Employment Inequality, Labour Income Dynamics and Household Consumption

by

Francis Chiparawasha

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Economics

University of Alberta

Abstract

This thesis has five chapters. Chapter 1 introduces the thesis. Chapter 2 uses a search and matching model that incorporates labour demand, labour supply, labour market efficiency, search effort and duration dependence of job finding rates to examine the leading cause(s) of a rise in employment inequality (divergence of employment rates) between low-educated and high-educated men in Canada. In the model, workers differ in education attainment, ability, and search effort, and firms post vacancies to attract workers. The widening employment gap results from shifts in how workers supply labour, exert job search effort, the skills firms prefer to hire and the ruggedness of the matching process. This model is calibrated to match the rise in employment inequality between 1990 and 2019 and perform counterfactual experiments. The results indicate that changes in labour supply (e.g., shifts in preference for leisure and generosity of government programs) is the primary cause of the rise in inequality, while improvements in job search technology reduced the employment gap.

Chapter 3 provides a guide to estimating the canonical income process widely used in macro and labour economics that includes autoregressive, transitory, and fixed effect components using quasidifferences. Estimation in quasidifferences has its advantages over estimations in levels and differences, yet it is rarely used in practice, and nothing is known about its effectiveness. The chapter provides a catalog of biases in the estimated parameters for various N, T, initial conditions and weighting schemes us-

ing Monte Carlo simulations. The chapter also applies the quasidifferences method to Danish administrative data on male earnings. The method yields divergent estimates of the autoregressive parameter for different weighting schemes, which conforms to the simulated results. In particular, the estimates are divergent when the variance of transitory shocks is higher than that of permanent shocks, true persistence is high, and the variance of the permanent component in the first sample year is nonzero. Lastly, the chapter applies quasidifferences to income and consumption data from a calibrated standard incomplete-markets model. The income and consumption data allow the joint estimation of the income process parameters and consumption insurance against permanent and transitory shocks. The results show that estimation in quasidifferences yields significant biases in the autoregressive parameter and the estimated insurance of shocks. All these findings warn against the routine use of estimation in quasidifferences. However, the income process parameters can be reliably estimated using quasidifferences only when the variance of permanent shocks is higher than that of transitory shocks, and one uses equal weighting of the moments.

Chapter 4 uses the methodology of Blundell et al. (2008) to estimate the degree of household consumption insurance and the primary sources of this insurance in South Africa during the period 2008-2017. Household income shocks can be detrimental to families' well-being, especially in developing countries where the level of consumption for many families is already close to subsistence. In the chapter, income shocks can either be transitory or permanent, and the ability of households to smooth these income shocks depends on available resources. Because families have access to limited resources, they can only partially cushion consumption from the

income shocks. The results indicate that households have almost full insurance against transitory shocks and can insure about 31 percent of the permanent income shocks. The degree of consumption insurance also differs by education, region of residence, race and age. The results also show that income from private transfers is more effective at smoothing consumption than public transfers. Lastly, Chapter 5 provides concluding remarks.

Preface

This dissertation consists of an introductory chapter, three main chapters and concluding remarks. The second chapter of the three main chapters, "A Guide to Estimating the Canonical Income Process in Quasidifferences", is a co-authored work with my supervisor Dr. Dmytro Hryshko.

Acknowledgements

I would like to use this opportunity to sincerely thank all the people who played a part during my PhD research journey. It has been a privilege to come across different people who inspired me to continue with my research. First and foremost, I would like to express my deepest gratitude to my supervisor Dr. Dmytro Hryshko for his continuous support and invaluable feedback. I am also thankful to him for providing me the opportunity to work as a Research Assistant on his SSHRC Insight Grant (*Improving the Measurement of Consumption and Income Dynamics*). I learnt a lot during the research assistantship. Without his continuous support and encouragement, the completion of this dissertation may have been a mammoth task. I am also grateful for the contributions from my dissertation committee members Dr. Malik Shukayev and Dr. Chetan Dave, for their continuous guidance and support.

I am grateful to Dr. Heather Eckert, our Graduate Chair at Economics Department, for constantly creating opportunities to share our work among colleagues as well as to receive constructive feedback during the PhD cohort meetings. Dr. Heather Eckert also organized workshops for us as PhD students to learn how to present our ideas. Her unwavering support during the COVID pandemic was comforting and inspirational. I would like to also thank Mrs. Christina Smith, our Graduate Program Advisor, for her administrative support and prompt response to my queries during my entire program. She was always available to help since I arrived at the department to commence my PhD studies.

I also wish to thank the following professors for teaching me insightful courses: Dr. Haifang Huang, Dr. Vera Brencic, Dr. Dmytro Hryshko, Dr. Tilman Klumpp, Dr. Dana Andersen, Dr. Joseph Marchand, Dr. Sebastian Fossati Pereira, Dr. Xingfei Liu, Dr. Chetan Dave and Dr. Malik Shukayev. In addition, I am grateful to Dr. Martha White (Computer Science), Dr. Yau Wong (Mathematics & Statistical Sciences), and Dr. Yaozhong Hu (Mathematics & Statistical Sciences) for allowing me to audit their Machine Learning, Numerical Methods and Statistical Processes courses, respectively. During the PhD research process, I also received four funding offers: University of Alberta Graduate Completion Scholarship, University of Alberta Graduate Fellowship, University of Alberta Graduate Scholarship and University of Alberta Doctoral Recruitment Scholarship. I am grateful and appreciative for this funding because it helped to make my research less stressful.

I also wish to thank my close friends Agamani Chakrabarty and Kwabena Bediako Boateng for their encouragement and friendship during my research tenure and all the PhD students who provided valuable feedback during our PhD cohort meetings. Finally, I am indebted to my mom Agness Chiparawasha, my siblings, Prosper, Priscilla, Alice and Evelyn, and my wife, Monica Amalia Franco Herrera for their love and support. Any errors or omissions are my sole responsibility.

Table of Contents

List of Tables

List of Figures

Abbreviations

DWMD Diagonally-weighted minimum distance.

EP Employment-to-population ratio.

EWMD Equally-weighted minimum distance.

Fig Figure.

GMM Generalized method of moments.

NAICS North American Industry Classification System.

NIDS National Income Dynamics Survey.

OMD Optimally-weighted minimum distance.

PSID Panel Study of Income Dynamics.

Glossary of Terms

Consumption insurance is the fraction of income shocks that do not pass through to consumption.

Employment inequality refers to the difference in employment rates between high school dropouts and graduates (individuals with at least high school diploma).

Nonemployment refers to all individuals who are out of employment, including unemployed and out of the labour force.

Prime-age includes individuals between the ages of 25 and 54, inclusive.

Chapter 1

Introduction

Many people aspire to acquire a good education, work during their prime, and accumulate wealth to insure consumption against income shocks over the life cycle. This dissertation examines; i) the causes of an increase in employment inequality in Canada, ii) the measurement of canonical income process parameters using quasidifferences and ii) the ability of South African households to smooth consumption from income shocks. The first issue relates to labour market changes that heterogeneously affect men of different educational attainment, leading to a divergence in employment rates. Second topic catalogs biases associated with estimating income process parameters using quasidifferences. Lastly, the third chapter focuses on the ability of households to cushion consumption from income shocks and the sources of the insurance.

The employment rate of prime-age men in Canada has been trending downwards for several decades; see Picot and Heisz (2000) and Galarneau et al. (2013). The evidence also shows that the rate of this decrease is inversely related to the level of education, dropping faster for low-educated men; see Fig. A-1. On average, high-educated men tend to fare better in the labour market. The employment rates of high-educated and low-educated men do not necessarily have to be the same, but it is concerning when the gap widens over time. There is no literature in Canada examining the existence and causes of this widening employment gap. Chapter 2 documents and analyzes the causes of this rise in the employment gap between low-educated and high-educated prime-age men in Canada.

The evidence shows that the employment gap has been rising over time, a phenomenon defined as employment inequality in the literature. The chapter identifies the driving factors of the rise in employment inequality by building a search and matching model augmented with heterogeneous ability among workers, variable job search effort, and education choice. The model contains five potential channels that can drive divergence in employment rates, namely: i) labour demand, ii) labour supply, iii) labour market frictions, iv) search effort and v) duration dependence of job finding rates.

The model is calibrated to match the increase in employment inequality between 1990 and 2019 in the Canadian data and perform counterfactual experiments to disentangle the role of each channel. Identifying the primary driver(s) of the rise in equality is crucial in designing and choosing the appropriate policy response. For example, a decrease in the demand for low-educated men by firms may be addressed by cutting payroll taxes or providing job subsidies.[1] A rise in employment inequality driven by trade competition or automation may require worker retraining programs. When the rise is driven by increasing generosity of the safety net, it may require reviewing and adjusting public programs to minimize their role in discouraging job search. Government programs such as unemployment insurance, social assistance and disability insurance may alter the incentive to work by providing liquidity relief to households during unemployment spells (Chetty, 2008). However, some shifts may be difficult to address. If the fall reflects a genuine change in preference for leisure, as documented in Aguiar et al. (2021), bringing workers back into employment will be challenging. The results from the chapter suggest that shifts in labour supply have been the main driver of rising employment inequality. This conclusion aligns with the evidence in Section 2.3 that shows an increase in leisure time among low-educated men, while it trended downwards for the high-educated men. Public programs (the safety net) are also becoming increas-

[1] See Nickell and Bell (1995) for more information on this policy response.

ingly generous to low-educated prime-age men.

Chapter 3 analyzes the measurement of income process parameters. The common approach in macro and labour economics involves modelling income as a combination of permanent and transitory components and fixed effects. Understanding how income evolves during an individual's working life involves estimating the parameters related to each component. In most instances, the estimation of these parameters relies on minimum-distance estimation and autocovariance moments of log-income in levels y_{it}, or differences $\Delta y_{it} = y_{it} - y_{i,t-1}$. However, there is an overlooked method of quasidifferences $\tilde{\Delta y_{it}} = y_{i,t} - \rho y_{it-1}$, where Δy_{it} is the growth in income y of individual i between period $t-1$ and t, and ρ is the persistence of the permanent component, as proposed by Blundell et al. (2015).

Compared to estimation in levels and differences, estimation using quasidifferences is simple and encompasses the advantages of these two methods. The method provides a way of estimating both the fixed effect and persistent parameters without the complications required to estimate the same parameters using level moments. The only challenge of using quasidifferences to estimate income parameters is that it is infeasible to jointly estimate the persistence parameter ρ together with the other parameters. Blundell et al. (2015) propose joint estimation by minimizing the distance between model and data moments using a pre-set persistence grid. This method is appealing due to its simplicity. However, not much is known regarding its accuracy when applied to data. Therefore, Chapter 3 provides a detailed analysis of the biases associated with estimating income parameters using this method.

The chapter uses Monte Carlo simulations to document the biases for different sample sizes N, length of observed income series T, initial conditions and weighting schemes.[2] Understanding the accuracy of this method is crucial because parameters that define income dynamics are often inputs in other applications. Examples of these applications include the design of optimal tax policy (Farhi and Werning, 2012), the relationship between income shocks and consumption (Blundell, 2014; Hryshko and Manovskii, 2022; Bryukhanov and Hryshko, 2020), the role of the welfare system and family labour supply as sources of smoothing consumption from income volatility (Blundell et al., 2015) and models on wealth accumulation (Carroll, 1997). The main results from the chapter show that equally-weighted estimation recovers the true persistence of permanent shocks ρ for different T, N, and initial conditions only when the variance of permanent shocks is higher than that of transitory shocks.

The chapter also applies the quasidifferences method to Danish administrative data on earnings as well as income and consumption data from a simulated incomplete-markets model. Using income and consumption data entails estimating additional so-called transmission parameters that measure the fraction of permanent and transitory shocks passing through to consumption. The results from these experiments indicate biases in the estimated persistence and the transmission parameters for permanent and transitory shocks. Biases exist for high and low values of persistence, different weighting schemes and different number of households.

