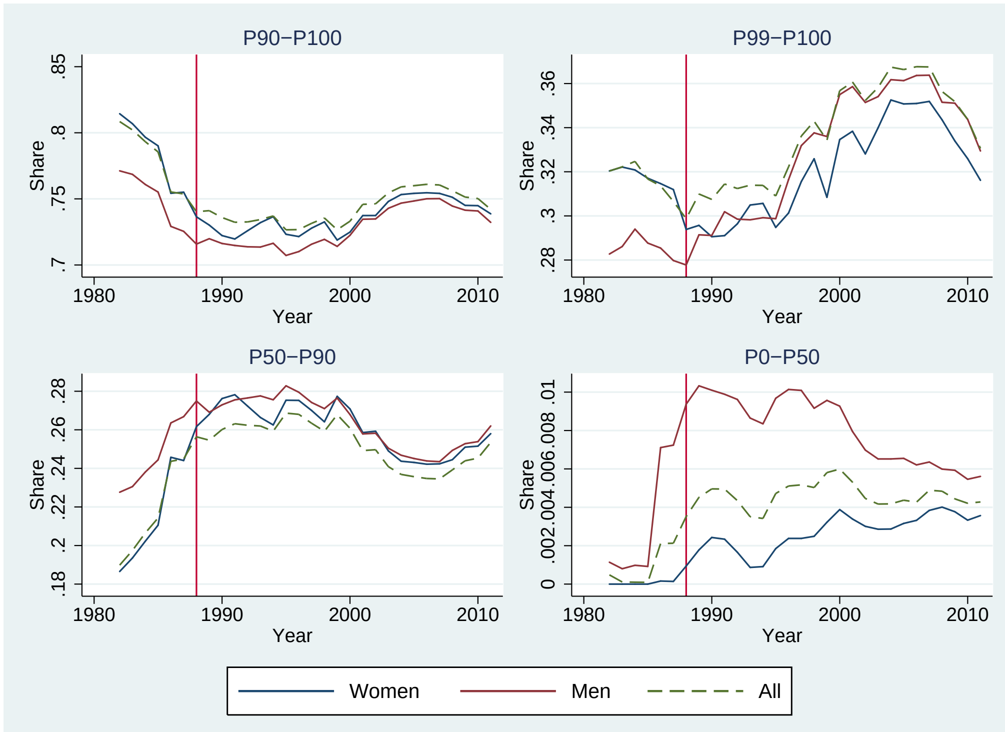

covered. For the top 1%, the increase in wealth shares that started in the mid-1990s was much more pronounced for men then for women. The stagnation in wealth shares in the first half of the 1990s also appears mostly driven by women; among men, we see a modest increase. Second, until 2000, the imbalance between the top 10% and the bottom 90% was more pronounced among women than among men. Finally, the poorest women have experienced gains between the late 1980s and 2011, whereas the situation for the poorest of men has deteriorated over the period. They were particularly affected around the time of the dot-com bubble.

The following three figures show trends in inequality within cohorts. Our data allows us to characterize patterns for people born between 1900 and 1989, split in ten-year groups. To simplify exposition, Figure 3.6 shows trends for the 1900-1909 and 1910-1919 cohorts ("older cohorts"), Figure 3.7 shows trends for the 1970-1979 and 1980-1989 cohorts ("younger cohorts"), and Figure 3.5 shows trends for the 1920-1929 to 1960-1969 cohorts ("middle cohorts"). We first consider trends in inequality in the middle cohorts. Generally, wealth shares held by the top have increased across cohorts. In earlier figures, we saw that the imbalance between the top 10% and the bottom 90% increased between 1999 and 2007. In Figure 3.5, it is apparent that that was largely driven by the 1950-1059 and 1960-1969 cohorts; more specifically, by people who were between 41 to 60 years old in 2000. In comparison the wealth shares held by people aged 61 to 80 in 2000 (cohorts 1920-1929 and 1930-1939) were roughly constant through the 2000s.

As for the older and younger cohorts, Figure 3.6 and Figure 3.7 show that they exhibit patterns

Figure 3.5: Trends in wealth concentration, within cohort: middle cohorts

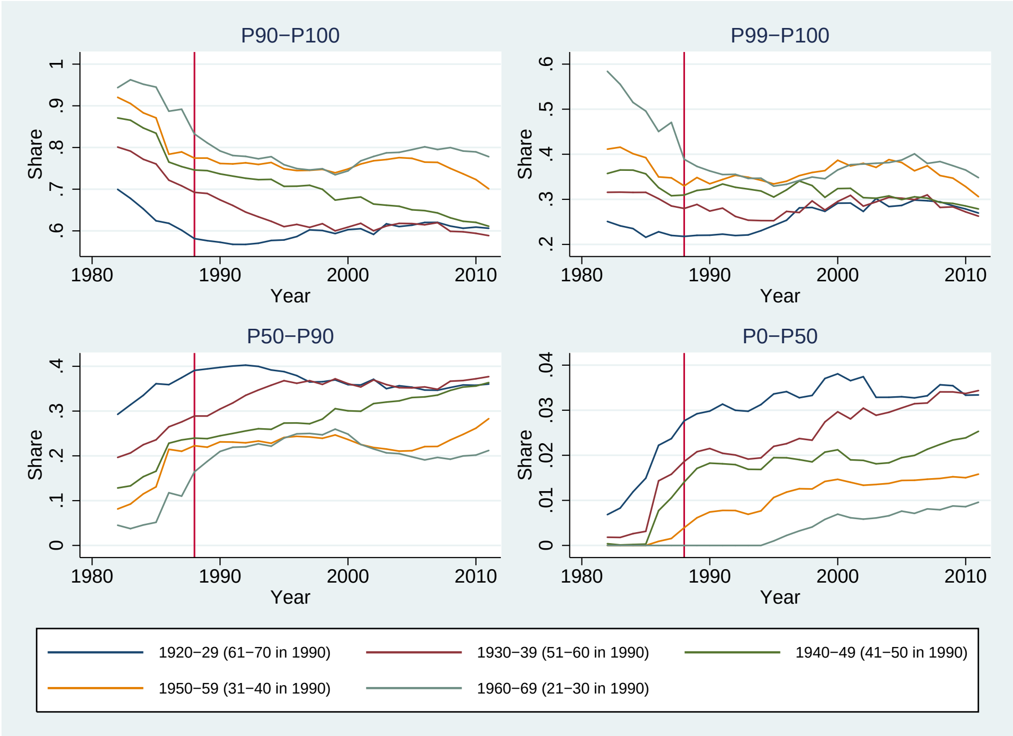

Figure 3.6: Trends in wealth concentration, within cohort: older cohorts

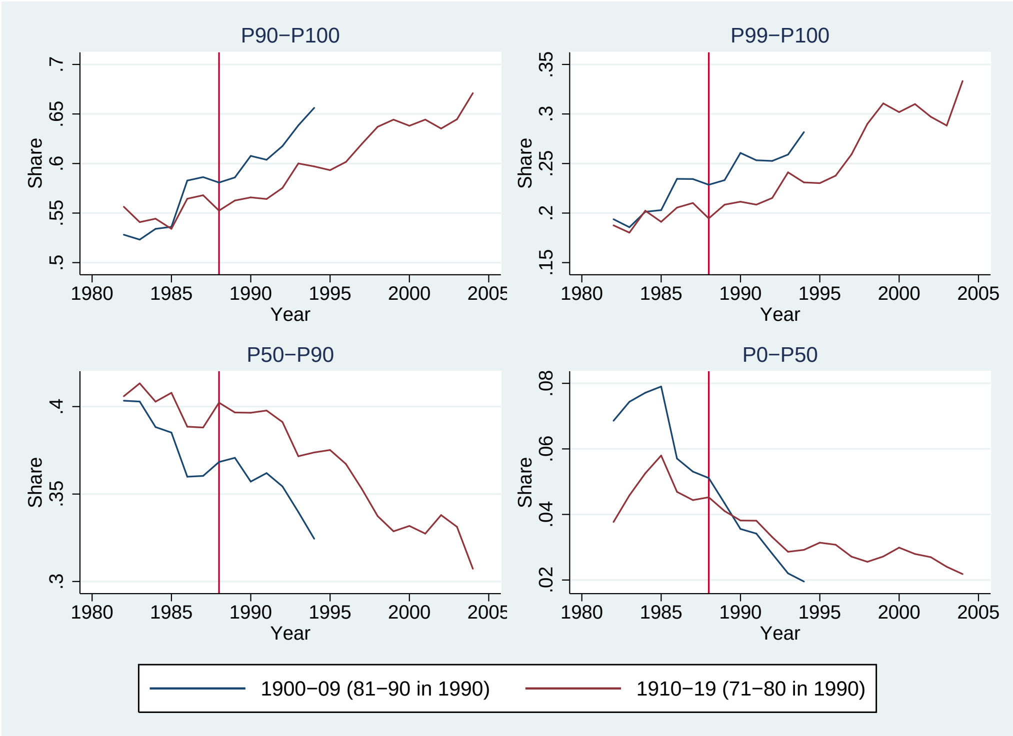

Figure 3.7: Trends in wealth concentration, within cohort: younger cohorts

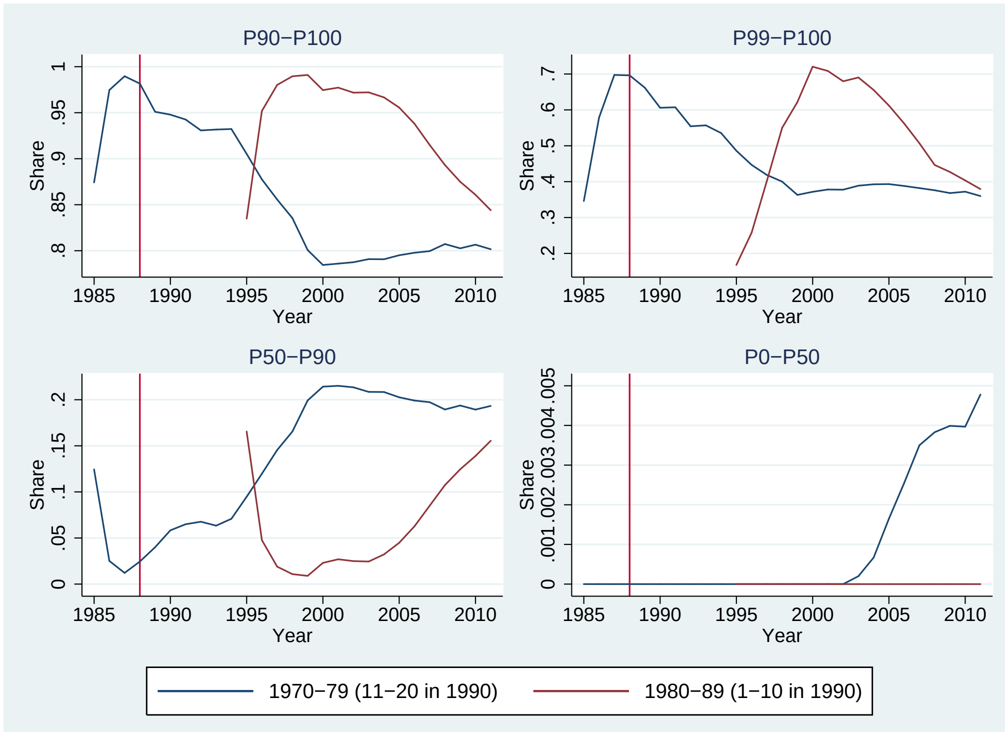

particularly characteristic of the period of life in which we observe them. On the one hand, Figure 3.6 shows higher wealth concentration in the older of the two cohorts, people born between 1900 and 1909. This contrasts with results from Figure 3.5, and likely reflects the two factors: first, at the point in time when they are observed, these cohorts are heavily selected; and second, we are probably also seeing a decumulation effect that exacerbates wealth concentration. To illustrate this point, suppose that extremely wealthy individuals consume the same amount in (very) old age as less wealthy individuals. In that case, the latter group will deplete their resources relatively faster than the former, thus exacerbating wealth concentration. Figure 3.7 shows similarly characteristic but very different patterns of inequality among the younger cohorts. Unsurprisingly, wealth inequality is very acute at a young age: fewer youths have any form of wealth, and so top shares are extremely high. The trend for the 1970-1979 cohort suggests that people start accumulating wealth in their mid-twenties, leading to a decrease in top shares. By the time people are in their thirties, wealth shares have become relatively stable. Interestingly, the increase in wealth imbalance between the top 10% and the bottom 90% that we observed for the 1950-1959 and 1960-1969 cohorts is almost non-existent in the younger groups, as it was in the older groups.