Chapter 4 answers the question: to what extent can South African households shield consumption from permanent and transitory income shocks? The income process is similar to the one discussed in Chapter 3, except that the permanent component is a random walk and fixed effects are set to zero. Imposing the random walk assumption makes estimation

[2]The chapter considers optimal, diagonal and equal weighting schemes.

of parameters that describe the comovement of income and consumption using differences Δy_{it} instead of quasidifferences $\tilde{\Delta y_{it}}$ preferable. Thus, the chapter estimates the transmission parameters using log-income differences.[3] These transmission parameters are directly related to the degree of consumption insurance. The analysis also documents the level of insurance across education, region of residence, race and age.

Besides estimating the degree of insurance, the chapter investigates the main insurance devices households utilize in the event of an income shock. In this framework, households encounter permanent and transitory shocks, and their ability to shield consumption from income shocks depends on their access to insurance devices. Insurance devices include personal savings, borrowing, family labour supply, safety net, taxation, private transfers and other sources that can stabilize consumption in the event of a shock. The chapter specifically focuses on the role of private and public transfers. Estimates of the magnitude and sources of consumption insurance can inform the design of programs such as tax rebates, food relief, agriculture input schemes and welfare transfers. The results indicate that families have full and partial insurance against transitory and permanent income shocks, respectively. Private transfers are more effective for consumption insurance than public transfers. The results also show that households with older and more-educated or white heads have more consumption insurance than the average family.

[3]The estimation is similar to the application of quasidifferences to standard incomplete-markets model data above, except that it uses a differences method and real data from South Africa.

Chapter 2

Employment Inequality Among Prime-Age Men in Canada

2.1 Introduction

The employment-to-population ratio of prime-age men has been trending downwards for several decades. This decrease is more pronounced for low-educated men, leading to a rise in employment inequality over time. Fig. 2.1 depicts the employment-to-population ratios of high school dropouts and graduates in Canada.[1] While the ratio for low-educated men dropped by more than 25 percent between 1976 and 2020, the decline for graduates is less than 10 percent. These diverging trends document a rise in employment inequality—a widening gap in employment-to-population ratios between dropouts and graduates. Although the employment rate of low-educated men is usually lower than that of high-educated men, it is concerning when the gap continues to diverge over time.

I analyze the potential drivers of this rise in employment inequality. The employment adjustments in Fig. 2.1 are likely to be driven by changes in the economy, such as the evolving effect of automation, trade, preference for leisure and job search technology. The effect of these changes may have an asymmetric impact across education groups. This chapter aims to answer the question: why is the employment rate of low-educated men falling faster than that of high-educated men?

Any effort to slow down or reverse this rise in inequality requires understanding the potential cause(s) in order to design an optimal policy response. For instance, a decrease in employment due to shifts in preference over leisure may be less concerning than a fall due to factors that hinder men from supplying labour even when they prefer to work. When

[1] Graduates include men with at least a high school diploma. This grouping is informed by the degree of a fall in employment rates as shown in Fig. A-1. The terms "low-educated" and "high school dropout" are used interchangeably throughout the chapter and similarly for "high-educated" and "graduate".

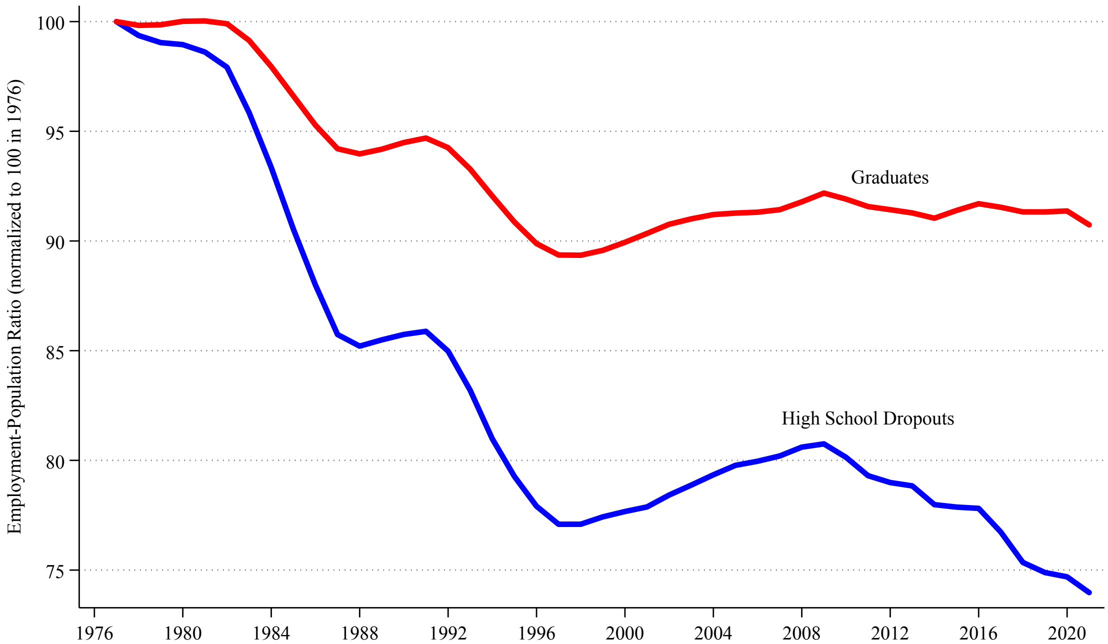

Figure 2.1: Employment-Population Rate of Prime-Age Men in Canada, Age 25-54

Notes: The data is from annual averages of monthly Canadian Labour Force Surveys (LFS). The trends are demographically adjusted for age and province of residence. High school dropouts include all men without high school diploma. Graduates include men with at least a high school diploma. The data is weighted by survey weights. The base employment rates in 1976 are 81% and 92% for high school dropouts and graduates, respectively.

the decline is due to factors beyond their control, it becomes vital to understand the causes. Better knowledge can inform the design of policy response that minimizes potential future social problems, such as populism and poor mental health, that are likely to emerge when some groups in society have perceived feelings of alienation.[2]

The downward trend in employment among prime-age men is also documented in the US (see Abraham and Kearney, 2020, for a survey) and earlier studies in Canada (Picot and Heisz, 2000; Galarneau et al., 2013). There is no consensus in the literature on why the decrease is significantly more for some groups and the leading causes of the divergence. Stud-

[2]See Nickell and Bell (1995), Farré et al. (2018) and Guriev (2018).

ies on this decline and the differences across education are scant in the Canadian literature. However, studies on the potential causes of the general downward trend in employment rates of prime-age men are vast.[3] The main causes of the fall in employment rates are broadly classified into three channels: i) demand-side factors, ii) supply-side factors and iii) labour market search frictions.

The advent of the internet is transforming how people search for work, and its benefits accrue to individuals who can optimally utilize the new tools. Although the three channels above are essential in explaining the changes in employment and employment inequality, most studies tend to overlook the direct effect of shifts in job search behaviour and the role of nonemployment duration. Better job search methods can help shorten nonemployment spells by increasing the probability of finding work, thus reducing the proportion of individuals in long-term nonemployment. Since there is evidence that firms discriminate against workers with long nonemployment spells (see Farber et al., 2016; Van Belle et al., 2018; Farber et al., 2019; Lyshol et al., 2021), a shorter nonemployment duration might reduce the discrimination effect of firms. When workers face less discrimination, they may increase their job search effort, which increases their chances of finding work. Employment inequality may rise over time due to differences in job search behaviour/intensity and nonemployment duration across education. Therefore, I also consider the role of job search effort and duration dependence of job-finding rates.

To ascertain which of the five channels (demand, supply, search fric-

[3]See Autor et al. (2008), Aguiar and Hurst (2008), Binder and Bound (2019), Acemoglu and Restrepo (2020), Grigoli et al. (2020), Abraham and Kearney (2020), Georgieff and Milanez (2021), Aguiar et al. (2021) and Wolcott (2021). Abraham and Kearney (2020) provide a review of the literature on the determinants of the decline in employment-population ratio and Perez-Arce and Prados (2021) review the literature on falling labour force participation.

tions, search effort and nonemployment duration dependence) dominantly explain the rise in inequality, I build a theoretical search and matching model with worker heterogeneity in ability, job search and education choice. The model also includes an exogenous job separation channel. I calibrate the model to match Canadian data moments. Understanding which of these channels is driving the rise in the employment gap can be crucial for policymakers in designing policies. For instance, if the divergence is due to labour demand factors, instituting reskilling training programs for the less educated may help to increase their employability. Alternatively, if the decline is due to government insurance programs (e.g., unemployment insurance, disability insurance, social assistance) becoming more generous over time, policymakers may need to restructure some programs. If it is due to changes in job search technology that disproportionately help high-educated men, the government can push for programs that reduce the cost of acquiring quality job-related information. Some factors are, however, uncertain on how the government should respond. For example, if the steeper decline in low-educated employment rate reflects changes in tastes between work and leisure due to the evolution of social norms or leisure technology (Eberstadt, 2016; Aguiar et al., 2021). I quantify the role of these endogenous channels to help understand what policies are necessary to address the rise in employment inequality.

This chapter makes three contributions. The first contribution is to document significant changes in labour-matching efficiency, job search and nonemployment duration between 1990 and 2019. The data shows that shifts in job search effort have been different across education and nonemployment duration. For instance, job search intensity barely changed for short-term (up to 24 weeks) unemployed dropouts but rose significantly for graduates. I also document changes in time spent on leisure between 1986 and 2016. The data show that graduates have significantly reduced time spent on leisure, while that of high school dropouts remains essentially constant (positive increase but insignificant).

The second contribution is that I use these facts to build a theoretical search and matching model in a labour market with search frictions as in Diamond (1982), Mortensen (1982) and Pissarides (1985). There are workers nonemployed in each period even though the economy has unfilled vacancies because of the search frictions. The model helps to match the evidence of a rising share of workers in nonemployment. To replicate the rise in employment inequality, I extend the model to include heterogeneity in ability, variable job search effort and educational choice. The model builds on Wolcott (2021) by adding variable search effort among nonemployed workers and nonemployment duration dependence of job-finding rates.[4] Workers choose whether to at least graduate from high school or drop out of high school.[5] Adding these components enriches the model in that workers can influence their chances of finding work by exerting more search effort. Because the job search process is time intensive, too much search effort is costly due to forgone leisure time. This trade-off ensures the existence of an optimal search effort that directly affects employment by altering job-finding rates and indirectly by changing the overall efficiency of the labour market. For example, advancement in job search technology (in-

[4]The model adds the role of nonemployment duration following the convention in the literature. Kroft et al. (2016), Jarosch and Pilossoph (2018) and Kroft et al. (2019) provide examples of how to directly control for nonemployment duration dependence of job-finding rates in the data and structural models.

[5]In Wolcott's model, workers have a binary choice of graduating high school or attending college; thus, she lumps dropouts with high school graduates. The trends in Fig. A-1 for Canada show that high school dropouts are different from high school graduates—the decrease in employment of high school graduates is closer to the college cohort than that of dropouts.

ternet, job search boards, etc.) makes it easier to increase the intensity of search but also improves the overall efficiency of the labour market. As the results show, shifts in search effort had a significant effect on employment inequality, even bigger than that of labour demand.

The final model has six channels; five endogenous and one exogenous, that can potentially explain the rise in employment inequality. The endogenous channels include three broad (labour demand, labour supply and labour market efficiency) and two specific (search effort and nonemployment duration dependence) and job separation as an exogenous channel. The third contribution involves calibrating the model to match labour market characteristics in Canada (wages, labour flows, tightness and share of graduates). With the calibrated model, I conduct counterfactual experiments to determine which of these five channels significantly explains the rise in employment inequality. The results indicate that the labour supply channel is the main contributor to the widening employment gap, while improvements in search frictions helped to moderate the rise. The role of labour supply is consistent with low-educated workers benefiting the most from a more generous safety net and preferring to spend more time on leisure. The inequality-reducing effect of matching efficiency suggests that advancements in job search technology were more beneficial to high school dropouts. Wolcott (2021) also reports that improvements in labour market efficiency contributed to reducing employment inequality in the US. However, she finds the labour demand channel as the primary driver of the rise in employment inequality. This difference in results implies that while trade competition and automation played a more prominent role in the US, shifts in labour supply were more prevalent in Canada. Over the last few decades, the loss of manufacturing and offshorable jobs has been more salient in the US than in Canada.

The rest of the chapter is structured as follows. Section 2.2 reviews the literature. Section 2.3 discusses data and documents the evidence illustrating different shifts in the labour market. Section 2.4 presents the details of the model and the static equilibrium. Section 2.5 calibrates the model to match Canadian data moments. Section 2.6 discusses the calibration and counterfactual results. Then Section 2.7 provides some robustness checks, and Section 2.8 concludes.

2.2 Related Literature

This section discusses the literature on how the five channels above can lead to declining employment rates and rising inequality. On the labour demand side, the proliferation of automation that disproportionately destroys jobs in which low-educated men are likely to find employment, e.g., manufacturing, will cause a faster fall in their employment rates. Acemoglu and Restrepo (2020) document evidence that automation displaces less-educated jobs while creating new tasks that are more likely to favour more-educated workers.[6] This effect means employment rates of high-educated workers fall relatively slower than that of low-educated workers as they are more likely to benefit from the new cognitive tasks.

Dixon et al. (2021) use firm data to analyze the impact of automation in Canada. They argue that robots led to polarization in employment along the skills dimension. While the employment of middle-skilled workers fell, the employment of low-skilled and high-skilled workers rose; thus, robots had both substitution and complementary effects on labour employment, depending on the level of skills. In this case, robots displace middle-skilled

[6] Autor et al. (2003) group tasks into manual, routine, and abstract (cognitive). The repetitive manual and routine tasks are more vulnerable to automation.

workers while raising the need for low and high-skilled workers. The trends in Fig. A-1 contradict these results. If it were true, the employment rate of high school dropouts (low-skilled workers) would not have fallen by a big margin.

Another source of changes in labour demand is import competition.[7] Autor and Dorn (2013) analyze the effect of import competition on US local labour markets and document a significant adverse effect on labour force participation, wages in non-manufacturing sectors, household earnings, and an uptick in public transfer payments. Malgouyres (2017) and Dauth et al. (2021) report similar significant spillover effects from Chinese import competition in France and Germany, respectively. Acemoglu et al. (2016) documents similar spillover effects in the US. Blinder and Krueger (2013) explore the effect of offshoring information tasks abroad. They estimate that about 25% of US jobs are vulnerable to offshoring.[8] The share of men employed in tradable sectors in Canada began to trend downwards around the early 2000s; see Fig. A-2, with the decrease steeper for low-educated men. The fall in tradable sector employment is consistent with trade competition displacing local jobs.

The second channel relates to labour supply — shifts in work preferences or institutional policies. Kuhn and Robb (1998) credit the fall in employment among low-skilled men in Canada and the US to a decline in real wages.[9] Moffitt (2012) finds that changes in institutional policies

[7]See studies on import competition from China in developed and developing countries such as Álvarez and Claro (2009), Autor et al. (2013), Bloom et al. (2015), Acemoglu et al. (2016) and Malgouyres (2017).

[8]They, however, assert that the proportion of high-educated workers in these offshorable jobs is high, and the sectors most vulnerable to offshoring are professional and technical services, manufacturing, finance, insurance, and information services.

[9]The authors find that wages of low-skilled workers in the US and Canada have fallen since the 1980s, but more so for the latter, where the social safety net is relatively generous. Fig. A-3 show a similar fall in high school dropouts' wages from the 1980s until the early 2000s, after which low-educated wages in Canada began to rise.

such as taxes, public transfers, and minimum wages cannot significantly explain the fall in the US employment, but Binder and Bound (2019) report that increase in access to disability insurance partially explains the decrease. Eberstadt (2016) and Eberstadt (2022) discuss the role of the safety net. Most of these safety net programs (e.g., disability insurance, unemployment insurance, social assistance) tend to benefit less-educated workers more than high-educated workers. Aguiar et al. (2021) use time-use data and report that the fall in employment of young men (ages 21-30) in the US since 2004 is mainly due to advancements in gaming technology. Better technology is making men enjoy more time on leisure. Similarly, Ye (2021) shows that less-educated men spend more time on leisure than high-educated men.

Krueger (2017) cites health challenges as the cause of declining labour force participation (in the US) among prime-age men. The author documents a rise in the proportion of men on pain medication and labour force participation being lower in states with higher opioids prescription. Similarly, Cheung et al. (2020) estimate the cost of productivity loss due to opioid abuse in Canada to be above $5 billion. Drug abuse (legal and illegal) makes it harder to maintain productivity and consistent employment. Park and Powell (2021) document a positive correlation between the rise in disability insurance claims and falling labour supply due to illicit opioids. The government can also influence labour supply via the criminal justice system. Larson et al. (2021) find a causal relationship between the rise in incarceration rates in the US and falling employment-to-population ratios. They report that a 10% increase in individuals with felony records (between ages 18 and 54) led to a rise in nonemployment by 3%. Individuals with a criminal record may find it harder to rejoin the labour market.

The third channel relates to labour market efficiency, also known as search frictions in the literature. For firms, frictions (search or information asymmetry frictions) complicate the process of identifying suitable workers (Lester et al., 2018). Changes in this efficiency can significantly affect employment; for example, Sedlacek (2016) finds that half of the drop in unemployment during the 2008 Great Recession was due to shifts in matching efficiency. Abraham and Kearney (2020) assert that rising real minimum wages can introduce inefficiencies in the labour market, contributing to falling employment rates. Because minimum wages predominantly affect jobs occupied by low-educated individuals, these frictions then affect their employment relatively more.

As a departure from most of the above studies, Wolcott (2021) quantitatively measures the role of labour demand, labour supply, matching efficiency and job separations in driving employment inequality in the US. The results show that changes in labour demand and job separation rates contributed the most to the rise in inequality between high school and college graduates. I follow the model in Wolcott (2021) but add two more channels: job search effort and nonemployment duration dependence of job-finding rates. By adjusting their search effort, workers can influence their probability of finding work and a higher job-finding rate entails a quicker exit from the nonemployment pool. Although the literature on labour market search with endogenous search effort is vast, the focus is mostly on understanding shifts along the business cycle.[10] I argue that changes in search technology over time can be the source of employment inequality if they disproportionately benefit one group.

Mussida and Zanin (2020) find that high-educated workers are more

[10]See, for example, Christensen et al. (2005), Lise (2013), Gomme and Lkhagvasuren (2015) and Çenesiz and Guimarães (2022).

likely to use a multi-method job search approach. The top panels of Fig. 2.2 depicts similar trends (number of search methods) by nonemployment duration and education in Canada. Mukoyama et al. (2018) show that changes in job search help to dampen the potential rise in labour market inefficiencies. In other words, workers change their search behaviour in response to any inefficiencies, which helps to moderate the impact of some shocks.[11] Faberman and Kudlyak (2019) show that job search intensity differs with unemployment duration, and workers search less in areas with tighter labour markets. Not only does search effort differ by duration, but firms can also discriminate against workers with long unemployment duration (Jarosch and Pilossoph, 2018). For instance, Pizzinelli and Speigner (2017) find that variations in the share of long-term unemployment explain a significant share of the variations in job search intensity in the UK.

Given that job-finding rates tend to fall with nonemployment duration (see Kroft et al., 2019; Jarosch and Pilossoph, 2018), employment inequality may rise if the composition of long-term nonemployment changes over time. I also differ from Wolcott (2021) in that I directly control for shifts in the nonemployment duration dependence of job-finding rates. Lastly, I focus on high school dropouts because lumping them with high school graduates (as in Wolcott, 2021) masks a steeper rise in employment inequality.

2.3 Empirical Evidence

This section documents changes in leisure time, labour market efficiency and job search behaviour between 1990 and 2019.

[11] In the search and matching with endogenous search effort literature, there is an ongoing debate on whether job search effort is pro- or counter-cyclical. This chapter abstracts from this debate because it focuses on a steady-state to steady-state comparison.

2.3.1 Data

The study relies on data from Labour Force Surveys (LFS) and General Household Surveys (GHS). The LFS is a nationally-representative monthly household survey conducted by Statistics Canada to gather information on individuals in employment, unemployment, and out of the labour force. The survey contains information on the characteristics of the working-age population, including job search behaviour, nonemployment duration, age, sex, marital status, educational attainment, and occupation characteristics. The survey has been running since 1945, with significant revisions in 1976 and 1997. Because the focus is on prime-age men, I only use the data for non-institutionalized men between the ages of 25 and 54. I use the LFS to compute employment ratios, wages, job-finding rates, job separation rates, search behaviour, and nonemployment duration dependence.

The LFS collected information on wages from 1997. Therefore, I complement wage data with values from three surveys: Survey of Work History, 1981; Survey of Union Membership, 1984; Labour Market Activity Survey, 1986-1990. These surveys record information on hourly wages. I then convert the nominal wages into real values using the aggregate consumer price index. Following Shimer (2005) and Shimer (2012), the job-finding rates f_t and separation rates δ_t are defined as

$$f_t = 1 - (n_{t+1} - n^s_{t+1})/n_t,$$

$$\delta_t = n^s_{t+1}/(1 - 0.5 f_t) e_t,$$

where n^s_t is the number of newly-nonemployed (less than 4 weeks), n_t is the total number of nonemployed workers, and e_t is the total number of employed individuals.

I measure job search effort s_t as the number of search methods unemployed workers use when looking for work. LFS asks interviewees about the methods they used to search for work during the month, and they can select up to six different search methods.[12] I then define *search effort* as the number of search methods an individual use in a given month.[13] The use of search methods rather than time spent searching for a job is common in the literature (see Bachmann and Baumgarten, 2013; Mussida and Zanin, 2020). I define long-term nonemployment as continuous out of employment for more than 24 weeks.

To measure the role of nonemployment duration dependence on job-finding rates, I utilize an approach common approach in the literature (Kroft et al., 2016; Jarosch and Pilossoph, 2018; Kroft et al., 2019; Kospentaris, 2021).[14] Let τ be nonemployment duration in months. Using the formula for f_t above, I group the nonemployed workers by nonemployment duration $\tau \in \{1, ..., 24\}$, and then calculate the average job-finding rate at each duration such that $f(\tau)$ is the job-finding rate for workers with τ months in nonemployment and $f(0)$ is the rate for the newly-nonemployed (less than a month).[15] I then normalize the average job-finding rates at each duration by the job-finding rate of newly-nonemployed workers such that $\Lambda(\tau) = \frac{f(\tau)}{f(0)}$, with the normalized rates decreasing over duration; see the bottom panels of Fig. 2.2. This normalization means that $\Lambda(\tau) \in [0, 1] \quad \forall \quad \tau \in \{1, ..., 24\}$.

I use time-use data to estimate changes in leisure. The time-use data is from the 1986 and 2015 General Social Surveys. The survey has a

[12]The LFS asks individuals to list the search methods they used to find work up to six different methods: i) employment agencies, ii) contacting employers directly, iii) networking, iv) looking at job ads, v) placed/answered job ads and vi) other methods. See Fig. A-4 that depicts the proportion of men using each method by education.

[13]In the model below, I normalize average search effort in 1990 and 2019 to be between zero and one such that $s_t \in \{0, 1\}$.

[14]It is important to note that these papers focus on unemployment duration while this chapter focuses on nonemployment duration.

[15]I calculate the job-finding rates for the different duration by adjusting the definition of newly-nonemployed workers n_t^s in the formula.

time-use component that tracks changes in how Canadians allocate time to different activities and is conducted once every ten years. Interviewees recall how they allocate time to various activities during the previous day. Following Aguiar and Hurst (2007), I derive four different measures of leisure. Leisure Measure 1 refers to time spent on: socializing, personal care, pet care and gardening. Leisure Measure 2 refers to time spent on Leisure Measure 1 plus time spent on sleeping, eating and personal activities (excluding own medical care). Leisure Measure 3 includes Leisure Measure 2 plus time spent on child care. Leisure Measure 4 is any time not allocated to market or nonmarket work.