Of course, trends by cohorts are indicative of changes in inequality across cohorts, as well as changes over the life-cycle. As an alternative illustration of the latter, we stitch together the life-cycle profiles from the various cohorts. Figure 3.8 presents these pseudo life-cycle patterns. As suggested by Figure 3.7, wealth inequality drops considerably between 20 and 30 years old, then decreases steadily

Figure 3.8: Trends in wealth concentration, pseudo life-cycle patterns

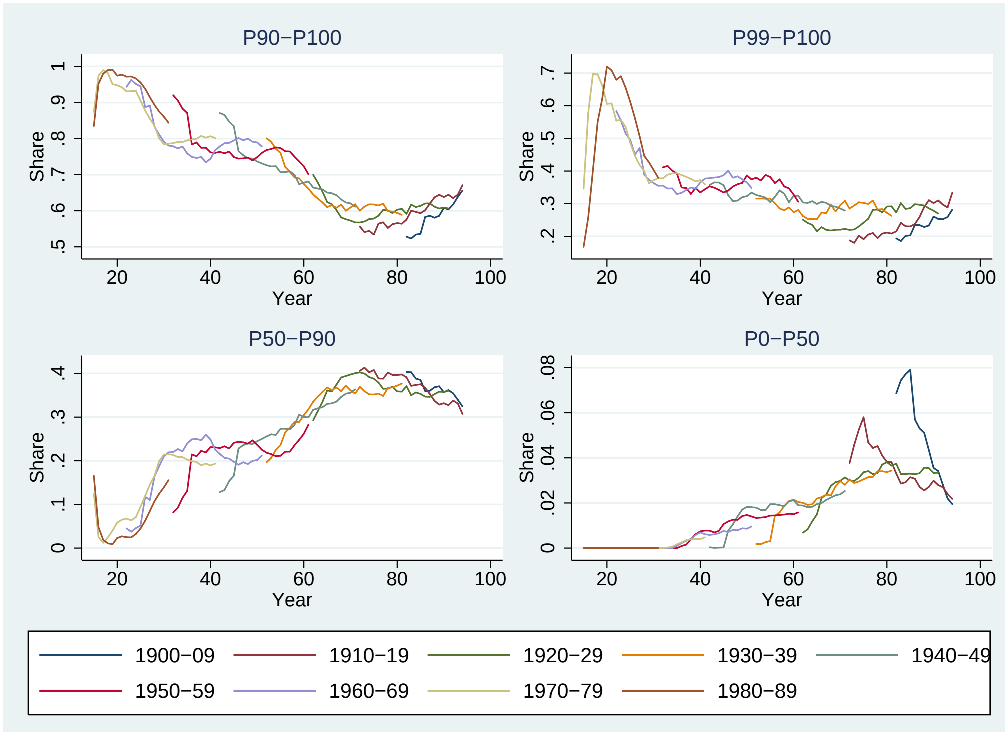

until 60 or 65 years old, at which point it plateaus and even picks up a little bit. Hence, the share of wealth held by the top 10% passes from almost 80% at age 30 to a little under 70% at age 60. Notably, the last panel of Figure 3.8 suggests that the bottom half of the population owns nothing until they're almost 40.

Finally, in Figure 3.9, we present trends in wealth inequality by region. First, Figure 3.9 shows that wealth concentration has increased substantially in the Prairies between 1988 and 2011, relative to the other provinces. It went from exhibiting the lowest to the highest top 10% share. The top 1% share in the Prairies also increased considerably between 1988 and 2005, but it experienced a relatively large drop after that. Second, wealth concentration was fairly constant in the Atlantic provinces throughout the period considered, around high levels. Even more striking, while accounting for pensions makes a large difference for the wealth share held by the poorest half in most provinces, the same does not hold in the Atlantic provinces. Third, in 1990 the bottom 50% fared relatively better in the Prairies, BC, Ontario and Quebec than in Canada at large, but that advantage largely disappeared in the first half of the 2000s.

3.5 Accounting for non-taxable capital income

To address the fact that some forms of wealth do not generate taxable income, we use the Survey of Financial Security (SFS) to impute owner-occupied housing, currency and deposits, and debt and non-

Figure 3.9: Trends in wealth concentration, within regions

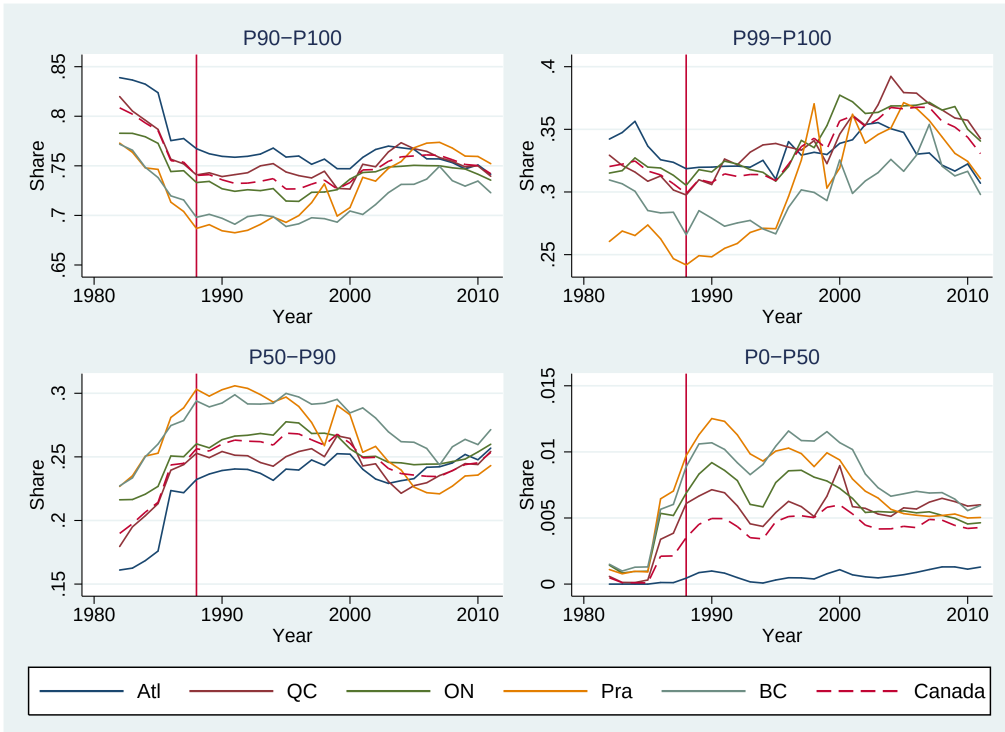

mortgage loans. In preliminary results not shown here, we find that accounting for non-taxable capital income impacts shares both at the top and at the bottom of the distribution. As expected, it results in lower estimates and higher estimates of top and bottom shares respectively. Furthermore, numbers that account for non-taxable capital income suggest a more clearly marked deterioration at the bottom over the period considered than observed in Section 3.4.

The Survey of Financial Security (SFS) is the main source of information on the assets and liabilities of Canadians. It offers a very detailed picture of the composition of wealth held by households in 1999, 2005, 2012 and 2016. Prior to the SFS, wealth data was collected through the asset and debt supplement to the Survey of Consumer Finances (SCF), incorporated every seven years between 1957 and 1984. We use the SCF/SFS (henceforth, SFS) to impute owner-occupied housing, currency and deposits, and consumer credit and non-mortgage loans for LAD sample members.

Our approach to the imputation of assets that do not generate taxable income follows very closely that used by Garbinti et al. (2017). Since the SFS collects information on assets and liabilities at the family level, we base the imputation on two family-level variables: the main income recipient's age and family market income.[3] First, we create ten age categories: less than 25 years old, 25-34, 35-39,

[3]The concept of family used in the SFS is the economic family, "a group of two or more persons who live in the same dwelling and are related to each other by blood, marriage, common-law union, adoption or a foster relationship" (Statistics Canada). This may include adults with their children, but also extended relatives such as grandparents or uncles and aunts. As such, it is broader than the LAD family. This is most likely to be an issue for the housing imputations, as people may choose to double up to cope with a relatively expensive housing market. For instance, in parts of Canada where housing

40-44, 45-49, 50-54, 55-59, 60-64, and 65 and older. Second, within each age category, we create six relative income categories: P0-P25, P25-P50, P50-P75, P75-P90, P90-P95, and P95-P100. For each of the resulting 60 groups, and for each of the three asset classes that do not generate taxable income, we compute two numbers: the proportion of a given asset that is held by people in that group, and the proportion of people in that group who hold that asset. For owner-occupied housing, we compute a third number, the proportion of housing value that is owed as mortgage.

Next, we use this information to allocate owner-occupied residential structures, currency and deposits, and consumer credit and non-mortgage loans from the NBSAs to LAD sample members. For each LAD family, we identify the main income recipient as the person with the highest pre-tax income, following the definition used in the SFS.[4] A useful feature of the LAD is that it allows to identify main income recipients even when they were not sampled; that is, if only one adult from a couple was sampled in the LAD, we have enough information to determine whether they are the main income recipient and, if they are not, we know the age of their partner. Having reproduced the 60 groups from the SFS in the LAD, we allocate aggregate wealth as follows. First, we randomly designate asset holders to match the proportion of SFS respondents who reported owning that asset in a given group. Second, we attribute to each asset holder the group-specific per capita value of that asset. For owner-occupied residential structures, we adjust that amount to reflect the proportion of housing that is owed as mortgage in a given group. Naturally, this approach can only account for inequalities in ownership rates and value across the groups that are explicitly used for the imputation; within these groups, inequality is assumed away by construction, implying that our estimates necessarily define a lower bound on inequality. Although not perfect, this is only really a problem if inequality within these groups has changed over time.

Because the imputations are done at the family level while we observe taxable income at the individual level, we take one more step to distribute wealth within families. For owner-occupied housing, we attribute wealth to adults only, following the proportion of adult-owned capitalized income held by each one. For currency and deposits and for consumer credit and non-mortgage loans, we allocate wealth following the proportion of family capitalized income held by each family member.

Incorporating these imputations decreases top shares, and the effect is quite large. Unsurprisingly, it has the greatest impact on the trends experienced by the bottom 50%. For that group, the results including the imputed assets and debts suggest a larger share overall, but also a more clearly marked decrease over the period considered.

is more expensive, extended family members may be more likely to stay together. When matching their housing to sample members in the LAD, this would understate the value of housing held by LAD sample members. As a result, this may lead to an understatement of wealth inequality in areas where there are greater incentives for extended family members to live together. Generally, if poorer households are more likely to double-up with extended family, the imputation procedure used would understate inequality.

[4] When identifying main income recipients in the LAD, we limit our search to adults, since children are included in LAD families only if they share some fiscal dependency relationship with the adults present. Furthermore, to follow what is done in the SFS, we use the older of two adults where they report the same income.

3.6 Conclusion

The literature on wealth inequality has produced quantitatively and qualitatively different results world-wide, often within given countries (eg. Bricker et al., 2016; Kopczuk and Saez, 2004; Saez and Zucman, 2016). To a large extent, the lack of consensus has stemmed from the challenge researchers face in measuring wealth, and from the consequent diversity of data sets that have been used (Kopczuk, 2015). In Canada, only one source of information has been used to inform the evolution of wealth inequality over the past three decades, the Survey of Financial Security (Brzozowski et al., 2010; Morissette et al., 2006; Morissette and Zhang, 2006; Oja, 1983). To circumvent under-coverage at the top, and mitigate the resulting under-estimation of wealth inequality, Davies and Di Matteo (2018) have complemented this source of information with lists of wealthy individuals. They find that between 1970 and 2012, the share of wealth held by the top 10% of Canadian households hovered around 56%, while the share held by the top 1% remained close to 23%. In other words, their results suggest no change in wealth inequality, as measured by wealth shares, over the last four decades. In comparison, earlier work found some evidence that wealth inequality may have decreased lightly between 1970 and 1977 (Morissette and Zhang, 2006; Oja, 1983, and increased moderately between 1984 and 1999 (Morissette and Zhang, 2006; Morissette et al., 2006). Hence, existing work in Canada suggests that there have been modest or no changes in wealth inequality in the recent period. This is particularly puzzling in light of the research that has suggested increases in several other countries, including the United States, United Kingdom, France and Sweden (Bricker et al., 2016; Lundberg and Waldenström, 2018; Roine and Waldenström, 2015; Saez and Zucman, 2016).

In this paper, we bring a new perspective on the evolution of wealth inequality in Canada. We exploit personal income tax records from the Longitudinal Administrative Databank (LAD) to capitalize income and obtain estimates of wealth inequality. Our results for the top 10% are consistent with earlier findings. We find that the top 10% share has been fairly stable over the period considered, with a very modest decrease from 1988 to 1999 and a modest increase thereafter. Our higher-frequency data allows us to determine that the latter arises from an increase in the top 10% share between 1999 and 2007, followed by an almost equal decrease between 2007 and 2011. Furthermore, whereas Davies and Di Matteo (2018) found that the top 1% share was fairly stable between 1984 and 2012, our results suggest that very wealthy people have in fact become relatively wealthier between 1995 and 2007. Their gains were subsequently halved between 2007 and 2011. Although we do not present the results here, we observe a similar trend for the top 0.1%. Importantly, using individuals as the unit of analysis yields higher estimates of inequality than using households, as the SFS does. We estimate that the top 10% share peaked at 76% in 2006 and 2007, compared to approximately 56% in the SFS Davies and Di Matteo (2018). Corresponding numbers for the top 1% share are 37% compared to 23%.

Our data also allows us to look at wealth inequality within groups, namely by sex, cohort and region. We find that the increase in the top 1% share initiated in the mid-1990s is most pronounced among men; and that the share held by the bottom 50% of men decreased over the period covered, while the share held by the bottom 50% of women increased. We also find evidence that inequality decreases over the life-cycle, and that it does so fairly steadily between ages 30 and 65. Lastly, we observe that wealth concentration increased most in the Prairies, between 1988 and 2005, and that the Atlantic provinces exhibited not only the highest top 10% share over the period covered, but also the most stable.

In conclusion, although the income capitalization method presents non-negligible limitations, it also offers attractive coverage at the top of the wealth distribution, and regular observations over long periods of time. In preliminary work, we find that the LAD is well-suited to imputations that allow us to account for wealth that does not generate taxable income. In addition to this, an obvious extension of the work presented here is to take advantage of the longitudinal nature of the data to improve the treatment of capital gains. Furthermore, the LAD is also an ideal data set to incorporate in our measures of inequality considerations surrounding social security wealth, relying on the work of Baker et al. (2003) and Milligan (2005). The fact that our results are in line with the SFS for the top 10% and that they differ in the expected direction for the top 1% is encouraging. It suggests that applying the income capitalization method to the LAD is a worthwhile endeavour to further our understanding of wealth inequality in Canada.

Conclusion

In this thesis, I present new evidence on economic vulnerability and inequality in the Canadian context. Both of the first chapters look at individual experiences, and more specifically at the role that family plays in shaping economic outcomes. In the first paper, I show that divorce and divorce risk impact the labour supply of women, from the time they get married to the period following marital dissolution. A 10% increase in the lifetime divorce risk faced by women when they get married leads to a 4.34 to 5.88% increase in their labour supply in every period throughout their marriage. This is similar in magnitude to the effect of a 10% increase in their lifetime wages. Furthermore, the effect of divorce on women who do experience it is conditional on the marriage they were in. Differences in spousal wage, quality of time spent with one's spouse and divorce risk while married have larger impacts on women's labour supply after than before divorce.

The reduced-form literature on divorce has extensively documented the fact that marital dissolution should be viewed as a process, starting well before the actual event. However, it has produced largely contradictory results on the direction of the impact on women's labour supply. Understanding how marital dissolution affects women at different points in their lifetime provides a stronger basis for policy making. Being divorced is correlated with a wide range of negative outcomes, including poverty and morbidity (eg. Curtis and Rybczynski, 2014; Joung et al., 1997). However, I find that divorce exacerbates pre-existing differences across women's marriages. Therefore, effective policies should target factors that contribute directly to economic security and health, since vulnerabilities are amplified and not driven by divorce.

In the second chapter, we investigate the role played by education in the intergenerational transmission of income in Canada. In the United States, Bowles and Gintis (2002) have found that education and cognitive and non-cognitive skills account for approximately half of intergenerational income correlation. Along many dimensions Canadian children face less unequal environments than children in the United States (Corak et al., 2011). In this context, it would be reasonable to expect not only that intergenerational income correlation would be smaller in Canada, but also that education would account for a smaller portion of the correlation that does exist. In fact, we find that a similar proportion of Canadian relative mobility is associated with education. Furthermore, our estimates show comparable experiences of intergenerational mobility across individuals with differing levels of education. Both findings suggest untapped opportunities in the education system to improve people's mobility.

We also find suggestive evidence that respondents with a Bachelor's degree may have successfully used their field of study to increase mobility; or that more mobile individuals choose better-paying fields. Field of study also plays a meaningful role among individuals who have completed graduate studies, and is associated with reduced mobility. As tertiary education has become increasingly widespread, uncovering the role of more specific decisions, like degree type and field of study, is also increasingly central to our understanding of the part played by education in people's experience of mobility.

In the third chapter, we adopt a broader perspective on inequality. We study the evolution of wealth inequality in Canada, from 1982 to 2011. Whereas it is well documented in the wealth inequality literature that different data sets tend to yield different results (eg. Kopczuk, 2015), until now only one data source has been used to inform Canadian trends. Ours is the first paper to bring a different perspective on the topic. We apply the income capitalization method to the Longitudinal Administrative Databank (LAD), which produces yearly estimates, with better coverage at the top of the wealth distribution. Consistent with earlier results from Davies and Di Matteo (2018), we find that the top 10% share has been fairly stable over the period considered, with a very modest decrease from 1988 to 1999 and a modest increase thereafter. Our higher-frequency data allows us to determine that the latter arises from an increase in the top 10% share between 1999 and 2007, followed by an almost equal decrease between 2007 and 2011. Furthermore, whereas Davies and Di Matteo (2018) found that the top 1% share was fairly stable between 1984 and 2012, our results suggest that very wealthy people have in fact become relatively wealthier between 1995 and 2007. Their gains were subsequently halved between 2007 and 2011. In general, our estimates of inequality are higher than in existing work: we calculate that the wealth share held by the top 10% and the top 1% respectively peak at 76 and 37%, compared to 56 and 23% in Davies and Di Matteo (2018).

Bibliography

Abbott, B. and G. Gallipoli (2019). Permanent-income inequality. Working paper. → page 103

Alessie, R., T. F. Crossley, and V. Hildebrand (2006). Estimating a collective household model with survey data on financial satisfaction. Technical report, IFS Working Papers, Institute for Fiscal Studies (IFS). → page 15

Atkinson, A. B., A. K. Maynard, C. G. Trinder, J. Corlyon, S. Jenkins, and H. Sutherland (1983). *Parents and children: incomes in two generations.* London Eng. Heinemann Educational Books 1983. → pages 3, 72, 78

Autor, D. H. and M. J. Handel (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of Labor Economics 31*(S1), S59–S96. → page 82

Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333. → pages 77, 82

Aydemir, A., W.-H. Chen, and M. Corak (2013). Intergenerational education mobility among the children of canadian immigrants. *Canadian Public Policy 39*(Supplement 1), S107–S122. → page 75

Bailey, M. J. and S. M. Dynarski (2011). Inequality in postsecondary education. In G. J. Duncan and R. J. Murnane (Eds.), *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances*. New York: Russell Sage Foundation. → page 75

Baker, M., J. Gruber, and K. Milligan (2003). The retirement incentive effects of canada's income security programs. *Canadian Journal of Economics/Revue canadienne d'économique 36*(2), 261–290. → pages 103, 116

Baldwin, J. R., W. Gu, and B. Yan (2007). User guide for statistics canadas annual multifactor productivity program. *Canadian Productivity Review Research Paper*. Statistics Canada Catalogue no. 15-206XIENo. 14. → page 104

Bargain, O. and O. Donni (2012). Expenditure on children: A rothbarth-type method consistent with scale economies and parents' bargaining. *European Economic Review 56*(4), 792–813. → page 15

Bargain, O., O. Donni, and P. Kwenda (2014). Intrahousehold distribution and poverty: Evidence from cote d'ivoire. *Journal of Development Economics 107*, 262–276. → page 15

Bargain, O., L. González, C. Keane, and B. Özcan (2012). Female labor supply and divorce: New evidence from ireland. *European Economic Review 56*(8), 1675–1691. → pages 1, 8, 11

Becker, G. S. and N. Tomes (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of political Economy 87*(6), 1153–1189. → page 74

Becker, G. S. and N. Tomes (1986). Human capital and the rise and fall of families. *Journal of labor economics 4*(3, Part 2), S1–S39. → page 74

Belley, P. and L. Lochner (2007). The changing role of family income and ability in determining educational achievement. *Journal of Human capital 1*(1), 37–89. → page 75

Björklund, A., M. Lindahl, and E. Plug (2006). The origins of intergenerational associations: Lessons from swedish adoption data. *The Quarterly Journal of Economics 121*(3), 999–1028. → pages 74, 75, 98

Black, S. E. and P. J. Devereux (2011). Recent developments in intergenerational mobility. In O. Ashenfelter and D. Autor (Eds.), *Handbook of Labor Economics, Volume 4b*. Elsevier. → page 75

Black, S. E., P. J. Devereux, and K. G. Salvanes (2005). Why the apple doesn't fall far: Understanding intergenerational transmission of human capital. *American economic review 95*(1), 437–449. → page 75

Blanden, J., P. Gregg, and L. Macmillan (2007). Accounting for intergenerational income persistence: noncognitive skills, ability and education. *The Economic Journal 117*(519), C43–C60. → pages 3, 4, 73, 74, 75, 76, 84

Blanden, J., P. Gregg, and L. Macmillan (2013). Intergenerational persistence in income and social class: the effect of within-group inequality. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 176*(2), 541–563. → page 85

Blau, P. M. and O. D. Duncan (1967). *The American occupational structure.* New York: Wiley. → pages 75, 76

Blishen, B. R. (1967). A socio-economic index for occupations in canada. *Canadian Review of Sociology/Revue canadienne de sociologie 4*(1), 41–53. → page 76

Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2007). Collective labour supply: Heterogeneity and non-participation. *The Review of Economic Studies 74*(2), 417–445. → page 15

Blundell, R., P.-A. Chiappori, and C. Meghir (2005). Collective labor supply with children. *Journal of political Economy 113*(6), 1277–1306. → page 15

Bowles, S. and H. Gintis (2002). The inheritance of inequality. *Journal of economic Perspectives 16*(3), 3–30. → pages 3, 4, 72, 73, 74, 75, 76, 98, 117

Brassiolo, P. (2016). Domestic violence and divorce law: When divorce threats become credible. *Journal of Labor Economics 34*(2), 443–477. → page 7

Bray, T. (1999). *The shadow education system: Private tutoring and its implications for planners*. UNESCO International Institute for Educational Planning. → page 75

Bricker, J., A. Henriques, J. Krimmel, and J. Sabelhaus (2016). Measuring income and wealth at the top using administrative and survey data. *Brookings Papers on Economic Activity 2016*(1), 261–331. → pages 100, 115

Browning, M. and P.-A. Chiappori (1998). Efficient intra-household allocations: A general characterization and empirical tests. *Econometrica 66*(6), 1241–1278. → page 15

Browning, M., P.-A. Chiappori, and A. Lewbel (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *Review of Economic Studies 80*(4), 1267–1303. → pages 15, 17, 18, 19

Brzozowski, M., M. Gervais, P. Klein, and M. Suzuki (2010). Consumption, income, and wealth inequality in canada. *Review of economic dynamics 13*(1), 52–75. → pages 101, 115

Burton, P., S. Phipps, and L. Zhang (2014). The prince and the pauper: Movement of children up and down the canadian income distribution. *Canadian Public Policy 40*(2), 111–125. → pages 1, 7

Carneiro, P. and J. J. Heckman (2002). The evidence on credit constraints in post-secondary schooling. *The Economic Journal 112*(482), 705–734. → page 75

Chen, W.-H., Y. Ostrovsky, and P. Piraino (2017). Lifecycle variation, errors-in-variables bias and nonlinearities in intergenerational income transmission: new evidence from canada. *Labour Economics 44*, 1–12. → pages 3, 72

Cherchye, L., B. De Rock, and F. Vermeulen (2012). Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information. *American Economic Review 102*(7), 3377–3405. → page 15

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? The geography of intergenerational mobility in the United States. *The Quarterly Journal of Economics 129*(4), 1553–1623. → pages 3, 72, 78

Chetty, R., N. Hendren, P. Kline, E. Saez, and N. Turner (2014). Is the United States still a land of opportunity? Recent trends in intergenerational mobility. *American Economic Review 104*(5), 141–47. → pages 3, 73

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica 56*(1), 63–89. → pages 12, 14, 15

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of political Economy 100*(3), 437–467. → pages 12, 14, 15, 130

Chiappori, P.-A. and I. Ekeland (2006). The micro economics of group behavior: general characterization. *Journal of Economic Theory 130*(1), 1–26. → page 15

Chiappori, P.-A., B. Fortin, and G. Lacroix (2002). Marriage market, divorce legislation, and household labor supply. *Journal of political Economy 110*(1), 37–72. → pages 3, 10, 12, 15, 36

Chiappori, P.-A., M. Iyigun, and Y. Weiss (2009). Investment in schooling and the marriage market. *American Economic Review 99*(5), 1689–1713. → page 15

Chiappori, P.-A. and M. Mazzocco (2017). Static and intertemporal household decisions. *Journal of Economic Literature 55*(3), 985–1045. → page 15

Chiappori, P.-A., S. Oreffice, and C. Quintana-Domeque (2012). Fatter attraction: anthropometric and socioeconomic matching on the marriage market. *Journal of Political Economy 120*(4), 659–695. → page 14

Chiappori, P.-A. and Y. Weiss (2006). Divorce, remarriage, and welfare: A general equilibrium approach. *Journal of the European Economic Association 4*(2-3), 415–426. → page 15

Chiappori, P.-A. and Y. Weiss (2007). Divorce, remarriage, and child support. *Journal of Labor Economics 25*(1), 37–74. → page 15

Connolly, M., C. Haeck, D. Lapierre, et al. (2019). Social mobility trends in canada: Going up the great gatsby curve. Working paper. → pages 3, 72, 73, 78

Corak, M. (2001). Death and divorce: The long-term consequences of parental loss on adolescents. *Journal of Labor Economics 19*(3), 682–715. → pages 1, 7

Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives 27*(3), 79–102. → pages 3, 4, 72, 73, 74, 75, 79, 88, 98

Corak, M. (2017). Divided landscapes of economic opportunity: The canadian geography of intergenerational income mobility. Technical report, Human Capital and Economic Opportunity Working Group Working Paper 2017-043. → pages 73, 78, 80