The model below requires information on job posting by firms. I follow Landais et al. (2018) and Kroft et al. (2019) to derive a “recruiter-producer ratio", which is a proxy for vacancies.[16] The proxy measures the recruiters-to-employed ratio of workers. However, the challenge is identifying individuals working in the recruitment industry from the available labour force survey data. Landais et al. (2018) use a 5-digit North American Industry Classification System (NAICS) in the United States and Kroft et al. (2019) utilize a 4-digit code for Canada. Kroft et al. (2019) show that the ratio from the 4-digit aggregation, which is a broader classification, performs relatively well to match the proxy from the 5-digit code. The recruiter-producer ratio is

$$v = \frac{\rho \times \text{rec}}{l - \rho \times \text{rec}},$$

where rec is the number of workers in the recruiting industry, ρ is a scaling

[16]Until recently, Canada did not collect detailed information on vacancies. The Job Vacancies Statistics (JVS) is a monthly vacancy series that commenced in 2011 of data collected at the establishment level. The Job Vacancy and Wage Survey (JVWS) is a quarterly series that started in 2015. The vacancy information is collected at the business location level. There is also a Help-wanted Index, which runs from 1981 to 2003. All these data series do not allow the calculation of vacancies by education without imposing some strong assumptions.

factor, and l is the number of workers employed in all private industries. The scaling factor is for adjusting the fact that only a small proportion of workers are in the recruiting industry and is set to $\rho = 2.6$ as in Kroft et al. (2019). This ratio is a good proxy for vacancies because there is a positive correlation between firms posting vacancies and resources allocated to recruitment (Kroft et al., 2019). I derive the vacancy proxy by education.

2.3.2 Shifts in the Labour Market

2.3.2.1 Leisure Time

As mentioned in Section 2.2, a rise in employment inequality may be due to differences in preference for leisure over time. If one group prefers to spend more time on leisure than in the past, its employment rate will fall by a more considerable margin. Table 2.1 checks whether there are significant changes in time spent on leisure by tabulating changes between 1986 and 2015. The table presents measures of leisure at four different aggregation levels as in Aguiar and Hurst (2007) and Aguiar and Hurst (2008). Leisure measure 1 is the narrowest, while Leisure measure 4 is the broadest. Columns 1-3 tabulate changes for high school dropouts, and columns 4-6 for the graduates.

All four measures indicate no significant increase in time spent on leisure for high school dropouts between 1986 and 2015. In comparison, time spent on leisure fell by 3.6 hours per week (measure 2) for the graduates. A test of whether the change is significantly different across the two groups (penultimate row in the table) rejects the null hypothesis that the shifts are the same. The fifth row shows changes in time spent on home production. While dropouts spent more time on home production in 1986 than graduates, the situation was reversed by 2015. The graduates

Table 2.1: Leisure Activities for Prime-Age Men (Hours per Week)

Education Group:	High School Dropouts (D)			Graduates (G)		
	1986	2015	Change (Δ_D)	1986	2015	Change (Δ_G)
Activity:	(1)	(2)	(3)	(4)	(5)	(6)
Leisure measure 1	37.1	39.3	2.2	30.3	28.4	-1.8
			(2.13)			(0.83)
Leisure measure 2	110.4	112.8	2.4	101.8	98.2	-3.6
			(2.55)			(1.05)
Leisure measure 3	112.4	115.2	2.8	104.5	102.2	-2.4
			(2.58)			(1.07)
Leisure measure 4	116.0	116.8	0.8	110.5	104.4	-6.0
			(0.79)			(1.07)
Home production	8.7	11.4	2.7	6.9	12.6	5.6
			(1.24)			(0.46)
Observations			704			1209
χ^2 test statistic on $\Delta_D - \Delta_G$ (Leisure 2)						7.67 (0.0056)
χ^2 test statistic on $\Delta_D - \Delta_G$ (home prod.)						0.5 (0.4809)

The calculation controls for marital status, age, and household size. Leisure Measure 1 refers to the time individuals spent socializing, in passive leisure, in active leisure, volunteering, in pet care, and gardening. Leisure Measure 2 refers to the time individuals spent in Leisure Measure 1 plus time spent sleeping, eating, and in personal activities (excluding own medical care). Leisure Measure 3 includes Leisure Measure 2 plus time spent in child care. Leisure Measure 4 is any time not allocated to market or nonmarket work. Home production measures time spent preparing meals, housework, home maintenance, gardening, and other household activities. Standard errors are in parentheses.

decreased time spent on leisure while increasing time spent on home production by 5.6 hours per week. These significant shifts in how men allocate time to various activities might be partly the reason behind the rise in the employment gap in Fig. 2.1.

Figs. A-5-A-6 in the Appendix also check whether there have been significant changes in access to the safety net as discussed in Eberstadt (2022). Fig. A-5 plots the share of men living in households receiving at least one government benefit regardless of their labour market status.[17] The plot shows a significant rise in dropouts receiving benefits. Similarly, Fig. A-6 focuses only on employment insurance. Nonemployed high school dropouts receiving employment insurance rose by 12% between 1990 and

[17] Some of the benefits include social assistance, employment insurance, child tax credit, and provincial transfers.

2014, compared to only 5% for graduates. These two figures indicate that the safety net has become increasingly more generous to high school dropouts.

The changes in leisure and generosity of the safety net above directly affect labour supply decisions. For instance, a more generous safety net favouring high school dropouts might lower their urgency to find new work. Such changes in labour supply decisions can lead to a rise in employment inequality. As Section 2.6.2 indicates, shifts in labour supply is the most significant contributor to the rising employment inequality between 1990 and 2019. Section 2.6.2.3 explores the potential role of government insurance programs in changing workers' labour supply over time.

2.3.2.2 Job Search and Nonemployment Duration Dependence

Workers can influence the probability of finding work by changing how they search for work. More so, advancements in search technology, e.g., the internet—job boards, introduced new and broader ways of searching for work. The new tools make searching for work effective, thereby improving the overall efficiency of the labour market in matching searching workers to hiring firms. Employment inequality may rise/fall over time if the benefit differs across education. To check whether there have been significant changes in search effort between 1990 and 2019, I use the number of search methods. This measure corresponds to the definition of search effort in the model below. Because search effort tends to differ across nonemployment duration, Fig. 2.2 (top panels) shows shifts in search effort by duration. Fig. A-7 provides an alternative measure of search effort that directly controls for nonemployment duration. The figure depicts a ratio of the number of search methods over the nonemployment duration.

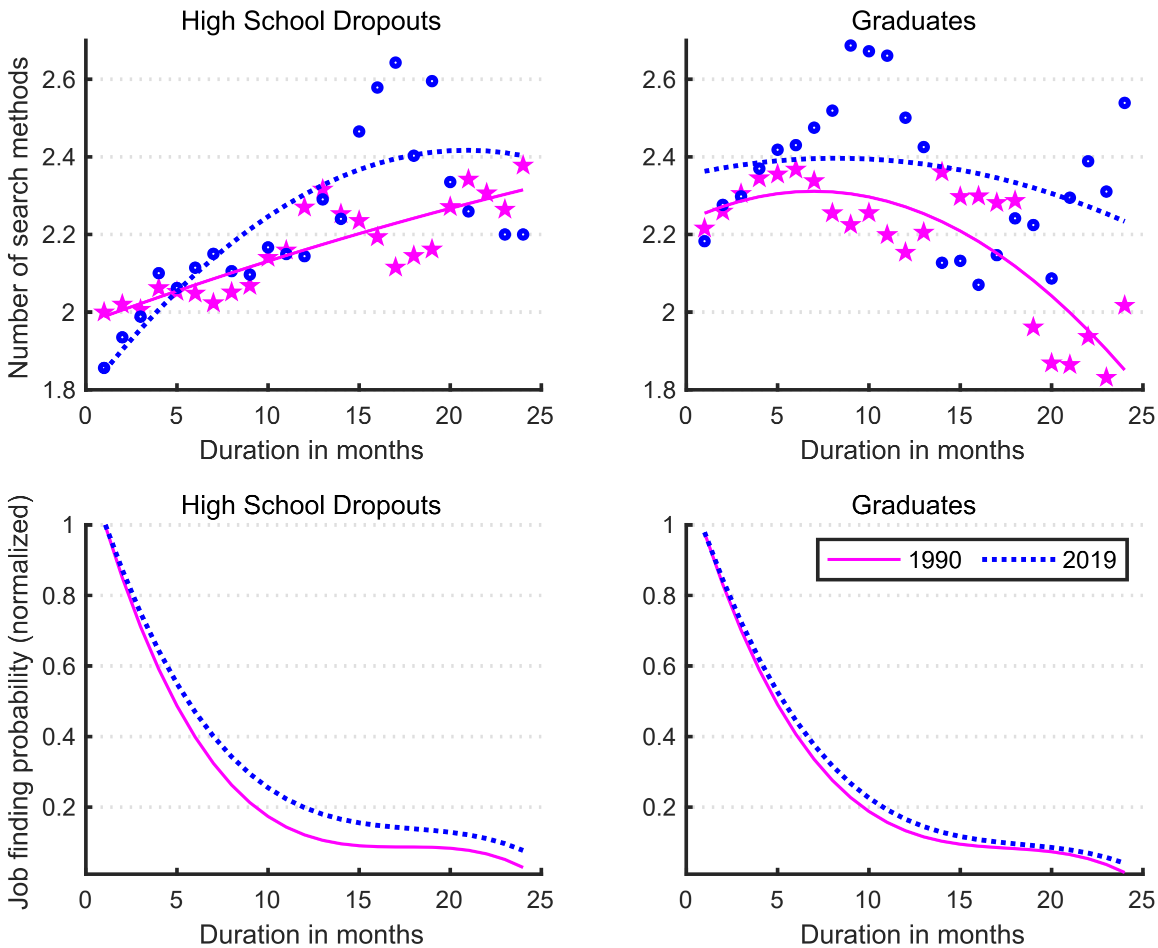

Figure 2.2: Job Search Effort and Job-finding Rates by Duration, Age 25-54

Notes: The data is annual averages of the monthly Canadian Labour Force Surveys (LFS). The plots are demographically adjusted for age. Graduates include high school graduates and dropouts include individuals with no formal schooling, or elementary education. The bottom panels depict average job-finding rate normalized by the rate for individuals with less than one month in nonemployment such that $\Lambda(\tau) = \frac{f(\tau)}{f(0)}$, where τ is duration, $f(0)$ is the job-finding rate for individuals in nonemployment for 1-4 weeks.

While search effort rises almost linearly with duration for high school dropouts, the effort for graduates is hump-shaped, rising initially and then falling when the duration is long. Also, the search effort of graduates within a year of nonemployment is higher than that of high school dropouts before dropping below when nonemployment duration is long enough. Comparing 1990 and 2019, the degree of search effort around early nonemployment remains relatively the same for the graduates but slightly decreased for dropouts. The search effort is almost the same for dropouts until 15 months of nonemployment, after which there was a rise.

The top panels in Fig. 2.2 indicate that there have been some changes in search effort between the two periods. However, the rise in search effort over time might not be as apparent due to the duration dimension. Fig. A-7 in the Appendix depicts an evident rise in job search, especially starting in the late 1990s. This increase in search effort coincides with the rise of the internet. What happens to employment inequality, however, depends on the relative changes. An increase in the job-finding rate for long-term nonemployed workers increases employment. There is also an indirect effect in that employment rises more if the number of individuals in long-term nonemployment decreases because the effect of firms that discriminate against long nonemployment spells becomes less salient (more below). The bottom panels plot the normalized job finding probabilities $\Lambda(\tau)$ by duration, education group, and over time (1990 and 2019).

Nonemployment duration can affect employment inequality in two ways: i) by changing the probability of receiving an offer or ii) by altering the search effort composition of nonemployed workers. If the composition of nonemployment duration changes over time, and the changes are different across education, this might cause a rise in employment inequality. The literature shows that firms discriminate against workers with long-term nonemployment duration, so job-finding rates decrease with duration. Consistent with the job search plots, the job-finding rate of individuals with longer nonemployment duration has risen (relative to the newly unemployed) over time. Secondly, the increase is significantly higher for high school dropouts. For instance, the probability of finding work, relative to the newly nonemployed high school dropout with ten months in nonemployment, has increased by about 7 percentage points, while it is only 2 percentage points for graduates. This significant rise in high school dropouts duration adjustment relative to graduates is likely to reduce employment inequality.

2.3.2.3 Labour Market Efficiency (Search Frictions)

One of the most significant changes over the last few decades has been the impact of the internet on labour markets. Not only did the internet change how the unemployed search for work, but it also improved the rate at which workers match with employers. What happens to employment inequality is *prima facie* unclear but depends on the relative shifts in this matching efficiency over time or improvements in search frictions.[18] Changes in the degree of specialization, skills match, institutional policies, etc., can alter the efficiency of the market. For example, occupations where most people with moderate education are usually employed now require onerous training and registration process.[19]

As Moscarini (2001) notes, this increased specialization may force the graduates to only apply to a few job opportunities that match their specific skills, while dropouts still apply to wider job openings. If this is the case, matching efficiency for graduates is likely to drop while that of high school dropouts rise, and the changes would help narrow the employment gap. Mandated minimum wage increases, predominantly affecting the low-educated, can also cause a decrease in matching efficiency. Because both factors that push efficiency down and up happen simultaneously, it is not

[18] Improvements in matching efficiency can also be interpreted as a reduction in search frictions.

[19] Examples include auctioneers, hair stylists, florists, and manicurists, which did not need formal training in the past.

Table 2.2: Matching Function Parameters Across Education Groups

Education Group:	High School Dropouts			Graduates		
	1990	2019	Full sample	1990	2019	Full sample
Activity:	(1)	(2)	(3)	(4)	(5)	(6)
Matching elasticity, α	0.8703	0.5694	0.8777	0.7435	0.7003	0.8262
	(0.1386)	(0.0887)	(0.0098)	(0.1053)	(0.0807)	(0.0075)
Matching efficiency, m_0	0.3107	0.8320	0.3880	0.3749	0.5365	0.4063
	(0.0800)	(0.1531)	(0.0087)	(0.0502)	(0.0659)	(0.0054)

Notes: The parameters are estimated from a labour market search and matching technology in Section 2.4.2 where m_0 and α are the efficiency and matching elasticity parameters, respectively. Standard errors are in parentheses. *Data Source:* Labour Force Survey.

immediately evident in which direction efficiency trends. Table 2.2 provides estimates of the matching efficiency and elasticities of matches by education. The parameter m_0 is a catchall estimate that tracks the net change in efficiency. I estimate results in the table using a matching function in Section 2.4.[20]

Interestingly, the matching elasticity, α of high school dropouts decreased by 35%, while the same estimate for graduates fell only by 6%. This change corresponds to a decrease in workers' bargaining power in the model below. Although a decrease in bargaining power may affect workers' compensation (wages), it may also help to keep the employment of low-educated workers high as firms find it more profitable to hire them. The higher profits encourage firms to hire more workers. In this case, the estimates indicate that the employment gap would have been higher had the low-educated workers' bargaining power not significantly adjusted downwards. The full sample estimates in Columns (3) and (6) are from data covering 1990 to 2019. I use them as inputs in robustness checks in the appendix, where I assume that the matching elasticity α remains constant.

The second row shows results for matching efficiency. The esti-

[20] The values in the table form part of the model parameters in the model below.

mates indicate that the matching efficiency, m_0 for high school dropouts more than doubled between 1990 and 2019, while the efficiency parameter for graduates increased by only 43%. These results indicate that changes in the labour market, either due to advancements in search technology or other factors such as specialization and diversification, were most beneficial to high school dropouts. Model results in Section 2.6 show that this channel effectively lowered employment inequality.

2.4 Model

This section outlines the model of search and matching under frictions, heterogeneity ability, and variable search effort. It is a search model where individuals choose their education and how much job search effort to exert in the labour market when out of employment. Workers can either be employed (e) or nonemployed (n), and the economy is populated by two types of agents: high school dropouts and graduates. The workers are heterogeneous in ability, which affects their education choice. In turn, education choice is closely related to the type of employment a worker can find. The chapter builds on Wolcott (2021) by adding variable search effort and nonemployment duration dependence. However, workers have to choose the optimal level of search because the process is costly in terms of foregone leisure.

2.4.1 Economic Environment

Workers are infinitely lived within a discrete-time framework, and the economy is populated by workers endowed with one unit of labour but differ in their ability such that there are $\mathcal{A}$ types of workers indexed by $x \in \{x_1 < x_2 < ... < x_{\mathcal{A}}\}$, where $x_1 = 0$. In the model, I calibrate the number of ability levels into deciles such that $\mathcal{A} = 10$. As in Wolcott (2021), worker ability is permanent and observable to employers and is approximated by a discrete log-normal distribution with A types of ability in the economy. Heterogeneity in ability implies that the economy has submarkets of workers and each submarket has $\mathcal{A}(x)$ workers who are either employed $e(x) \in [0, 1]$, or nonemployed $n(x) \in [0, 1]$, and aggregate population is normalized to one such that $\sum_x [e(x) + n(x)]\mathcal{A}(x) = 1$. The ability levels match the number of submarkets in the economy, which help to avoid crowding out between workers of different ability. Workers search for new job opportunities when they are out of employment. As mentioned above, each worker j chooses an occupation depending on the level of education attainment such that $j \in \{L, H\}$ where L is high school dropouts and H is graduates.[21]

Because worker ability is observable, firms post vacancies $v(x)$ specific to a submarket x and a successful match produces output $y(x)$. Posting vacancies specific to a submarket ensures that a firm minimizes expenditure on recruitment costs. There are two occupation-specific production technologies; output from these technologies is perfect substitutes. This means that workers can be distinguished by their productivity, and firms pay a worker-type-specific cost to post a vacancy. I assume that low-educated workers are employed in sectors that mainly involve routine tasks, e.g., manufacturing production line, where ability plays a limited role in determining output. With this assumption, the ability of low-educated workers is normalized to one, while production for the high-educated workers depends on their ability such that the production function is:

$$y_{jt}(x) = \begin{cases} A_L & \text{if} \quad j = L \\ A_H x & \text{if} \quad j = H. \end{cases} \tag{2.1}$$

[21]The assumption is that education attainment determines the industries individuals find work.

When there is a successful match in this period, firms produce output and pay wages in the next period.

2.4.2 Matching Function

For production to occur, nonemployed workers and hiring firms match via a Cobb-Douglas constant returns to scale technology. Successful job matches in each submarket x and occupation $j \in \{L, H\}$ are given by $M_j(x) = m_{0j}S_j(x)_t^{\alpha}v_t(x)^{(1-\alpha)}$ where m_{0j} is a measure of matching efficiency, v_j are open vacancies, $\alpha \in [0,1]$ is the elasticity of matches, and S_j is the total job search effort by nonemployed workers.[22] Shifts in matching efficiency parameters m_{0j} capture changes in search frictions in the labour market. If this parameter changes more for the low-educated (see Table 2.2), it will contribute to reducing employment inequality. Let the degree of labour market tightness be given by $\theta_j = v_j/S_j$, then it takes $q_j = M_j/v_j = m_{0j}\theta_t^{-\alpha}$ for a firm to fill an open vacancy. The probability of finding a match for a unit of search is $f(q_j) = m_{0j}\theta_t(x)^{(1-\alpha)} = \theta q_{jt}(\theta)$ and the probability that a worker of type j finds work in submarket x is $s_j(x)f(q_j)$, where s is search effort and normalized such that $s_j(x) \in [0,1]$.

This matching function differs from Wolcott (2021) because it includes the direct effect of search effort s and nonemployment duration dependence of job-finding rates Λ. Matches in the economy are a function of aggregate search effort $S_j = s_j(x)n_j(x)$ instead of the number of nonemployed workers $n_j(x)$. The job search process for each worker involves costs $c(s_j(x))$ due to forgone leisure. Exerting more effort in job search increases the chance of finding employment, but it means forgone leisure. This trade-off and the assumption that search costs are convex in s ensures the ex-

[22]This matching function is common in the literature. See Nakajima (2012), Gomme and Lkhagvasuren (2015), Kudlyak and Hornstein (2017), Pizzinelli and Speigner (2017) and Setty and Yedid-Levi (2020).

istence of an optimal search effort in each submarket. The duration of nonemployment spells affects the rate of successful matches through its effect on job-finding rates. The job-finding rates in the economy are determined by

$$\tilde{f}(\Lambda_j, s_j, q_j) = s_j(x) f(q_j) \Lambda_j,$$

where $\Lambda_j = \frac{1}{24}\sum_{\tau=1}^{24} \Lambda_j(\tau)$. As defined in Section 2.3.1, $\Lambda_j(\tau) \in [0,1] \quad \forall \tau \in \{1, 2, 3, ..., 24\}$ implies that $\Lambda_j \in [0,1]$. When $\Lambda_j = 1$, it means all workers in nonemployment are newly-nonemployed and nonemployment duration does not affect job-finding rates. As the number of workers in long-term nonemployment increases, Λ_j decreases towards zero, which lowers the overall job-finding rate.

2.4.3 Value functions

2.4.3.1 Firms

Given that there are submarkets, firms post ability-specific vacancies. They decide whether to post a vacancy in the L or H market. Because L market is independent of ability, firms treat it as a single labour market. Let $\mathcal{J}_{jt}(x)$ be the present discounted value for a firm that match to a worker of ability x and occupation $j \in \{L, H\}$, and $\mathcal{V}_{jt}(x)$ be the present discounted value of a firm searching for a worker. In a stationary equilibrium,

$$\mathcal{J}_{jt}(x) = \underbrace{y_{jt}(x) - \omega_{jt}(x)}_{\text{profit from match}} + \beta \underbrace{\left[\delta_j \mathcal{V}_{jt+1}(x) + (1-\delta_j)\mathcal{J}_{jt+1}(x)\right]}_{\text{expected benefit from current match}}, \tag{2.2}$$

where $\beta = (1+r)^{-1}$ is the discount factor. When there is a match, the firm gains $y_{jt}(x)$ and pays an endogenous wage $\omega_{jt}(x)$ to worker j of ability x, and in the next period the match can either be destroyed by exogenous occupation-specific job destruction rate δ_j, or continue with probability (1-

δ_j). On the other hand, the value for a firm searching for a worker is

$$\mathcal{V}_{jt}(x) = \underbrace{-\kappa_j}_{\text{cost of posting a vacancy}} + \beta \underbrace{\left[(1-q_{jt})\mathcal{V}_{jt+1}(x) + q_{jt}\mathcal{J}_{jt+1}(x)\right]}_{\text{expected benefit of successful match}}, \tag{2.3}$$

where κ_j is the cost of posting a vacancy. The firm gains nothing but still pays the cost of searching for a worker κ_j, and in the next period, either matches with a worker and receives J_{jt+1} or continues to search.