Corak, M., L. Curtis, and S. Phipps (2011). Economic mobility, family background, and the well-being of children in the united states and canada. In T. Smeeding, R. Erikson, and M. Jäntti (Eds.), *Persistence, privilege, and parenting: The comparative study of intergenerational mobility*, pp. 73–108. New York: Russell Sage Foundation. → pages 76, 98, 117

Corak, M. and A. Heisz (1999). The intergenerational earnings and income mobility of canadian men: Evidence from longitudinal income tax data. *Journal of Human Resources 34*(3), 504533. → pages 79, 88

Corak, M. and P. Piraino (2011). The intergenerational transmission of employers. *Journal of Labor Economics 29*(1), 37–68. → pages 77, 98

Couprie, H. (2007). Time allocation within the family: Welfare implications of life in a couple. *The Economic Journal 117*(516), 287–305. → pages 15, 18

Couprie, H., E. Peluso, and A. Trannoy (2010). Is power more evenly balanced in poor households? *Journal of Public Economics 94*(7-8), 493–507. → page 15

Cunha, F. and J. Heckman (2007). The technology of skill formation. *American Economic Review 97*(2), 31–47. → page 75

Cunha, F., J. J. Heckman, L. Lochner, and D. V. Masterov (2006). Interpreting the evidence on life cycle skill formation. *Handbook of the Economics of Education 1*, 697–812. → pages 74, 75, 98

Currie, J. and E. Moretti (2003). Mother's education and the intergenerational transmission of human capital: Evidence from college openings. *The Quarterly journal of economics 118*(4), 1495–1532. → page 75

Currie, J. and E. Moretti (2007). Biology as destiny? Short-and long-run determinants of intergenerational transmission of birth weight. *Journal of Labor economics 25*(2), 231–264. → page 75

Curtis, L. J. and K. Rybczynski (2014). Exiting poverty: Does sex matter? *Canadian Public Policy 40*(2), 126–142. → pages 1, 7, 71, 117

Davies, J. B. and L. Di Matteo (2018). Filling the gap: Long run canadian wealth inequality in international context. Technical report, University of Western Ontario, London, ON, Department of Economics Research Reports, 2018-1. → pages 5, 6, 100, 101, 102, 103, 107, 115, 118

Davies, J. B., N. M. Fortin, and T. Lemieux (2017). Wealth inequality: Theory measurement and decomposition. *Canadian Journal of Economics/Revue canadienne d économique 50*(5), 1224–1261. → page 102

De Haan, M. (2011). The effect of parents schooling on childs schooling: a nonparametric bounds analysis. *Journal of Labor Economics 29*(4), 859–892. → page 75

Dickson, M., P. Gregg, and H. Robinson (2016). Early, late or never? When does parental education impact child outcomes? *The Economic Journal 126*(596), F184–F231. → page 75

Diewert, W. E. and E. Yu (2012). A Canadian business sector data base and new estimates of Canadian TFP growth. Technical report, University of British Columbia, Working Paper. → page 104

Donni, O. (2008). Labor supply, home production, and welfare comparisons. *Journal of Public Economics 92*(7), 1720–1737. → page 18

Dunbar, G. R., A. Lewbel, and K. Pendakur (2013). Children's resources in collective households: identification, estimation, and an application to child poverty in malawi. *American Economic Review 103*(1), 438–71. → page 15

Duncan, G. J. and R. J. Murnane (2011). *Whither opportunity? Rising inequality, schools, and children's life chances*. Russell Sage Foundation. → page 75

Erikson, R. and J. H. Goldthorpe (1992). *The constant flux: A study of class mobility in industrial societies*. Oxford University Press, USA. → page 75

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2016). Heterogeneity in returns to wealth and the measurement of wealth inequality. *American Economic Review 106*(5), 651–55. → page 101

Frenette, M. et al. (2017). Which families invest in registered education savings plans and does it matter for postsecondary enrolment? Technical report, Statistics Canada, Analytical Studies Branch. → page 75

Frenette, M., D. Green, and G. Picot (2004). Rising income inequality in the 1990s: An exploration of three data sources. Technical report, Statistics Canada, Analytical Studies Branch. → page 37

Frenette, M., D. A. Green, and K. Milligan (2007). The tale of the tails: Canadian income inequality in the 1980s and 1990s. *Canadian Journal of Economics/Revue canadienne d'économique 40*(3), 734–764. → page 37

Garbinti, B., J. Goupille-Lebret, and T. Piketty (2017). Accounting for wealth inequality dynamics: Methods, estimates and simulations for france (1800-2014). Technical report, Centre for Economic Policy Research, London, UK, Discussion Paper DP11848. → pages 100, 113

Goussé, M., N. Jacquemet, and J.-M. Robin (2017). Marriage, labor supply, and home production. *Econometrica 85*(6), 1873–1919. → page 14

Gray, J. S. (1998). Divorce-law changes, household bargaining, and married women's labor supply. *The American Economic Review 88*(3), 628–642. → pages 2, 8, 11

Green, D. A., W. C. Riddell, and F. St-Hilaire (2016). Income inequality in canada: Driving forces, outcomes and policy. In D. A. Green, W. C. Riddell, and F. St-Hilaire (Eds.), *Income inequality: The Canadian story*, pp. 1–73. Monteal: Institute for Research on Public Policy. → pages 1, 4, 6, 102

Harding, J. F., P. A. Morris, and D. Hughes (2015). The relationship between maternal education and children's academic outcomes: A theoretical framework. *Journal of Marriage and Family 77*(1), 60–76. → page 45

Hemeon, J. (2016). *Historical Data Linkage Quality: The Longitudinal and International Study of Adults, and Tax Records on Labour and Income*. Statistics Canada= Statistique Canada. → page 77

Holmlund, H., M. Lindahl, and E. Plug (2011). The causal effect of parents' schooling on children's schooling: A comparison of estimation methods. *Journal of Economic Literature 49*(3), 615–51. → page 75

Hout, M. (1988). More universalism, less structural mobility: The american occupational structure in the 1980s. *American Journal of sociology 93*(6), 1358–1400. → pages 73, 76, 92

Jacob, H. (1989). Another look at no-fault divorce and the post-divorce finances of women. *Law and Society Review 23*(1), 95–115. → pages 56, 58

Johnson, W. R. and J. Skinner (1986). Labor supply and marital separation. *The American Economic Review 76*(3), 455–469. → pages 2, 8, 11, 55, 60, 61

Johnson, W. R. and J. Skinner (1988). Accounting for changes in the labor supply of recently divorced women. *Journal of Human Resources 23*(4), 417–436. → pages 2, 8, 11, 12

Joung, I. M., K. Stronks, H. Van De Mheen, F. W. van Poppel, J. Van der Meer, and J. Mackenbach (1997). The contribution of intermediary factors to marital status differences in self-reported health. *Journal of Marriage and the Family 59*(2), 476–490. → pages 1, 7, 71, 117

Kalil, A., R. Ryan, and M. Corey (2012). Diverging destinies: Maternal education and the developmental gradient in time with children. *Demography 49*(4), 1361–1383. → page 45

Kaushal, N., K. Magnuson, and J. Waldfogel (2011). How is family income related to investments in childrens learning? In G. J. Duncan and R. J. Murnane (Eds.), *Whither Opportunity? Rising Inequality, Schools, and Children's Life Chances*. New York: Russell Sage Foundation. → page 74

Knudsen, E. I., J. J. Heckman, J. L. Cameron, and J. P. Shonkoff (2006). Economic, neurobiological, and behavioral perspectives on building americas future workforce. *Proceedings of the National Academy of Sciences 103*(27), 10155–10162. → page 75

Kopczuk, W. (2015). What do we know about the evolution of top wealth shares in the united states? *Journal of Economic Perspectives 29*(1), 47–66. → pages 5, 100, 115, 118

Kopczuk, W. and E. Saez (2004). Top wealth shares in the united states, 1916-2000: Evidence from estate tax returns. *National Tax Journal 57*(2), 445–487. → pages 5, 100, 115

Kornrich, S. and F. Furstenberg (2013). Investing in children: Changes in parental spending on children, 1972–2007. *Demography 50*(1), 1–23. → page 74

Kramarz, F. and O. N. Skans (2014). When strong ties are strong: Networks and youth labour market entry. *Review of Economic Studies 81*(3), 1164–1200. → page 77

Kuhn, M., M. Schularick, and U. Steins (2018). Income and wealth inequality in america, 1949-2016. Technical report, Opportunity and Inclusive Growth Institute, Minneapolis, USA, Institute Working Paper 9. → page 102

Lewbel, A. and K. Pendakur (2008). Estimation of collective household models with engel curves. *Journal of Econometrics 147*(2), 350–358. → page 15

Lindquist, M. J., J. Sol, and M. Van Praag (2015). Why do entrepreneurial parents have entrepreneurial children? *Journal of Labor Economics 33*(2), 269–296. → page 77

Lise, J. and S. Seitz (2011). Consumption inequality and intra-household allocations. *The Review of Economic Studies 78*(1), 328–355. → page 15

Low, H., C. Meghir, L. Pistaferri, and A. Voena (2017). Marriage, labor supply and the social safety net. Technical report, NBER Working Paper No. 24356. → page 13

Lucas, S. R. (2001). Effectively maintained inequality: Education transitions, track mobility, and social background effects. *American Journal of Sociology 106*(6), 1642–1690. → page 90

Lundberg, J. and D. Waldenström (2018). Wealth inequality in sweden: What can we learn from capitalized income tax data? *Review of Income and Wealth 64*(3), 517–541. → pages 101, 115

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of political Economy 89*(6), 1059–1085. → pages 16, 24, 28, 29, 49

MaCurdy, T. E. (1985). Interpreting empirical models of labor supply in an intertemporal framework with uncertainty. In J. Heckman and B. Singer (Eds.), *Longitudinal analysis of labor market data*, pp. 111–155. Cambridge: Cambridge University Press. → pages 16, 21, 24, 28, 29, 34

Maddala, G. S. (1983). *Limited-dependent and qualitative variables in econometrics*. Number 3. Cambridge University Press. → pages 43, 44

Martínez-Toledano, C. (2017). Housing bubbles, offshore assets and wealth inequality in spain (1984-2013). Working paper. → pages 100, 101

Mazzocco, M. (2004). Saving, risk sharing, and preferences for risk. *American Economic Review 94*(4), 1169–1182. → page 15

Mazzocco, M. (2007). Household intertemporal behaviour: A collective characterization and a test of commitment. *The Review of Economic Studies 74*(3), 857–895. → page 15

Mazzocco, M., C. Ruiz, and S. Yamaguchi (2013). Labor supply , wealth dynamics , and marriage decisions. Working paper. → pages 13, 14, 21, 71

Michaud, P.-C. and F. Vermeulen (2011). A collective labor supply model with complementarities in leisure: Identification and estimation by means of panel data. *Labour Economics 18*(2), 159–167. → page 15

Milligan, K. (2005). Life-cycle asset accumulation and allocation in canada. *Canadian Journal of Economics/Revue canadienne d'économique 38*(3), 1057–1106. → pages 103, 116

Morissette, R. and X. Zhang (2006). Revisiting wealth inequality. *Perspectives on Labour and Income 7*(12), 5–16. Statistics Canada Catalogue no. 75-001-XIE. → pages 5, 101, 115

Morissette, R., X. Zhang, and M. Drolet (2006). The evolution of wealth inequality in canada, 1984–99. In E. N. Wolff (Ed.), *International perspectives on household wealth*, pp. 151–192. Cheltenham, UK: Edward Elgar Publishing. → pages 5, 101, 115

Mueller, R. E. (2005). The effect of marital dissolution on the labour supply of males and females: Evidence from canada. *The journal of socio-economics 34*(6), 787–809. → pages 2, 8, 11, 55, 62, 63

Mullen, A. L., K. A. Goyette, and J. A. Soares (2003). Who goes to graduate school? Social and academic correlates of educational continuation after college. *Sociology of Education 76*(2), 143–169. → page 90

Neyman, J. and E. L. Scott (1948). Consistent estimates based on partially consistent observations. *Econometrica 16*(1), 1–32. → page 28

Oja, G. (1983). The distribution of wealth in canada. *Review of Income and Wealth 29*(2), 161–173. → pages 5, 101, 115

Oreopoulos, P. (2003). The long-run consequences of living in a poor neighborhood. *The quarterly journal of economics 118*(4), 1533–1575. → page 79

Oreopoulos, P., M. Page, and A. H. Stevens (2008). The intergenerational effects of worker displacement. *Journal of Labor Economics 26*(3), 455–483. → page 79

Oreopoulos, P., M. E. Page, and A. H. Stevens (2006). The intergenerational effects of compulsory schooling. *Journal of Labor Economics 24*(4), 729–760. → page 75

Ornstein, M. D. (1981). The occupational mobility of men in ontario. *Canadian Review of Sociology/Revue canadienne de sociologie 18*(2), 183–215. → pages 76, 77

Parkman, A. M. (1992). Unilateral divorce and the labor-force participation rate of married women, revisited. *The American Economic Review 82*(3), 671–678. → pages 2, 8, 11

Passeron, J.-C. and P. Bourdieu (1964). *Les héritiers: Les étudiants et la culture*. Éditions de Minuit. → page 75

Pendakur, K. (2002). Taking prices seriously in the measurement of inequality. *Journal of Public Economics 86*(1), 47–69. → page 38

Peters, H. E. (1986). Marriage and divorce: Informational constraints and private contracting. *The American Economic Review 76*(3), 437–454. → pages 2, 8, 11, 55, 57, 58, 59, 60

Ramey, G. and V. A. Ramey (2010). The rug rat race. *Brookings Papers on Economic Activity, Spring*, 129–176. → page 75

Reinhold, S., T. Kneip, and G. Bauer (2013). The long run consequences of unilateral divorce laws on childrenevidence from sharelife. *Journal of Population Economics 26*(3), 1035–1056. → page 7

Rivera, L. A. (2016). *Pedigree: How elite students get elite jobs*. Princeton University Press. → page 77

Roine, J. and D. Waldenström (2009). Wealth concentration over the path of development: Sweden 1873-2006. *Scandinavian Journal of Economics 111*(1), 151–187. → page 101

Roine, J. and D. Waldenström (2015). Long-run trends in the distribution of income and wealth. In *Handbook of income distribution*, Volume 2, pp. 469–592. Elsevier. → page 115

Rothstein, J. (2019). Inequality of educational opportunity? Schools as mediators of the intergenerational transmission of income. *Journal of Labor Economics 37*(S1), S85–S123. → page 85

Saez, E. and G. Zucman (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics 131*(2), 519–578. → pages 5, 100, 103, 115

Schneider, D., O. P. Hastings, and J. LaBriola (2018). Income inequality and class divides in parental investments. *American Sociological Review 83*(3), 475–507. → page 75

Solon, G. (1992). Intergenerational income mobility in the united states. *The American Economic Review 82*(3), 393–408. → pages 3, 72, 78

Statistics Canada (n.d.-a). Table 17-10-0060-01 Estimates of population as of July 1st, by marital status or legal marital status, age and sex (table). CANSIM (database). Last updated April 25 2019. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1710006001. → page 36

Statistics Canada (n.d.-b). Table 18-10-0003-01 Inter-city indexes of price differentials of consumer goods and services, annual (table). CANSIM (database). Last updated July 17 2019. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1810000301. → page 38

Statistics Canada (n.d.-c). Table 18-10-0004-01 Consumer Price Index, monthly, not seasonally adjusted (table). CANSIM (database). Last updated September 18 2019. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1810000401. → page 38

Statistics Canada (n.d.-d). Table 36-10-0107-01 Household final consumption expenditure, quarterly, Canada (table). CANSIM (database). Last updated August 30 2019. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610010701. → page 104

Statistics Canada (n.d.-e). Table 36-10-0510-01 National balance sheet, persons and unincorporated business, annual, 1961 - 2011 (table). CANSIM (database). Last updated October 15 2012. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610051001. → page 104

Statistics Canada (n.d.-f). Table 39-10-0023-01 Age-specific divorce rates per 1,000 legally married females (table). CANSIM (database). Last updated December 17 2015. https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3910002301. → pages 1, 7

Stevenson, B. (2007). The impact of divorce laws on marriage-specific capital. *Journal of Labor Economics 25*(1), 75–94. → pages 7, 11, 55, 61, 62

Stevenson, B. and J. Wolfers (2006). Bargaining in the shadow of the law: Divorce laws and family distress. *The Quarterly Journal of Economics 121*(1), 267–288. → page 7

Torche, F. (2011). Is a college degree still the great equalizer? Intergenerational mobility across levels of schooling in the United States. *American Journal of Sociology 117*(3), 763–807. → pages 73, 76, 90, 92

Voena, A. (2015). Yours, mine, and ours: Do divorce laws affect the intertemporal behavior of married couples? *American Economic Review 105*(8), 2295–2332. → pages 3, 10, 12, 13, 15, 20

Waldfogel, J. and E. Washbrook (2011). Early years policy. *Child Development Research 2011*. → page 75

Wanner, R. A. and B. C. Hayes (1996). Intergenerational occupational mobility among men in canada and australia. *Canadian Journal of Sociology/Cahiers canadiens de sociologie 21*(1), 43–76. → pages 73, 76, 92

Wolff, E. (2017). Deconstructing household wealth trend in the united stetes, 1983 to 2016. For presentation at the First WID World Conference, Paris, France. → page 102

Zarifa, D. (2012a). Choosing fields in an expansionary era: Comparing two cohorts of baccalaureate degree-holders in the united states and canada. *Research in Social Stratification and Mobility 30*(3), 328–351. → pages 75, 90

Zarifa, D. (2012b). Persistent inequality or liberation from social origins? Determining who attends graduate and professional schools in Canada's expanded postsecondary system. *Canadian Review of Sociology/Revue canadienne de sociologie 49*(2), 109–137. → page 90

Zimmerman, D. J. (1992). Regression toward mediocrity in economic stature. *The American Economic Review 82*(3), 409–429. → pages 3, 72, 78

Appendix A

Appendix to Chapter 1

A.1 Labour supply equations: derivations

A.1.1 Intertemporal decision-making of the divorced household

Let $Q_{1,i,t} = \{l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}, A_{1,i,t+1}\}$, the set of choice variables for spouse 1 at time t when they are divorced. Then, their problem is

$$\max_{\{Q_{1,i,t}\}_{t=0}^{T}} E_0 \sum_{t=0}^{T} \beta^t U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t} C_{1,i,t})$$

subject to the budget constraint

$$\mathrm{w}_{1,i,t}(l_{1,i,t} + L_{1,i,t}) + c_{1,i,t} + p_t C_{1,i,t} + A_{1,i,t+1} = \mathrm{w}_{1,i,t} + Y_{i,t} + (1 + r_t) A_{1,i,t}$$

Alternatively, the problem can be expressed recursively, defining the Bellman equation as:

$$V^D_{1,i}(A_{1,i}) = \max_{Q_{1,i}} U_{1,i}(l_{1,i}, L_{1,i}, c_{1,i}, C_{1,i}) + \beta E V^D_{1,i}(A'_{1,i})$$

subject to the budget constraint:

$$\mathrm{w}_{1,i}(l_{1,i} + L_{1,i}) + c_{1,i} + pC_{1,i} + A'_{1,i} = \mathrm{w}_{1,i} + Y_i + (1 + r) A_{1,i}$$

Hence, using the same Stone-Geary utility function yields the Lagrangian:

$$\begin{aligned}\mathcal{L} = {} & \gamma_{c,1,i} ln(c_{1,i} - \bar{c}_{1,i}) + \gamma_{C,1,i} ln(C_{1,i} - \bar{C}_{1,i}) + \gamma_{l,1,i} ln(l_{1,i} - \bar{l}_{1,i}) + \gamma_{L,1,i} ln(L_{1,i} - \bar{L}_{1,i}) \\ & + \beta E V^D_{1,i}(A'_{1,i}) + \nu_{1,i}\left(\mathrm{w}_{1,i} + Y_i + (1 + r) A_{1,i} - \mathrm{w}_{1,i}(l_{1,i} + L_{1,i}) - c_{1,i} - pC_{1,i} - A'_{1,i}\right)\end{aligned}$$

The first-order conditions for $t = 0, \ldots, T$ are then:

$$l_{1,i,t} : \gamma_{l,1,i} \frac{1}{l_{1,i,t} - \bar{l}_{1,i,t}} = \nu_{1,i,t} \mathrm{w}_{1,i,t}$$

$$L_{1,i,t} : \gamma_{L,1,i} \frac{1}{L_{1,i,t} - \bar{L}_{1,i,t}} = \nu_{1,i,t} \mathrm{w}_{1,i,t}$$

$$c_{1,i,t} : \gamma_{c,1,i} \frac{1}{c_{1,i,t} - \bar{c}_{1,i,t}} = \nu_{1,i,t}$$

$$C_{1,i,t} : \gamma_{1,i,t} \frac{1}{C_{1,i,t} - \bar{C}_{1,i,t}} = \nu_{1,i,t} p_t$$

$$A_{1,i,t+1} : \beta E_t \left[\omega_{1,i,t+1} \frac{\partial V^D_{1,i}(A_{1,i,t+1})}{\partial A_{1,i,t+1}} \right] = \nu_{1,i,t} \tag{A.1}$$

$$\nu_{1,i,t} : \mathrm{w}_{1,i,t}(l_{1,i,t} + L_{1,i,t}) + c_{1,i,t} + p_t C_{1,i,t} + A_{1,i,t+1} = \mathrm{w}_{1,i,t} + Y_{1,i,t} + (1 + r_t) A_{1,i,t}$$

where $\nu_{1,i,t}$ is the Lagrange multiplier. As usual, $\frac{\partial V^D_{1,i}(A_{1,i,t+1})}{\partial A_{1,i,t+1}}$ is unknown but the Envelope theorem can be used to find its value.

$$\begin{aligned}
\frac{\partial V^D_{1,i}(A_{1,i})}{\partial A_{1,i}} &= \beta E \frac{\partial V^D_{1,i}(A'_{1,i})}{\partial A'_{1,i}} \frac{\partial A'_{1,i}}{\partial A_{1,i}} + \nu_{1,i} \left((1+r) - \frac{\partial A'_{1,i}}{\partial A_{1,i}} \right) \\
&= \frac{\partial A'_{1,i}}{\partial A_{1,i}} \left[\beta E \frac{\partial V^D_{1,i}(A'_{1,i})}{\partial A'_{1,i}} - \nu_{1,i} \right] + \nu_{1,i}(1+r)
\end{aligned}$$

in which from Equation A.1 we know that $\beta E \frac{\partial V^D_{1,i}(A'_{1,i})}{\partial A'_{1,i}} - \nu_{1,i} = 0$, and so:

$$\frac{\partial V^D_{1,i}(A'_{1,i})}{\partial A'_{1,i}} = \nu'_{1,i}(1 + r')$$

Therefore, Equation A.1 becomes

$$\beta E_t \left[\nu_{1,i,t+1}(1 + r_{t+1}) \right] = \nu_{1,i,t}$$

A.1.2 Intertemporal decision-making of the married household

From Chiappori (1992), remember that the household problem can be solved as a problem faced by one individual in the household, constrained both by the household's budget and the fact that his or her spouse cannot be made worse off. Browning et al. (2014) show that this set-up is easily extended to the intertemporal context. Adding assets to their specification yields, and defining $Q_{i,t} = \{l_{1,i,t}, l_{2,i,t}, L_{1,i,t}, L_{2,i,t}, c_{1,i,t}, c_{2,i,t}, C_{1,i,t}, C_{2,i,t}, A_{i,t+1}\}$, the set of choice variables for the household, the married household's problem is:

$$V_i^M(A_{i,t}) = \max_{\{Q_{i,t}\}_{t=0}^T} E_0 \sum_{t=0}^{T} \beta^t U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) \tag{A.2}$$

subject to the constraints:

$$E_0 \sum_{t=0}^{T} \beta^t U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) \geq \bar{U}_{2,i,t}$$

$$\begin{aligned} \mathrm{w}_{1,i,t} l_{1,i,t} + \mathrm{w}_{2,i,t} l_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) \kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t \kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1 + r_t) A_{i,t} \end{aligned}$$

where we've used the consumption and leisure technology functions.

This is equivalent to

$$V_i^M(A_{i,t}) = \max_{\{Q_{i,t}\}_{t=0}^T} E_0 \sum_{t=0}^{T} \beta^t U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) + \mu_{i,0} E_0 \sum_{t=0}^{T} \beta^t U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) \tag{A.3}$$

subject to the budget constraint:

$$\begin{aligned} \mathrm{w}_{1,i,t} l_{1,i,t} + \mathrm{w}_{2,i,t} l_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) \kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t \kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1 + r_t) A_{i,t} \end{aligned}$$

where $\mu_{i,0}$ is the Lagrange multiplier on the Pareto optimality constraint.