2.4.3.2 Workers

Let $\mathcal{W}_{jt}(x)$ be the discounted value of an employed worker, and $\mathcal{U}_{jt}(x)$ be the present discounted value of being in nonemployment such that

$$\mathcal{W}_{jt}(x) = \omega_{jt}(x) + \beta \left[\delta_{jt}\mathcal{U}_{jt+1}(x) + (1-\delta_{jt})\mathcal{W}_{jt+1}(x)\right]. \tag{2.4}$$

An employed worker receives wage $\omega_{jt}(x)$ in the current period, and in the next period, can either remain employed, or the job is destroyed with an exogenous probability δ_{jt}. The value of nonemployment $\mathcal{U}_{jt}(x)$ is defined as

$$\mathcal{U}_{jt}(x) = \max\left[\mathcal{U}^c_{L,t}(x), \mathcal{U}^c_{H,t}(x)\right], \tag{2.5}$$

where $\mathcal{U}^c_{L,t}(x)$ and $\mathcal{U}^c_{H,t}(x)$ represents the continuation value of nonemployment for low-educated and high-educated workers, respectively. An alternative interpretation is that $\mathcal{U}^c_{L,t}(x)$ is the nonemployment value for workers searching for work in the low education submarket.[23]

The continuation value of nonemployment for a worker $j \in \{L, H\}$ searching for work in submarket x is defined by the condition

$$\mathcal{U}^c_{jt}(x) = b_{jt} - c(s_j(x)) + \beta\left[\tilde{f}(\Lambda_j, s_j, q_j)(x)\mathcal{W}_{jt+1}(x) + (1-\tilde{f}(\Lambda_j, s_j, q_j)(x))\mathcal{U}_{jt+1}(x)\right], \tag{2.6}$$

[23] There are no benefits for low-educated workers to search for work in the high-education submarkets, and because the low-education submarket pays lower wages, high-educated workers have no incentive to search for jobs in that submarket.

where b_j accounts for a worker's outside option that encompasses leisure, home production, or government insurance programs such as employment insurance, disability insurance, means-tested benefits and progressivity of the tax system.[24] The term $b_{jt} - c(s_j(x))$ is the value of nonemployment relative to market work, with $c(s_j(x))$ as the cost of job search. As a catchall parameter, changes in b_j mirror shifts in labour supply in the model. A rise in employment inequality driven by changes in labour supply preferences happens when the changes are different across groups. For example, a generous safety net that disproportionally benefits high school dropouts may lead to a relatively bigger drop in their labour supply.

As in Wolcott (2021), high-educated workers can switch across occupations, or perceive them as two different people at different times, but with the same ability. A nonemployed worker chooses search effort $s_j(x)$ subject to an increasing convex cost function $c(s(x))$ that maximises expected value of nonemployment

$$\mathcal{U}_j(x) = \max_{s_j(x)} \left\{ b_{jt} - c(s_j(x)) + \beta \left[\tilde{f}(\Lambda_j, s_j, q_j)(x) \mathcal{W}_{jt+1}(x) + (1 - \tilde{f}(\Lambda_j, s_j, q_j)(x)) \mathcal{U}_{jt+1}(x) \right] \right\}, \tag{2.7}$$

with average job search effort in each submarket $\bar{s}_j(x)$. Effectively changes in labour supply correspond to shifts in the net value of nonemployment $b_{jt} - c(s_j(x))$. This setup makes it possible to directly measure the effect of changes in search effort on the employment gap. The search cost function is convex in search effort, twice differentiable, and $c(0) = c'(0) = 0$.[25]

[24]This parameter can represent many more factors such as value of acquiring more education, health shocks, childcare (Goensch et al., 2021). In the model, b is a catchall parameter that represents all factors that are likely to cause change in labour supply.

[25]There is no on-the-job search in the model.

2.4.3.3 Wage bargaining

Workers and firms in each submarket share the match surplus and negotiate wages via Nash bargaining:

$$\omega_{jt}(x) = \text{argmax}\left(\mathcal{W}_{j,t}(x) - \mathcal{U}_{j,t}(x)\right)^{\phi_j}\left(\mathcal{J}_{j,t}(x) - \mathcal{V}_{j,t}(x)\right)^{(1-\phi_j)}, \quad 0 < \phi_j < 1 \quad (2.8)$$

where ϕ is bargaining power of workers. Let $\Xi_{jt}(x) = \max\{\mathcal{W}_{j,t}(x) + \mathcal{J}_{j,t}(x) - \mathcal{V}_{j,t}(x) - \mathcal{U}_{j,t}(x), 0\}$ be the total surplus from a match in occupation j and submarket x. The firm receives share $(1-\phi_j)\Xi_{jt}(x)$ of the surplus with the rest going to the worker, and a match continues if $\Xi_{jt}(x) > 0$. The maximization problem in Eqn. (2.8) yields $\mathcal{W}_{j,t}(x) - \mathcal{U}_{j,t}(x) = \frac{\phi_j}{1-\phi_j}[\mathcal{J}_{j,t}(x) - \mathcal{V}_{j,t}(x)]$.

2.4.4 Steady state equilibrium

In steady-state, I assume that the free entry condition holds such that the present discounted value for a firm should be zero, $\mathcal{V}_j(x) = 0$. This condition is essential because if $\mathcal{V}_j(x) > 0$, more firms join the submarket to post vacancies, and there is no incentive to search for a worker when $\mathcal{V}_j(x) < 0$. For instance, if $\mathcal{V}_j(x) > 0$ firms post more vacancies, labour market tightness increases, and the job filling rate falls until $\mathcal{V}_j(x) = 0$. I assume that each occupation has an infinite number of firms free to join each ability submarket and post vacancies. As in Wolcott (2021), there are no submarkets in the occupations where the high school dropouts find employment, so firms treat it as a single market. The assumption is that ability does not matter in occupations where high school dropouts find employment.

2.4.4.1 Optimal search effort

Workers take $q_j(\bar{s}, \theta)$ as given so the first order condition from Eqn. (2.7) is

$$c'(s_j(x)) = \beta f_j(\bar{s}, \theta) \left[\mathcal{W}_j(x) - \mathcal{U}_j(x)\right].$$

Combining the first order conditions Eqns. (2.4)-(2.6) yields the optimal search effort in each submarket as

$$\omega_j(x) - b_j = \left[\frac{1}{\beta} + \delta_j - 1\right] \frac{c'(s_j(x))}{\theta_j q_j} + s_j(x)c'(s_j(x)) - c(s_j(x)), \tag{2.9}$$

with the twice differentiable search cost function as $c(s) = \frac{s_0}{\gamma} s^{\gamma}$, $s_0 > 0$ and $\gamma > 1$. The parameter s_0 represents a scale parameter and is set to one in the model. The term on the right of Eqn. (2.9) represents the marginal cost of search effort and the left term is the discounted value of employment. Eqn. (2.9) implies that nonemployed workers equate the discounted value of employment to the marginal cost of search effort. Endogenous search effort helps to amplify the reaction of nonemployment and vacancies to shifts in other variables, e.g., productivity. As Gomme and Lkhagvasuren (2015) note, endogenous search effort can amplify nonemployment and vacancies in three ways. First, when productivity y_j rises, firms post more vacancies, and workers increase the intensity of search, which incentivizes firms to create more vacancies. Second, when workers increase search effort, it lowers the value of nonemployment (search is costly—$c(s)$) without necessarily reducing profits, so firms respond by posting more vacancies. Lastly, a rise in search effort increases the probability of finding a successful match and vice versa. These three points emphasize the fact that endogenous search effort amplifies the responses of nonemployment, vacancies, and labour market tightness when productivity shifts.[26] As Table 2.3 below shows,

[26] Labour market tightness moves when nonemployment and/or vacancies shift by definition given that $\theta = v/S$.

the model results indicate a significantly higher increase in productivity for graduates between 1990 and 2019 than that of high school dropouts. These different shifts can then be the source of rising employment inequality.

2.4.4.2 Job creation curve

Eqns. (2.2), (2.3), and the free entry condition give the condition that explains the behaviour of firms in creating jobs (job creation curve). The cost of posting vacancies in the market creates inefficiencies because it reduces surplus from a match; thus, the marginal product differs from the paid wage. Rather, firms aim to equate the expected marginal benefit of hiring a worker to the cost as

$$y_j(x) - \omega_j(x) - \frac{\kappa_j \left(\frac{1}{\beta} + \delta_j - 1\right)}{q_j(\theta)} = 0. \tag{2.10}$$

Ceteris paribus, if the costs of posting vacancies rise, either the marginal product of workers $y_j(x)$ has to rise or the agreed wage must fall. Similarly, a rise in wages or labour market tightness implies that the marginal product of workers must also rise. Changes in $y_j(x)$ over time correspond to shifts in labour demand. Employment inequality rises if the gap in workers' marginal product increases over time.

2.4.4.3 Wages

Using the Nash wage bargain condition and Eqns. (2.2)-(2.6) yields the wage equation

$$\omega_j(x) = (1 - \phi_j)[b_j - c(s_j)] + \phi_j[y_j(x) + s_j \kappa_j \theta_j(x)]. \tag{2.11}$$

A tight labour market helps workers to bargain for higher wages. Although increasing search effort lowers the utility flow of nonemployment due to the rise in search cost, it also increases the probability of finding a job. These opposing effects mean that there must be an optimal level of search where the loss to the value of nonemployment matches the gain from increasing the chance of finding work.

2.4.4.4 Nonemployment

Steady state happens when the number of workers flowing out of nonemployment in each ability submarket matches the inflow of workers into the state. In each period, $\tilde{f}(\Lambda_j, \bar{s}_j, q_j)n_j(x)$ workers flow from nonemployment to employment, while $\delta_j e_j(x)$ transition from employment to nonemplyment. Equating these two flows plus the fact that $e_j(x)+n_j(x) = 1$ gives the nonemployment rate in steady state

$$n_j(x) = \frac{\delta_j}{\delta_j + \tilde{f}(\Lambda_j, \bar{s}_j, q_j)}, \tag{2.12}$$

such that this rate is a function of job destruction rate, search effort and tightness.

2.4.4.5 Education choice

Nonemployed workers maximize the discounted future value of being employed in either low-educated or high-educated occupations to choose the optimal level of education. Thus, education choices is an endogenous decision that requires solving

$$U_j(x) = \max_j \left[\frac{\left(\frac{1}{\beta} + \delta_j - 1\right)(b_j - c(s_j))}{(1-\beta)\left(\frac{1}{\beta} + \delta_j + \tilde{f}(\Lambda_j, \bar{s}_j, q_j) - 1\right)} \right] + \left[\frac{\tilde{f}(\Lambda_j, \bar{s}_j, q_j)\omega_j(x)}{(1-\beta)\left(\frac{1}{\beta} + \delta_j + \tilde{f}(\Lambda_j, \bar{s}_j, q_j) - 1\right)} \right] \tag{2.13}$$

in steady state. This education choice eqn. is derived from substituting Eqn. 2.6 in Eqn. 2.5. The steady-state equilibrium for the model is described by Eqns (2.9)-(2.13).

2.5 Calibration

Conducting a steady-state analysis requires comparing periods that have similar characteristics. This chapter compares employment rates in 1990 and 2019, which have similar characteristics because they correspond to business cycle peaks in Canada. The exercise involves calibrating the model above to match data moments from these two periods to determine the causes of the rise in employment inequality. The model incorporates five endogenous channels: shifts in labour supply, shifts in labour demand, shifts in labour market efficiency (search frictions), shifts in nonemployment duration dependence, shifts in job search effort and one exogenous channel: shifts in job separation rates. In the model, shifts in parameters corresponding to different channels translate to either a rise or decrease in the employment gap.

First, I estimate both the matching efficiency m_o and elasticity of matches α using nonemployment, vacancies, search effort, and job-finding rates data from the Labour Force Surveys.[27] The results are in Table 2.2, and already discussed in Section 2.3.2.3. The second step involves jointly estimating search effort, the value of nonemployment (supply-side parameters), and workers' productivity (demand-side parameters).The last stage is to estimate the mean and standard deviation of ability that I calibrate by matching the share of high-educated men in the economy. Workers of different ability are group into $\mathcal{A} = 10$ deciles.[28]

Identifying the demand, supply and search effort channels hinges on Eqns. (2.9)-(2.11), which are the optimal search effort, job creation

[27] To estimate the parameters, I use data covering 5 years at a monthly frequency. This selection corresponds to 1986-1990 for 1990 estimates and 2015-2019 for 2019.

[28] I also tried to group ability into percentiles $\mathcal{A} = 100$. Increasing the ability levels significantly increases the time to solve the model without changing the main results.

curve and wage equation, respectively. The three equations with three unknown parameters are enough to untwine the role of search effort, labour demand, and labour supply. Following Wolcott (2021), I use minimum distance estimation by first guessing the initial values of $\{A_L, A_H, b_L, b_H\}$ in each period. I then estimate tightness, search effort, and wages using Eqns. (2.9)-(2.11), minimizing the squared percent difference between model moments and true values. Unlike the data, the model includes heterogeneous ability and search effort for the graduates. Therefore, I find average model moments over ability and search effort before computing the squared percent difference. The last step involves calibrating the parameters (mean and standard deviation) that describe ability. The distribution of ability is assumed to be a discrete log-normal. I calibrate the two parameters by targeting the share of high-educated men, which corresponds to 73% in 1990 and 92% in 2019.