In turn, the weights on each household member's utility can be reconstructed for a more intuitive interpretation:

$$V_i^M(A_{i,t}) = \max_{\{Q_{i,t}\}_{t=0}^T} \lambda_{i,0} E_0 \sum_{t=0}^{T} \beta^t U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) + (1 - \lambda_{i,0}) E_0 \sum_{t=0}^{T} \beta^t U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) \tag{A.4}$$

subject to the budget constraint:

$$\begin{aligned} \mathrm{w}_{1,i,t} l_{1,i,t} + \mathrm{w}_{2,i,t} l_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) \kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t \kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1 + r_t) A_{i,t} \end{aligned}$$

Alternatively, this can be expressed using Bellman equations

$$V_i^M(A) = \max_{Q_i} \lambda_i \left[U_{1,i}(l_{1,i}, L_{1,i}, c_{1,i}, C_{1,i}) + \beta E V_{1,i}^M(A_i')\right] + (1-\lambda_i)\left[U_{2,i}(l_{2,i}, L_{2,i}, c_{2,i}, C_{2,i}) + \beta E V_{2,i}^M(A_i')\right]$$

subject to the budget constraint:

$$\begin{aligned}
\mathrm{w}_{1,i} l_{1,i} + \mathrm{w}_{2,i} l_{2,i} + (\mathrm{w}_{1,i} + \mathrm{w}_{2,i}) g_{\Lambda,i}(L_{1,i} + L_{2,i}) - (\mathrm{w}_{1,i} + \mathrm{w}_{2,i}) \kappa_{\Lambda,i} \\
+ c_{1,i} + c_{2,i} + p_t g_{K,i}(C_{1,i} + C_{2,i}) - p\kappa_{K,i} + A_i' \\
= \mathrm{w}_{1,i} + \mathrm{w}_{2,i} + Y_i + (1+r)A_i
\end{aligned}$$

Finally, this problem must also account for the fact that next period's problem changes if a person becomes divorced:

$$\begin{aligned}
V_i^M(A_i) = \max_{Q_i} \lambda_i \left[U_{1,i}(l_{1,i}, L_{1,i}, c_{1,i}, C_{1,i}) + \beta E\left(\omega_i' V_{1,i}^M(A_i') + (1-\omega_i') V_{1,i}^D(A_i')\right)\right] \\
+ (1-\lambda_i)\left[U_{2,i}(l_{2,i}, L_{2,i}, c_{2,i}, C_{2,i}) + \beta E\left(\omega_i' V_{2,i}^M(A_i') + (1-\omega_i') V_{2,i}^D(A_i')\right)\right]
\end{aligned}$$

subject to the budget constraint:

$$\begin{aligned}
\mathrm{w}_{1,i} l_{1,i} + \mathrm{w}_{2,i} l_{2,i} + (\mathrm{w}_{1,i} + \mathrm{w}_{2,i}) g_{\Lambda,i}(L_{1,i} + L_{2,i}) - (\mathrm{w}_{1,i} + \mathrm{w}_{2,i}) \kappa_{\Lambda,i} \\
+ c_{1,i} + c_{2,i} + p_t g_{K,i}(C_{1,i} + C_{2,i}) - p\kappa_{K,i} + A_i' \\
= \mathrm{w}_{1,i} + \mathrm{w}_{2,i} + Y_i + (1+r)A_i
\end{aligned}$$

where ω_i' is the probability that the household is still together next period, and

$$V_{s,i}^M(A_i) = \max_{Q_i} U_{s,i}(l_{s,i}, L_{s,i}, c_{s,i}, C_{s,i}) + \beta E\left(\omega_i' V_{s,i}^M(A_i') + (1-\omega_i') V_{s,i}^D(A_i')\right)$$

for $s = 1,2$. Note that $V_{1,i}^D(A_i)$ and $V_{2,i}^D(A,i)$ are just the problems faced by the divorced households and described previously.

Next, assume the same Stone-Geary utility function and consumption technology function. In addition, note that implicit in this specification is that assets are treated as public, in the sense that both household members contribute to and can draw from them: $A = A_1 + A_2$.

The corresponding first-order conditions are

$$l_{1,i,t} : \lambda_{i,t} \gamma_{l,1,i} \frac{1}{l_{1,i,t} - \bar{l}_{1,i,t}} = \nu_{i,t} \mathrm{w}_{1,i,t}$$

$$l_{2,i,t} : (1-\lambda_{i,t}) \gamma_{l,2,i} \frac{1}{l_{2,i,t} - \bar{l}_{2,i,t}} = \nu_{i,t} \mathrm{w}_{2,i,t}$$

$$L_{1,i,t}: \lambda_{i,t}\gamma_{L,1,i}\frac{1}{L_{1,i,t}-\bar{L}_{1,i,t}} = \nu_{i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}$$

$$L_{2,i,t}: (1-\lambda_{i,t})\gamma_{L,2,i}\frac{1}{L_{2,i,t}-\bar{L}_{2,i,t}} = \nu_{i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}$$

$$c_{1,i,t}: \lambda_{i,t}\gamma_{c,1,i}\frac{1}{c_{1,i,t}-\bar{c}_{1,i,t}} = \nu_{i,t}$$

$$c_{2,i,t}: (1-\lambda_{i,t})\gamma_{c,2,i}\frac{1}{c_{2,i,t}-\bar{c}_{2,i,t}} = \nu_{i,t}$$

$$C_{1,i,t}: \lambda_{i,t}\gamma_{C,1,i}\frac{1}{C_{1,i,t}-\bar{C}_{1,i,t}} = \nu_{i,t}p_t$$

$$C_{2,i,t}: (1-\lambda_{i,t})\gamma_{C,2,i}\frac{1}{C_{2,i,t}-\bar{C}_{2,i,t}} = \nu_{i,t}p_t$$

$$\begin{aligned} A_{i,t+1}: &\lambda_{i,t}\left[\beta E_t\left(\omega_{i,t+1}\frac{\partial V_{1,i}^M(A_{i,t+1})}{\partial A_{i,t+1}} + (1-\omega_{i,t+1})\frac{\partial V_{1,i}^D(A_{i,t+1})}{\partial A_{i,t+1}}\right)\right] \\ &+ (1-\lambda_{i,t})\left[\beta E_t\left(\omega_{i,t+1}\frac{\partial V_{2,i}^M(A_{2,i,t+1})}{\partial A_{2,i,t+1}} + (1-\omega_{i,t+1})\frac{\partial V_{2,i}^D(A_{2,i,t+1})}{\partial A_{2,i,t+1}}\right)\right] = \nu_{i,t} \end{aligned} \tag{A.5}$$

$$\begin{aligned} \nu_{i,t}: \mathrm{w}_{1,i,t}l_{1,i,t} + \mathrm{w}_{2,i,t}l_{2,i,t} + (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}(L_{1,i,t}+L_{2,i,t}) - (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})\kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t}+C_{2,i,t}) - p_t\kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1+r_t)A_{i,t} \end{aligned}$$

The terms $\frac{\partial V_{1,i}^M(A_{i,t+1})}{\partial A_{i,t+1}}$, $\frac{\partial V_{1,i}^D(A_{i,t+1})}{\partial A_{i,t+1}}$, $\frac{\partial V_{2,i}^M(A_{i,t+1})}{\partial A_{i,t+1}}$ and $\frac{\partial V_{2,i}^D(A_{i,t+1})}{\partial A_{i,t+1}}$ in Equation A.5 are unknown. However, they can be obtained as usual by relying on the Envelope theorem. In particular, $\frac{\partial V_{1,i}^D(A_i')}{\partial A_i'}$ and $\frac{\partial V_{2,i}^D(A_i')}{\partial A_i'}$ come from the respective problems of the two household members, which yield:

$$\frac{\partial V_{1,i}^D(A_i')}{\partial A_i'} = \nu_{1,i}'(1+r') \tag{A.6}$$

$$\frac{\partial V_{2,i}^D(A_i')}{\partial A_i'} = \nu_{2,i}'(1+r') \tag{A.7}$$

Therefore, all that remains to be found is

$$\omega_i' \left(\lambda_i \frac{\partial V_{1,i}^M(A_i')}{\partial A_i'} + (1-\lambda_i) \frac{\partial V_{2,i}^M(A_i')}{\partial A_i'} \right) \tag{A.8}$$

which is equal to

$$\omega_i' \frac{\partial V_i^M(A_i')}{\partial A_i'} \tag{A.9}$$

The problem of the married households yields the following

$$\begin{aligned}
\frac{\partial V_i^M(A_i)}{\partial A_i} = {} & \lambda_i \left[\beta E \left(\omega_i' \frac{\partial V_{1,i}^M(A_i')}{\partial A_i'} \frac{\partial A_i'}{\partial A_i} + (1-\omega_i') \frac{\partial V_{1,i}^D(A_i)}{\partial A_i} \frac{\partial A_i'}{\partial A_i} \right) \right] \\
& + (1-\lambda_i) \left[\beta E \left(\omega_i' \frac{\partial V_{2,i}^M(A_i')}{\partial A_i'} \frac{\partial A_i'}{\partial A_i} + (1-\omega_i') \frac{\partial V_{2,i}^D(A_i)}{\partial A_i} \frac{\partial A_i'}{\partial A_i} \right) \right] \\
& + \nu_i \left((1+r) - \frac{\partial A_i'}{\partial A_i} \right)
\end{aligned}$$

$$\begin{aligned}
\frac{\partial V_i^M(A_i)}{\partial A_i} = {} & \frac{\partial A_i'}{\partial A_i} \left(\lambda_i \left[\beta E \left(\omega_i' \frac{\partial V_{1,i}^M(A_i')}{\partial A_i'} + (1-\omega_i') \frac{\partial V_{1,i}^D(A_i)}{\partial A_i} \right) \right] \right. \\
& \left. + (1-\lambda_i) \left[\beta E \left(\omega_i' \frac{\partial V_{2,i}^M(A_i')}{\partial A_i'} + (1-\omega_i') \frac{\partial V_{2,i}^D(A_i)}{\partial A_i} \right) - \nu_i \right] \right) \\
& + \nu_i (1+r)
\end{aligned} \tag{A.10}$$

From Equation A.5, the first term on the right-hand side of Equation A.10 is equal to zero, so that:

$$\frac{\partial V_i^M(A_i)}{\partial A_i} = \nu_i (1+r)$$

and therefore:

$$\frac{\partial V_i^M(A_i')}{\partial A_i'} = \nu_i' (1+r')$$

Substituting this back into the first-order condition for A', along with Equation A.6 and Equation A.7 iterated one period forward, yields:

$$A' : \beta E \left[\omega_i' \nu_i' (1+r') + (1-\omega_i') \{ \lambda_i \nu_{1,i}' + (1-\lambda_i) \nu_{2,i}' \} (1+r') \right] = \nu_i$$

Denote the labour supply of spouse 1, $N_{1,i,t} = 1 - l_{1,i,t} - L_{1,i,t}$. Then,

$$N_{1,i,t} = \bar{N}_{1,i,t} - \frac{\lambda_{i,t}\gamma_{l,1,i}}{\nu_{i,t}\mathrm{w}_{1,i,t}} - \frac{\lambda_{i,t}\gamma_{L,1,i}}{\nu_{i,t}g_{\Lambda,i,t}\left(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t}\right)}$$

A.2 Data

A.2.1 Development of the family files

To understand the development of the family files, it is useful first to discuss the way in which personal income tax statements are used by Statistics Canada to create the T1 Family Files. In Canada, individuals file their taxes by submitting a T1 form to the Canada Revenue Agency (CRA). As part of filling their T1, Canadians are asked to submit not only their Social Insurance Number (SIN), but also their spouse's SIN, or SPSIN, and their marital status. When Statistics Canada receives T1 data from the CRA, it matches spouses' T1s, and constructs tax families, identified by a Family Identification Number, or FIN. These families form the basis of the T1 Family Files, or T1FF. Spouses are matched based on a complex algorithm that starts off looking for perfectly consistent SIN-SPSIN pairs; that is, pairs of people where each reported their SIN and SPSIN so as to form a coherent match. Where there are incongruencies, or where one or both spouses in an actual couple failed to report their partner's SIN, Statistics Canada relies on a combination of demographic variables, including gender and age, and on addresses to form matches. Importantly, that process is done every year independently from matches made in previous years; i.e., information from a given year is never used in subsequent years to aid in the matching of spouses. The T1 Family Files also includes dependants, usually but not exclusively corresponding to children. If children are themselves tax filers, the children's SIN is also listed as part of the parent's FIN. However, non-earning children are also linked to their parents, using information contained in parents' filing for Child Benefits. Note that children may be of any age, so long as they display a fiscal relationship with their matched parents. The result is then a tax family, similar to the concept of census family, but constructed based on fiscal relationships.

As part of the original development of the LISA and the administrative files, Statistics Canada obtained SINs for all LISA respondents that had consented to the administrative linkage. This list was they used to obtain their T1FF, as well as their T4, PPIC and IMDB information. It also constituted the starting point for the family files. For each year in which a respondent was observed in the T1FF, their FIN was retrieved, as well as the list of all other SINs observed in their FIN in that year. This produced a list of SINs, representing all individuals who were ever in the tax family of a LISA respondent, between 1982 and 2013.

As mentioned before, the T1FF is a mostly cross-sectional file. It has been used longitudinally before, namely in the construction of the Longitudinal Administrative Databank (LAD). However, the unit of observation in the LAD is the individual. That is, even though the LAD provides information on the family structure of sample members, it is built around individuals; i.e., it follows individuals over time, not families. For the purpose of this paper, but also for other research interested in complex family dynamics, it was important to construct the LISA family files so that families could be followed over time. The main challenges in doing so stemmed from the fact that tax records are not primarily collected for research purposes, and present corresponding limitations. First, people may change SIN over their lifetime; typically because they lost their original number, or because they were given a temporary number as immigrants. The latter case is further complicated by the fact that immigrant spouses may not transition to their new SIN at the same moment. In constructing the family files, it was therefore necessary to track down old SIN numbers, so that changes in identification were not mistaken for breaks in family relationships. Second, and as will be discussed further next, the cross-sectional nature of T1FF processing implies that simple processing mistakes can be mistaken for changes in family structure.

A.2.2 Marital trajectories

While marital status is considered an important determinant of several outcomes of interest to economists, it is not an innocuous variable. First, a person may have more than one marital status. For instance, one may be all at once legally married, separated from their spouse, and in a common-law relationship with a third person. In that context, a question arises as to whether they should be considered married, separated, or in a common-law relationship; i.e., which status is most relevant for the understanding of their choices. Second, this is further complicated by the fact that in most surveys and even administrative data sets, people are asked to report their marital status. As a result, what they report is informative not only of their actual situation but also of how they perceive their situation. Finally, people may be inclined to misreport or may simply not remember.

In taxes, marital status is self-reported and recorded as *mstco*; however, and as might be expected, it is evident from the data that people do not always report accurately. For instance, it is not uncommon for people to claim being ”single, never married” after having been married for a number of years; or for people to switch at a high frequency between married and divorced/widowed. The fact that people may both report a spousal SIN and fail to report being married suggests that there is more to it than trying to minimize tax payments. The ambiguity of marital status and the complexity of eliciting the information from people also comes through in the marital history section of the LISA. For example, some people report overlapping marriages; others specify a start date on a previous marriage without indicating the end date or how it ended. Unfortunately, there is little evidence on how these inconsistencies should be interpreted or dealt with.

The linkage of LISA respondents to their past and current family members is exploited to pin down marital trajectories as best as possible. This is done in two steps: first, T1FF processing variables are used to produce a trajectory of the family roster from 1982 to 2013; and second, inconsistencies are removed from these trajectories by combining information across years.

Establishing baseline trajectories. T1FF processing yields three variables that inform the structure of families, *rctp* (Record type), *fflag* (Family flag) and *fcmp* (Family composition). The record type indicates whether the individual is a tax filer or not, in which case they had to be imputed, and whether they are living or deceased; the family flag documents the position of the individual in the family, that is, whether they are an adult or a child and, if the former, whether they are unmatched, married or in a common-law relationship; and family composition records whether the family the individual is a part of is single-headed and whether there are children present. The algorithm by which the T1FF processing team produces these variables is complex and highly non-linear; however, it suffices to know that it uses the spousal SIN reported by tax filers on the T1 form, the children's SINs used to claim CTB, as well as individual characteristics, such as age and address.

All three T1FF processing variables are used to produce a baseline longitudinal family roster for family respondents, for the years 1982 to 2013. Relationships are established every year by comparing the *fflag*, *rctp* and *fcmp* for each pair or family members. Recall that the family members that can be informed with this data are people who have a form of economic dependence from a fiscal point of view. For instance, spouses who file taxes together, or a child for whom a parent claims the Child Tax Benefit. As a result, relationships are essentially unambiguous: adult-adult pairs are recorded as couples, adult-child pairs, as parents and children, and child-child pairs as siblings.

This step counts with two important limitations, which highlight the importance of the longitudinal linkage with family members. First, while the T1FF processing algorithm is impressive, it nonetheless results in a number of unmatched filers, a portion of whom are likely to be people who do live in families. Combined with years where all or some family members do not file their taxes, this results in patchy trajectories. Second, the T1FF processing algorithm can also yield false positives, where two people are matched who aren't really part of the same family. This has two particular manifestations in the data. Over time, family members may be recorded as having more than one type of (mutually exclusive) relationship with LISA respondents. For instance, a person may be categorized as a parent in one year and a spouse in another. Alternatively, LISA respondents sometimes appear to have an unlikely high number of past spouses, who are often in their family for only one period. These limitations are respectively addressed by the next step.

Exploiting the longitudinal nature of the roster. In a second step, the information from multiple years is used to fix the two problems outlined previously. First, where relationships conflict the most prevalent relationship is assumed. Using the information available over the 32 years spanned by the T1FF, I calculate the total number of years for each relationship type. The relationship is then recoded as corresponding to the relationship type that appears most often. For instance, if someone is coded as the spouse of a LISA respondent from 1982 to 1988, as their parent in 1989, then as their spouse again, their final type is recorded as spouse. Second, where a relationship could not be established between two individuals in a given year, data from other years is used. This is most straightforward for parents, children, and siblings, relationships which are set for life, even though bonds may vary in strength across families and over time. Any period where a person is an unmatched filer in a LISA respondent's family but appears before, after, or in between observations where they were identified as either of these family members, can be confidently recoded. Spousal relationships are slightly more complicated because they are naturally subject to change. A person who later becomes a spouse is coded as a potential spouse in any period previous to that first occurrence where they are observed in the LISA respondent's family. Where a spouse appears as an unmatched filer in between periods were they were coded as a spouse, they are assumed to be married. While it is possible for a married couple to divorce, reacquaint themselves and re-marry, it seems reasonable to assume these cases as sufficiently rate. Finally, when a person appears as an unmatched filer after having appeared last as a spouse, they are assumed to still be married, up to the point where they can be identified as either deceased or where they appear in a different household. In the latter case, they are recorded as separated/divorced. Recall that unmatched years refer to years where a person is observed in a LISA respondent's household, but their relationship to that individual could not be ascertained through T1FF processing.

A.3 Marginal utility of wealth: derivations

A.3.1 Marginal utility of wealth in the married problem, $\nu_{i,t}$

When household member 1 is still married they must consider the possibility that they or their spouse may dissolve the household in the following period. As a result, $\nu_{i,t}$ depends on $E(\nu_{i,t+1})$, but also on $E(\nu_{1,i,t+1})$ and $E(\nu_{2,i,t+1})$. To see this, recall the first-order condition for assets in the married problem

$$\nu_{i,t} = \beta E_t \left\{ (1+r_{t+1}) \left[\omega_{i,t+1} \nu_{i,t+1} + (1-\omega_{i,t+1}) \{ \lambda_{i,t} \nu_{1,i,t+1} + (1-\lambda_{i,t}) \nu_{2,i,t+1} \} \right] \right\}$$

Let

$$\nu_{i,t} = \exp(\ln \nu_{i,t}) \tag{A.11}$$

and define

$$\ln \nu_{i,t} = E_{t-1}(\ln \nu^*_{i,t}) + \varepsilon^*_{i,t} \tag{A.12}$$

where $\nu^*_{i,t} = \omega_{i,t} \nu_{i,t} + (1-\omega_{i,t})[\lambda_{i,t-1} \nu_{1,i,t} + (1-\lambda_{i,t-1}) \nu_{2,i,t}]$, and $\varepsilon^*_{i,t}$ is the one-period forecast error from unanticipated realizations of wages, prices, interest rates, and leisure and consumption needs. It follows that

$$\begin{aligned}
\varepsilon_{i,t}^{*} &= \ln \nu_{i,t} - E_{t-1}\left(\ln \nu_{i,t}^{*}\right) \\
\exp(\varepsilon_{i,t}^{*}) &= \frac{\exp(\ln \nu_{i,t})}{\exp(E_{t-1}\left(\ln \nu_{i,t}^{*}\right))} \\
E_{t-1}\left(\exp(\varepsilon_{i,t}^{*})\right) &= E_{t-1}\left(\frac{\exp(\ln \nu_{i,t})}{\exp(E_{t-1}(\ln \nu_{i,t}^{*})}\right) \\
E_{t-1}\left(\exp(\varepsilon_{i,t}^{*})\right) &= \frac{E_{t-1}\left(\nu_{i,t}\right)}{\exp(E_{t-1}\ln \nu_{i,t}^{*})} \\
E_{t-1}\left(\nu_{i,t}\right) &= \exp\left[E_{t-1}\left(\ln \nu_{i,t}^{*}\right)\right] E_{t-1}\left(\exp(\varepsilon_{i,t}^{*})\right) \\
E_{t-1}\left(\nu_{i,t}\right) &= \frac{\exp(\ln \nu_{i,t})}{\exp(\varepsilon_{i,t}^{*})} E_{t-1}\left(\exp(\varepsilon_{i,t}^{*})\right) \\
\nu_{i,t} &= \frac{E_{t-1}(\nu_{i,t})\exp(\varepsilon_{i,t}^{*})}{E_{t-1}\left(\exp(\varepsilon_{i,t}^{*})\right)} \qquad \text{(A.13)}
\end{aligned}$$

Assuming that there is no uncertainty with respect to interest rates and the probabilities of being part of the household next period, the first-order condition for assets gives that

$$\nu_{i,t} = \beta\left(1 + r_{t+1}\right)\left\{\omega_{i,t+1} E_t\left(\nu_{i,t+1}\right) + \left(1 - \omega_{i,t+1}\right)\left[\lambda_{i,t} E_t\left(\nu_{1,i,t+1}\right) + \left(1 - \lambda_{i,t}\right) E_t\left(\nu_{2,i,t+1}\right)\right]\right\} \qquad \text{(A.14)}$$

Note that $E_t(\nu_{1,i,t+1})$ and $E_t(\nu_{2,i,t+1})$ are evaluated with subsets of the same information that is used by household members to form $E_t(\nu_{i,t+1})$. Hence, define

$$\begin{aligned}
E_t(\nu_{1,i,t+1}) &= \phi_{1,t} E_t(\nu_{i,t+1}) \\
E_t(\nu_{2,i,t+1}) &= \phi_{2,t} E_t(\nu_{i,t+1})
\end{aligned}$$

That is, $\phi_{s,i,t}$ determines the comparison between the expected value of the marginal utility of wealth in the married couple and the marginal utility of wealth of spouse s if they were to be divorced.

Equation A.14 can therefore be rewritten as

$$\nu_{i,t} = \beta\left(1 + r_{t+1}\right)\left\{\omega_{i,t+1} + \left(1 - \omega_{i,t+1}\right)\left[\lambda_{i,t}\phi_{1,i,t} + \left(1 - \lambda_{i,t}\right)\phi_{2,i,t}\right]\right\} E_t\left(\nu_{i,t+1}\right)$$

Re-organizing and lagging one period gives

$$E_{t-1}(\nu_{i,t}) = \frac{\nu_{i,t-1}}{\beta(1 + r_t)\Omega_{i,t}}$$

where $\Omega_{i,t} = \omega_{i,t} + (1-\omega_{i,t})[\lambda_{i,t-1}\phi_{1,i,t-1} + (1-\lambda_{i,t-1})\phi_{2,i,t-1}]$.

Substituting this back into Equation A.13 gives

$$v_{i,t} = \frac{v_{i,t-1}}{\beta(1+r_t)\Omega_{i,t}} \frac{\exp(\varepsilon^*_{i,t})}{E_{t-1}\left(\exp(\varepsilon^*_{i,t})\right)}$$

or

$$v_{i,t} = v_{i,t-1}\left(\Omega_{i,t}\right)^{-1}\left(\exp(b^*_{i,t})\right)^{-1}\exp(\varepsilon^*_{i,t}) \tag{A.15}$$

where $b^*_{i,t} = \ln(\beta(1+r_t)) + \ln(E_{t-1}(\exp(\varepsilon^*_{i,t})))$.

Taking the logs of both sides, and repeatedly substituting for increasingly older values of $v_{i,t}$ yields

$$\ln v_{i,t} = \ln v_{i,0} - \sum_{j=1}^{t} b^*_{i,j} - \sum_{j=1}^{t}\ln\Omega_{i,t} + \sum_{j=1}^{t}\varepsilon^*_{i,j} \tag{A.16}$$

where $v_{i,0}$ is the marginal utility of wealth at time of marriage. From Section A.3,

$$\ln v_{1,i,t} = \ln v_{i,t^d-1} - \sum_{j=t^d}^{t} b_{1,i,j} + \sum_{j=t^d}^{t}\varepsilon_{1,i,j} \tag{A.17}$$

If women foresee the possibility of divorce (the baseline case), then v_{i,t^d-1} is just the value passed down from marriage; i.e.,

$$\ln v_{i,t^d-1} = \ln v_{i,0} \tag{A.18}$$

Under myopia however, there is a break between the married and the divorced problem and that equality does not hold. Instead, the divorced problem is evaluated for the first time at the time of dissolution.

A.3.2 Marginal utility of wealth in the divorced problem, $v_{1,i,t}$

An expression for $v_{1,i,t}$ can be obtained similarly. Let

$$v_{1,i,t} = \exp(\ln v_{1,i,t})$$

and

$$\ln v_{1,i,t} = E_{t-1}\left(\ln v_{1,i,t}\right) + \varepsilon_{1,i,t} \tag{A.19}$$

where $\varepsilon_{1,i,t}$ is the one-period forecast error for the single problem.