2.6 Results

2.6.1 Calibration Results

To conduct the analysis, I calibrate the parameters in equations (2.9)-(2.13), which summarizes the steady state. The calibrated parameters for 1990 and 2019 are in Table 2.3. The discount factor is calibrated to 0.9967 to match an annualized interest rate of 4%. The curvature parameter of the search cost function, γ is set at 2, which is from the literature.[29] On the other hand, the scale parameter s_0 of the search cost function is set to 1 for both groups (high school dropouts and college graduates).

To estimate the parameters for the matching function (m_0, α), I use

[29] This parameter is consistent with the value used in the literature. See Christensen et al. (2005), Lise (2013), Gomme and Lkhagvasuren (2015) and Çenesiz and Guimarães (2022).

Table 2.3: Steady State Parameter Calibration

Education Group:	Dropouts		Graduates		
	1990	2019	1990	2019	Source
Parameter:	(1)	(2)	(3)	(4)	(5)
Discount factor, β	0.9967	0.9967	0.9967	0.9967	monthly rate
Vacancy posting cost, κ	0.20	0.20	0.66	0.66	share of earnings (1990)
Elasticity of search cost, γ	2.00	2.00	2.00	2.00	Lise (2013)
Matching elasticity, α	0.8703	0.5694	0.7435	0.7003	estimation
Job finding rate, δ	0.2331	0.3818	0.2438	0.3734	LFS
Bargaining weight, ϕ	0.8703	0.5694	0.7435	0.7003	Hosios condition
Separation rate, δ	0.0252	0.0618	0.0086	0.0140	LFS
Matching efficiency, m_0	0.3107	0.8320	0.3749	0.5365	estimation
Duration dependence, $\Lambda(\tau)$	0.2134	0.2720	0.1925	0.2387	LFS
Productivity, A	0.9456	0.9493	2.6842	3.5681	calibrated
Search effort, $\bar{s}$	0.8618	0.8750	0.8959	0.9267	calibrated
Value of outside option, b	0.9529	1.2419	0.7447	0.7104	calibrated
Mean of logarithmic ability, μ_x			0.3	0.3	calibrated
SD of logarithmic ability, σ			0.4	0.4	calibrated

Notes: The table tabulates parameters from the Canadian data or internally calibrated values over the two periods. HS - High School, LFS - Labour Force Survey.

a minimum-distance procedure that minimizes the sum of squared differences between the observed job-finding rate and the job-finding rate predicted by the model. The minimum distance estimates are both in Tables (2.2)-(2.3), and discussed in Section 2.3. The matching efficiency parameter for low-educated workers more than doubled while that of high-educated workers increased by only 43%. This rise in matching efficiency is more likely to reduce inequality because it improves the matching rate between low-educated workers and firms. The chapter assumes that the allocation of less-educated and high-educated labour is efficient; the Hosios (1990) condition holds. When the Hosios condition holds, the worker's bargaining power parameters ϕ_j are equal to the matching elasticity parameters α_j. Changes in $\{\phi_L, \phi_H\}$ correspond to shifts in the bargaining power of workers. This interpretation shows that the bargaining power of low-educated workers decreased by 35% while the power of high-educated workers fell only by 6%. Lower bargaining power implies that firms receive a bigger share of the surplus from a successful match in 2019 than in 1990. A steeper downward shift in the bargaining power of low-educated workers relative to high-educated workers is likely to moderate the rise in employment inequality.

The evidence on the cost of posting vacancies for firms in Canada is scant. I assume that the cost of posting vacancies varies with the level of education; that is, it is more costly to hire a high-educated worker than a less-educated worker. Hiring high-educated workers costs more due to the general cost of recruitment and the cost related to the screening process to identify a worker with the correct ability. The cost of posting a low-educated vacancy is set at 20% of the 1990 wages of low-educated workers. The wage of low-educated workers in 1990 is normalized to one. This calibration follows Wolcott (2021). The cost of posting a vacancy for a high-educated worker is 66% of the 1990 wages of low-educated workers.[30] The job-finding and separation rates are derived from LFS as in Shimer (2005). The separation rates for both groups increased during the two periods, but more so for the low-educated workers. The sharp rise in separation rates of less-educated men relative to high-educated men indicates one of the potential causes of the increase in the employment gap.

There has been an upward shift in the nonemployment duration dependence of job-finding rates parameters for both groups. The nonemployment duration dependence parameters Λ_L indicate that changes in the

[30]Increasing the costs of posting low-educated vacancies while holding the costs of posting high-educated vacancies constant slightly changes the role of the demand channel. However, the results are insensitive to different costs of posting vacancies if I simultaneously shift both costs.

composition of nonemployment spells have helped push job-finding rates for low-educated workers by about 27%. However, the increase is lower for high-educated workers.[31] This upward shift is consistent with the trends in Fig. 2.2 that show more significant improvements in job-finding rates across nonemployment duration. These results are also consistent with the fact that the share of individuals in long-term nonemployment (> 24 weeks) decreased between 1990 and 2019.

The rest of the parameters are calibrated from equations (2.9)-(2.13). On the demand side, the productivity of low-educated workers remained almost the same (increase by only 0.3%) while that of high-educated workers increased by 33%. These results would be consistent with the hypothesis that automation complements cognitive tasks, thereby boosting the productivity of high-educated workers at the expense of low-educated workers. However, the graduates (high-educated) group also includes high school graduates who are usually employed in occupations equally vulnerable to automation. Whether these changes translate to automation and trade competition reducing the opportunities for low-educated workers is not immediately clear. The job search effort channel shows an increase in search effort between 1990 and 2019. This rise is consistent with advancement in search technology that makes searching more effective. The advent of the internet and job boards opens opportunities for more options on how to search for a job. The increase in search effort in Table 2.3 also matches the increase in Fig. A-7.

On the supply side, the parameters are consistent with government programs being more generous to high school dropouts over time or a rise

[31] It is crucial to remember that Λ and Fig. 2.2 show shifts in finding rates relative to the job-finding rates of the newly nonemployed.

in preference for leisure.[32] The parameter b shows a rise for low-educated workers, while that of high-educated workers decreased from 0.7447 to 0.7104. These trends are in line with the shifts in Table 2.1 and Fig. A-5. Table 2.1 shows a rise in time spent on leisure for the less-educated and vice-versa for graduates, but also, both men spent more time in home production. An increase in time spent on leisure or home production is equivalent to a rise in the utility flow of nonemployment in the model. Fig. A-5 indicates that public programs have become increasingly available to high school dropouts while the share of graduate men in households receiving benefits slightly fell between 1990 and 2019 (see more discussion in Section 2.6.2.3). All these results indicate a greater incentive for high school dropouts to reduce time spent working, leading to an upward trend in employment inequality.

I calibrate the ability parameters μ_x and σ_ξ by matching the rise in the share of graduates between 1990 and 2019. Fig. 2.3 provides a graphic illustration of the existence of the threshold ability x_ξ around the third decile. This figure uses the optimal parameters in Table 2.3. Fig. A-8 and Fig. A-9 in the appendix show that the threshold x_ξ exists even for different mean, standard deviation and different classifications of ability (from deciles to percentiles).

In the calibration process, I target labour market tightness, wages, and the share of college men in the economy. The results are in Table 2.4. The model matches the targeted moments relatively well, especially tightness in 2019 and the changes in the share of high-educated men. Table A-1 tabulates a comparison between model and data employment rates

[32]The parameter b is a catchall parameter for factors that affect labour supply, and the net value of the outside option excludes search cost. For example, the net value in 1990 for high school dropouts is $b_{L,90} - s_0 \frac{s_{L,90}^{\gamma}}{\gamma} = 0.9529 - 0.8618^2/2 = 0.5816$.

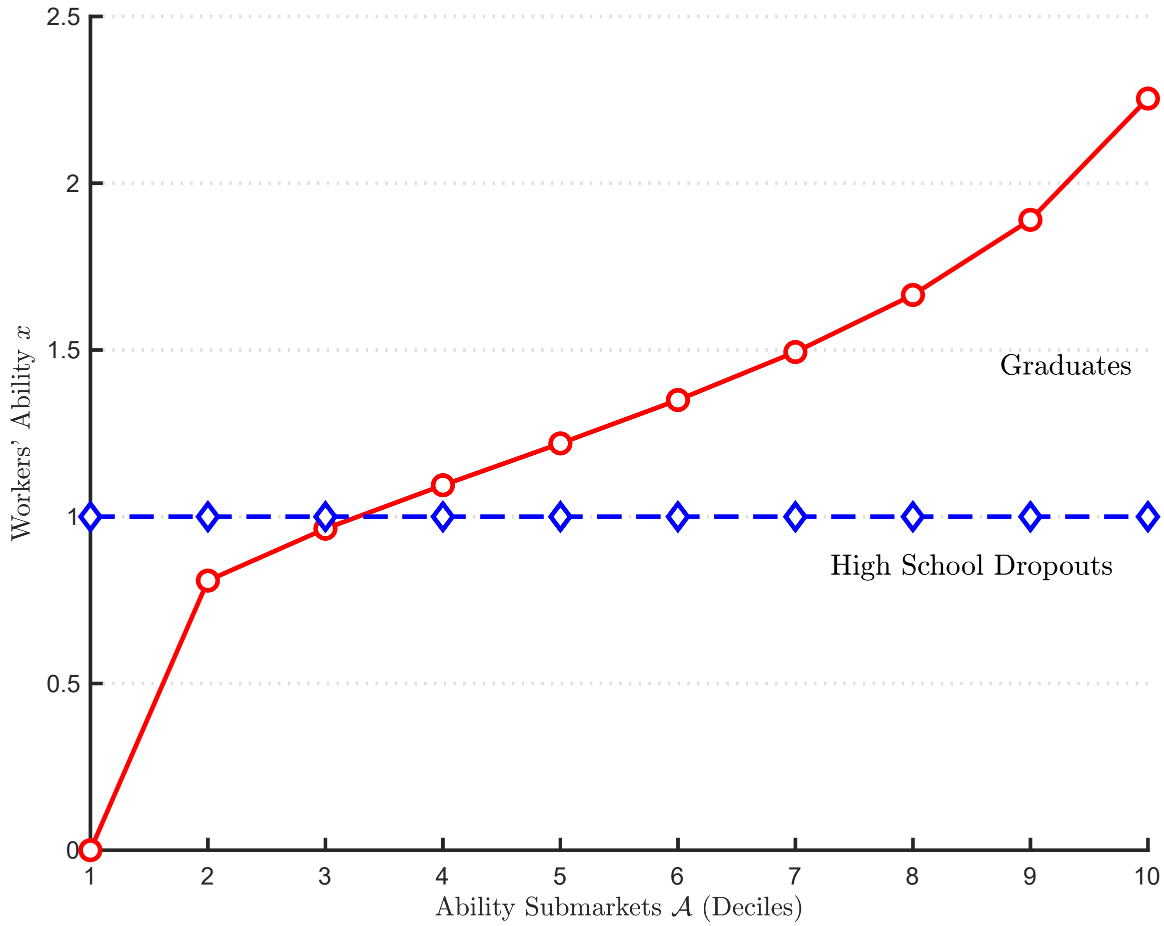

Figure 2.3: Distribution of Ability Across Deciles Using Calibrated Parameters

Notes: The figure depicts the ability values in each submarket (deciles) for high school dropouts and graduates using the optimal ability parameters $\mu_x = 0.3$ and $\sigma_x = 0.4$. Note that there is one submarket for low-educated jobs, which means ability remains constant across deciles (at the normalized value of 1).