Equation A.19 can be rewritten as follows

$$\begin{aligned}
\varepsilon_{1,i,t} &= \ln v_{1,i,t} - E_{t-1}(\ln v_{1,i,t}) \\
\exp(\varepsilon_{1,i,t}) &= \frac{\exp(\ln v_{1,i,t})}{\exp(E_{t-1}(\ln v_{1,i,t}))} \\
E_{t-1}(\exp(\varepsilon_{1,i,t})) &= E_{t-1}\left(\frac{\exp(\ln v_{1,i,t})}{\exp(E_{t-1}\ln v_{1,i,t})}\right) \\
E_{t-1}(\exp(\varepsilon_{1,i,t})) &= \frac{E_{t-1}(v_{1,i,t})}{\exp(E_{t-1}\ln v_{1,i,t})} \\
E_{t-1}(v_{1,i,t}) &= \exp[E_{t-1}(\ln v_{1,i,t})]E_{t-1}(\exp(\varepsilon_{1,i,t})) \qquad \text{(A.20)}
\end{aligned}$$

Using Equation A.19 again, the expression for $E_{t-1}(\ln v_{1,i,t})$ can be replaced in Equation A.20 by $\ln v_{1,i,t} - \varepsilon_{1,i,t}$. This yields

$$\begin{aligned}
E_{t-1}(v_{1,i,t}) &= \exp[\ln v_{1,i,t} - \varepsilon_{1,i,t}]E_{t-1}(\exp(\varepsilon_{1,i,t})) \\
E_{t-1}(v_{1,i,t}) &= \frac{\exp(\ln v_{1,i,t})}{\exp(\varepsilon_{1,i,t})}E_{t-1}(\exp(\varepsilon_{1,i,t})) \\
v_{1,i,t} &= \frac{E_{t-1}(v_{1,i,t})\exp(\varepsilon_{1,i,t})}{E_{t-1}(\exp(\varepsilon_{1,i,t}))} \qquad \text{(A.21)}
\end{aligned}$$

Recall the first-order condition for assets,

$$v_{1,i,t} = \beta E\omega_{1,i,t}(v_{1,i,t+1}(1+r_{t+1}))$$

Again, assuming that there is no uncertainty about interest rates and the probability of being alive in the following period, the first-order condition for assets gives

$$\begin{aligned}
v_{1,i,t} &= \beta(1+r_{t+1})E_t(v_{1,i,t+1}) \\
E_t(v_{1,i,t+1}) &= \frac{v_{1,i,t}}{\beta(1+r_{t+1})} \\
E_{t-1}(v_{1,i,t}) &= \frac{v_{1,i,t-1}}{\beta(1+r_t)} \qquad \text{(A.22)}
\end{aligned}$$

Substituting this into Equation A.21 yields

$$v_{1,i,t} = \frac{v_{1,i,t-1}}{\beta(1+r_t)}\frac{\exp(\varepsilon_{1,i,t})}{E_{t-1}(\exp(\varepsilon_{1,i,t}))}$$

or

$$\nu_{1,i,t} = \nu_{1,i,t-1} \exp(b_{1,i,t}) \exp(\varepsilon_{1,i,t}) \tag{A.23}$$

where

$$b_{1,i,t} = -\ln(\beta(1+r_t)) - \ln(E_{t-1}(\exp(\varepsilon_{1,i,t})))$$

Taking the logs of both sides and repeatedly substituting for increasingly older values of $\nu_{1,i,t}$ gives the following equation

$$\ln \nu_{1,i,t} = \ln \nu_{i,t^d-1} + \sum_{j=t^d}^{t} b_{1,i,j} + \sum_{j=t^d}^{t} \varepsilon_{1,i,j} \tag{A.24}$$

where the iterative substitution can only be performed up to the time of marital dissolution, t^d. The marginal utility of wealth, ν_{i,t^d-1}, is that which comes out of the dissolved married problem. Substituting for its value in Equation A.24 gives

$$\ln \nu_{1,i,t} = \ln \nu_{i,0} - \sum_{j=1}^{t^d-1} \ln \Omega_j + \sum_{j=1}^{t^d-1} b^*_{i,j} + \sum_{j=1}^{t^d-1} \varepsilon^*_{i,j} + \sum_{j=t^d}^{t} b_{1,i,j} + \sum_{j=t^d}^{t} \varepsilon_{1,i,j}$$

A.4 Labour supply with myopia: derivations

Under myopia, the married problem is

$$\begin{aligned} V_{1,M}(A_{i,t}) = \max_{Q_{i,t}} \, & \lambda_{i,t} U_{1,i,t}(l_{1,i,t}, L_{1,i,t}, c_{1,i,t}, C_{1,i,t}) + \lambda_{i,t} \beta E_t [V_{1,M}(A_{i,t+1})] \\ & + (1-\lambda_{i,t}) U_{2,i,t}(l_{2,i,t}, L_{2,i,t}, c_{2,i,t}, C_{2,i,t}) + (1-\lambda_{i,t}) \beta E_t [V_{2,M}(A_{i,t+1})] \end{aligned}$$

subject to the budget constraint

$$\begin{aligned} \mathrm{w}_{1,i,t} l_{1,i,t} + \mathrm{w}_{2,i,t} l_{2,i,t} + (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) g_{\Lambda,i,t}(L_{1,i,t} + L_{2,i,t}) - (\mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t}) \kappa_{\Lambda,i,t} \\ + c_{1,i,t} + c_{2,i,t} + p_t g_{K,i,t}(C_{1,i,t} + C_{2,i,t}) - p_t \kappa_{K,i,t} + A_{i,t+1} \\ = \mathrm{w}_{1,i,t} + \mathrm{w}_{2,i,t} + Y_{i,t} + (1+r_t) A_{i,t} \end{aligned}$$

The corresponding first-order conditions are

$$l_{1,i,t} : \lambda_{i,t} \gamma_{l,1,i} \frac{1}{l_{1,i,t} - \bar{l}_{1,i,t}} = \nu_{i,t} \mathrm{w}_{1,i,t}$$

$$l_{2,i,t} : (1-\lambda_{i,t})\gamma_{l,2,i}\frac{1}{l_{2,i,t}-\bar{l}_{2,i,t}} = \nu_{i,t}\mathrm{w}_{2,i,t}$$

$$L_{1,i,t} : \lambda_{i,t}\gamma_{L,1,i}\frac{1}{L_{1,i,t}-\bar{L}_{1,i,t}} = \nu_{i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}$$

$$L_{2,i,t} : (1-\lambda_{i,t})\gamma_{L,2,i}\frac{1}{L_{2,i,t}-\bar{L}_{2,i,t}} = \nu_{i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}$$

$$c_{1,i,t} : \lambda_{i,t}\gamma_{c,1,i}\frac{1}{c_{1,i,t}-\bar{c}_{1,i,t}} = \nu_{i,t}$$

$$c_{2,i,t} : (1-\lambda_{i,t})\gamma_{c,2,i}\frac{1}{c_{2,i,t}-\bar{c}_{2,i,t}} = \nu_{i,t}$$

$$C_{1,i,t} : \lambda_{i,t}\gamma_{C,1,i}\frac{1}{C_{1,i,t}-\bar{C}_{1,i,t}} = \nu_{i,t}p_t$$

$$C_{2,i,t} : (1-\lambda_{i,t})\gamma_{C,2,i}\frac{1}{C_{2,i,t}-\bar{C}_{2,i,t}} = \nu_{i,t}p_t$$

$$A_{i,t+1} : \lambda_{i,t}\beta E_t\left(\frac{\partial V^M_{1,i}(A_{i,t+1})}{\partial A_{i,t+1}}\right) + (1-\lambda_{i,t})\beta E_t\left(\frac{\partial V^M_{2,i}(A_{2,i,t+1})}{\partial A_{2,i,t+1}}\right) = \nu_{i,t} \tag{A.25}$$

$$\begin{aligned}\nu_{i,t} : \mathrm{w}_{1,i,t}l_{1,i,t} + \mathrm{w}_{2,i,t}l_{2,i,t} &+ (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})g_{\Lambda,i,t}(L_{1,i,t}+L_{2,i,t}) - (\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})\kappa_{\Lambda,i,t} \\ &+ c_{1,i,t}+c_{2,i,t}+p_t g_{K,i,t}(C_{1,i,t}+C_{2,i,t}) - p_t\kappa_{K,i,t} + A_{i,t+1} \\ &= \mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t}+Y_{i,t}+(1+r_t)A_{i,t}\end{aligned}$$

The terms $\frac{\partial V^M_{1,i}(A_{i,t+1})}{\partial A_{i,t+1}}$ and $\frac{\partial V^M_{2,i}(A_{i,t+1})}{\partial A_{i,t+1}}$ in Equation A.25 are unknown. However, they can be obtained as usual by relying on the Envelope theorem.

The individual problem of the married individuals yields the following

$$\begin{aligned}\frac{\partial V^M_{1,i}(A_i)}{\partial A_i} &= \lambda_i\beta E\left(\frac{\partial V^M_{1,i}(A'_i)}{\partial A'_i}\frac{\partial A'_i}{\partial A_i}\right) + (1-\lambda_i)\beta E\left(\frac{\partial V^M_{2,i}(A'_i)}{\partial A'_i}\frac{\partial A'_i}{\partial A_i}\right) + \nu_i\left((1+r) - \frac{\partial A'_i}{\partial A_i}\right) \\ &= \frac{\partial A'_i}{\partial A_i}\left[\lambda_i\beta E\left(\frac{\partial V^M_{1,i}(A'_i)}{\partial A'_i}\right) + (1-\lambda_i)\beta E\left(\omega'_i\frac{\partial V^M_{2,i}(A'_i)}{\partial A'_i}\right) - \nu_i\right] + \nu_i(1+r)\end{aligned} \tag{A.26}$$

From Equation A.25, the first term on the right-hand side of Equation A.26 is equal to zero, so that:

$$\frac{\partial V_{1,i}^{M}(A_i)}{\partial A_i} = \nu_i(1+r)$$

and therefore:

$$\frac{\partial V_{1,i}^{M}(A_i')}{\partial A_i'} = \nu_i'(1+r')$$

Substituting this back into the first-order condition for A' yields:

$$A' : \beta E\left[\nu_i'(1+r')\right] = \nu_i$$

Denote the labour supply of spouse 1, $N_{1,i,t} = 1 - l_{1,i,t} - L_{1,i,t}$. Then,

$$N_{1,i,t} = \bar{N}_{1,i,t} - \frac{\lambda_{i,t}\gamma_{l,1,i}}{\nu_{i,t}\mathrm{w}_{1,i,t}} - \frac{\lambda_{i,t}\gamma_{L,1,i}}{\nu_{i,t}g_{\Lambda,i,t}(\mathrm{w}_{1,i,t}+\mathrm{w}_{2,i,t})}$$

Appendix B

Appendix to Chapter 2

Table B.1: Child outcomes by parental total family income quintile

Parent family income quintile	Employed	Married or cohabitating
Bottom	80.1	58.6
Q2	86.1	64.0
Q3	85.7	66.0
Q4	91.3	74.0
Top	93.4	72.0
Total	87.3	66.9

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

Table B.2: Possible values of skill intensity index score

Complexity of score	Importance of skill				
	0	1	2	3	4
1	0	1	2	3	4
2	0	2	4	6	8
3	0	3	6	9	12
4	0	4	8	12	16
5	0	5	10	15	20
6	0	6	12	18	24
7	0	7	14	21	28

Table B.3: Decomposition of relative intergenerational mobility

	Total income				Market income				Employment income			
	Child couple		Child individual		Child couple		Child individual		Child couple		Child individual	
	Effect	%	Effect	%	Effect	%	Effect	%	Effect	%	Effect	%
Total effect	0.273	100.00	0.229	100.0	0.249	100.0	0.201	100.0	0.233	100.0	0.193	100.0
Explained	0.170	62.3	0.145	63.4	0.142	57.0	0.119	59.3	0.132	56.5	0.112	57.8
Education	0.111	40.5	0.110	48.1	0.095	38.2	0.090	44.8	0.090	38.6	0.082	42.6
Education (emp. status)	0.013	4.7	0.021	9.3	0.006	2.4	0.010	5.1	0.006	2.7	0.010	5.3
Education (marital status)	0.008	3.0	0.002	0.9	0.002	1.0	0.001	0.3	0.003	1.1	0.001	0.4
Education (net)	0.090	32.8	0.087	37.9	0.087	34.8	0.079	39.4	0.081	34.8	0.071	36.9
Emp. status (uncorr. w/ ed.)	0.014	5.0	0.023	9.9	0.012	5.0	0.021	10.4	0.014	6.1	0.023	11.9
Marital status (uncorr. w/ ed.)	0.046	16.8	0.012	5.4	0.034	13.8	0.008	4.2	0.027	11.7	0.006	3.3
Unexplained	0.103	37.7	0.084	36.6	0.107	43.0	0.082	40.7	0.102	43.5	0.081	42.2
N	6,289,700		6,289,700		5,898,700		5,898,700		5,720,800		5,720,800	

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.

Table B.4: Decomposition of relative intergenerational mobility, employment income (employed respondents only)

	Child couple income		Child individual income		Child and parent individual income	
	Effect	%	Effect	%	Effect	%
Total effect	0.211	100.0	0.168	100.0	0.141	100.0
Explained	0.124	58.9	0.099	58.6	0.094	66.5
Education, total	0.084	39.9	0.077	45.6	0.072	51.4
Education, through job char.	0.041	19.5	0.037	22.1	0.036	25.3
Skill use	0.038	18.3	0.034	20.3	0.032	23.0
Job quality	0.001	0.7	0.003	1.5	0.003	1.9
Education, through marital status	0.001	0.6	0.001	0.3	0.000	0.3
Education, net	0.043	20.3	0.040	23.5	0.037	26.1
Job char., uncorr. w/ ed.	0.021	10.2	0.019	11.4	0.019	13.5
Skill use	0.016	7.7	0.012	7.2	0.011	8.1
Job quality	0.005	2.5	0.007	4.2	0.008	5.4
Marital status, uncorr. w/ ed.	0.019	8.8	0.003	1.6	0.002	1.6
Unexplained	0.087	41.1	0.070	41.4	0.047	33.5
N	5,236,200		5,236,200		5,236,200	

Source: Longitudinal and International Study of Adults, 2014, and T1 Family File, 1982-2013.