(non-targeted moments), which are the moments of interest in measuring employment inequality. The model values are from the steady-state nonemployment rate Eqn. (2.12), a function of separation rates, job-finding rates and job search effort. The employment gap in the data is 6.9 percentage points, and the model predicts a gap of 6.8 percentage points. This result confirms that the model performs well at matching the rise in employment inequality observed in the data. Given that the model matches the employment gap, the next step is to perform counterfactual experiments. In

Table 2.4: Model Targeted Moments

Moment	Explaination	Year	Model (1)	Data (2)	Model Gap (%) (3)	Data Gap (%) (4)
$\theta_{L,90}$	L tightness	1990	0.27	0.25		
$\theta_{H,90}$	H tightness	1990	0.50	0.47	85	88
$\theta_{L,19}$	L tightness	2019	0.31	0.31		
$\theta_{H,19}$	H tightness	2019	0.61	0.61	97	97
$\bar{\omega}_{L,90}$	L wages (normalized)	1990	0.94	1.00		
$\bar{\omega}_{H,90}$	H wages	1990	1.24	1.39	38	39
$\bar{\omega}_{L,19}$	L wages	2019	1.16	1.16		
$\bar{\omega}_{H,19}$	H wages	2019	1.56	1.57	34	35
$\frac{M-\xi_{90}}{M}$	H share	1990	70%	73%		
$\frac{M-\xi_{19}}{M}$	H share	2019	90%	92%		

Notes: The table tabulates the performance of the model against moments from the data. The table summaries labour market tightness, wages, and the share of graduates in the economy.

other words, the aim is to assess what would have happened to employment inequality if only one of the six channels had changed between 1990 and 2019. Fig. 2.4 displays the data gap, model gap and results of the counterfactual exercise.

2.6.2 Counterfactual Results

2.6.2.1 One-Channel Counterfactual Results

The dark blue bar corresponds to a rise in the employment gap in the model (all channels turned on), and the red bar represents a similar rise in the data. These values are the same as the gap in Table A-1. The light blue bars show changes in employment inequality when only one channel is turned on; that is, only one of the parameters in Table 2.3 shifted between 1990 and 2019. The light blue bars indicate that two channels contributed to the rise in employment inequality: shifts in separation rates and labour supply. All else constant, if only separation rates had changed between the

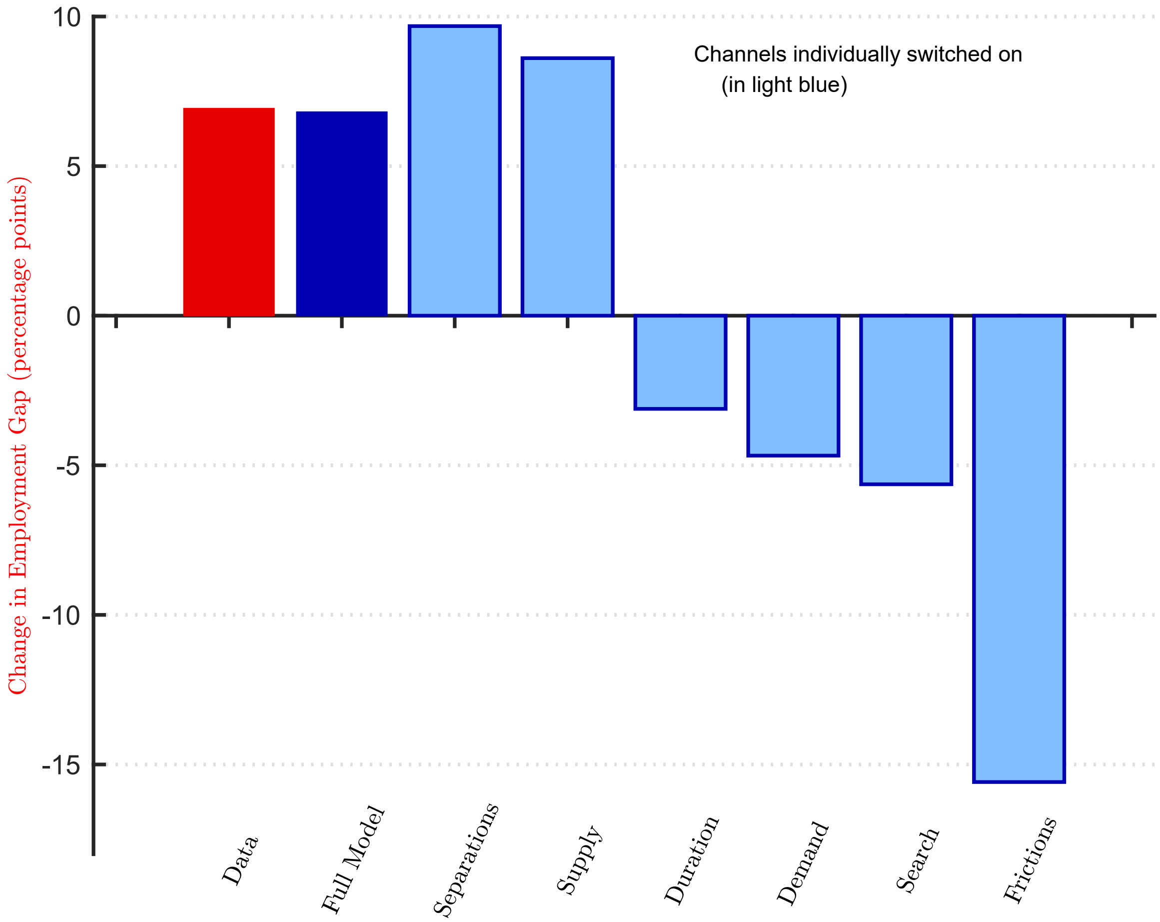

Figure 2.4: Employment Inequality by Channel: Counterfactual Experiments

Notes: The red bar shows the rise in the gap of employment rates (in percentage points) between high school dropouts and graduates between 1990 and 2019 in the data. The dark blue bar is the corresponding gap in the model when all channels are allowed to change between the two periods, whereas the remaining light blue bars captures the results from the counterfactual experiments where only one channel is allowed to change between 1990 and 2019. Search - search effort, Frictions - search frictions or labour market efficiency.

two periods, the gap in employment rates between dropouts and graduates would have risen by almost 10 percentage points.

Similarly, if only labour supply parameters shifted while all other channels remained at 1990 values, employment inequality would have risen by 8 percentage points. On the other hand, changes in labour demand parameters A alone reduce employment inequality by 5 percentage points. These results differ from Wolcott (2021), who reports the opposite for the

US. She finds shifts in labour demand to be the primary driver of the rise in employment inequality in the US between high school and college graduates and that changes in labour supply reduced the gap.

Given the similarities between the US and Canadian economies, the differences in the role of labour supply and demand might seem surprising. However, it can be reconciled by two potential reasons. First, Wolcott (2021) compares the rise in inequality between high school and college graduates, while this chapter contrasts the difference between high school dropouts and graduates. This focus on different groups matters. For instance, Juhn (1992) uses a demand-supply framework to study the causes of the drop in participation rates of high school dropouts in the US and shows that the major contributor to declining employment for this group was labour supply factors. Trade competition and automation might have a muted effect on widening the gap because the graduates group includes high school graduates who are equally likely to be affected by trade competition and automation. In this case, the impact of shifts in labour demand is likely to be similar. It also depends on who is most affected by automation and trade competition. If the assumption that automation heavily affects middle-skilled jobs while complementing low-skilled and high-skilled jobs (as in Dixon et al., 2021) holds, it might explain why shifts in demand-side factors led to a fall in inequality. Automation would increase opportunities for those with university degrees while at the same time lowering opportunities for those with high school diploma. Because both groups are part of the graduates in the model, the net effect depends on which group is affected the most.

The second reason the results are consistent with a muted effect of demand-side shifts is what happened to wages during this period. Shifts in demand that push the employment gap up should have been accompanied by a fall in wages of high school dropouts, but the evidence in Fig. A-3 shows the opposite. A decrease in demand for high school dropouts should lead to downward pressure on their wages. Instead, high school dropout wages have been keeping pace with the rise in graduate wages since the early 2000s, when trade competition began to intensify. The evidence in Table 2.1 (shifts in leisure) and Fig. A-5 (government benefits) is also consistent with significant shifts in labour supply factors. As Binder and Bound (2019) and Eberstadt (2022) note, a rise in the generosity of the safety net can gradually push employment rates down.

The far-right bar shows what happens to the employment gap when only labour market friction parameters change. Table 2.3 already indicates the potential role of shifts in labour market efficiency. There is a more considerable shift for less-educated men than high-educated men. The matching parameters of low-educated men are higher than that of high-educated men in 2019, but it was the opposite in 1990. These shifts entail that improvements in job search technologies and the matching process in the market were more beneficial to high school dropouts. Therefore, improvements in the efficiency of the labour market in matching workers with employers reduced the employment gap by approximately 15.6 percentage points. Wolcott (2021) finds similar results in the US. She argues that the shift in matching efficiency might be due to an increase in specialization for high-education-specific occupations, making it more difficult to find employment.

The changes in the efficiency parameters m_0 measure the role of overall improvements in the job market. The bar second from the right focuses specifically on changes in job search effort. An increase in search effort inherently improves labour market matching efficiency. It is important to note that job applicants in 1990 and 2019 might allocate the same time per week to search for work, but the 2019 job seekers will be at an advantage given that they can apply for more jobs within the same amount of time as their 1990 counterparts. The plot indicates that shifts in job search alone reduced employment inequality by 6 percentage points. This estimate implicitly implies that almost 40% of the improvement in labour market efficiency was due to shifts in job search effort. Fig. A-4 depicts trends in the share of men using a specific job search method by education.[33] Except for the other methods, the proportion of less-educated men using each method seems to get closer to that of high-educated men, especially from around the 2000s.[34] The narrowing down of the shares supports the result that shifts in job search effort led to a decrease in the employment gap.

The middle bar depicts the effect on employment inequality if only the nonemployment duration dependence channel is switched on. The figure shows that shifts in nonemployment duration dependence reduced employment inequality by around 3 percentage points. This result is in line with the increase in Λ tabulated in Table 2.3. Moreover, in the data, the share of high school dropouts in long-term nonemployment decreased by 4 percentage points between 1990 and 2019 (from 29% to 25%). In comparison, the same share for graduates only decreases from 26% to 24%.[35] As noted earlier, an increase in the nonemployment duration dependence parameter can affect inequality by increasing job-finding rates or reducing the potential impact of firms discriminating against long-term nonem-

[33] A more accurate measure of job search effort is time spent on job applications or seeking work, but such information is not available in Canada.

[34] The decrease in the use of more conventional methods like contacting agencies and employers directly is also very evident in the figure.

[35] To reduce clutter, the plot showing these trends is not included in the chapter but is available from the author upon request.

ployed workers. The second channel is not directly modelled in this chapter because that requires richer firm-level data which is difficult to find in Canada.

2.6.2.2 Multiple-Channel Counterfactual Results

Another potential counterfactual experiment is to check what happens to employment inequality if more than one channel changes at a time. Because separations and supply alone contributed to the rise in inequality, I perform counterfactual exercises of interacting these two channels with the other channels that reduce the gap. The results are in Fig. 2.5. The first bar shows the rise in employment inequality when all channels are switched on, the next four depict the interaction between separation rates and the other channels, and then the interaction between supply and the same channels. In the last bar on the right, I ask the question: "would shift in market efficiency alone have managed to cancel the rise in employment inequality driven by shifts in separations and supply factors?"

Although results in Fig. 2.4 are essential in isolating the role of different channels, shifts in the actual economy happen concurrently. A high inflow rate (into nonemployment) without a corresponding rise in the outflow rate is likely to increase the proportion of individuals in long-term nonemployment. As the share of long-term nonemployed rises, job-finding rates fall, and the discrimination effect of firms becomes more salient. Thus, the results show that if shifts in separation rates or labour supply happen concurrently with changes in duration dependence, inequality would have risen by 8 and 2 percentage points (second bar), respectively.

The interaction between separation rates and search frictions or matching efficiency would have led to a rise in inequality by around 4 percentage points, whereas it is the opposite when efficiency interacts with

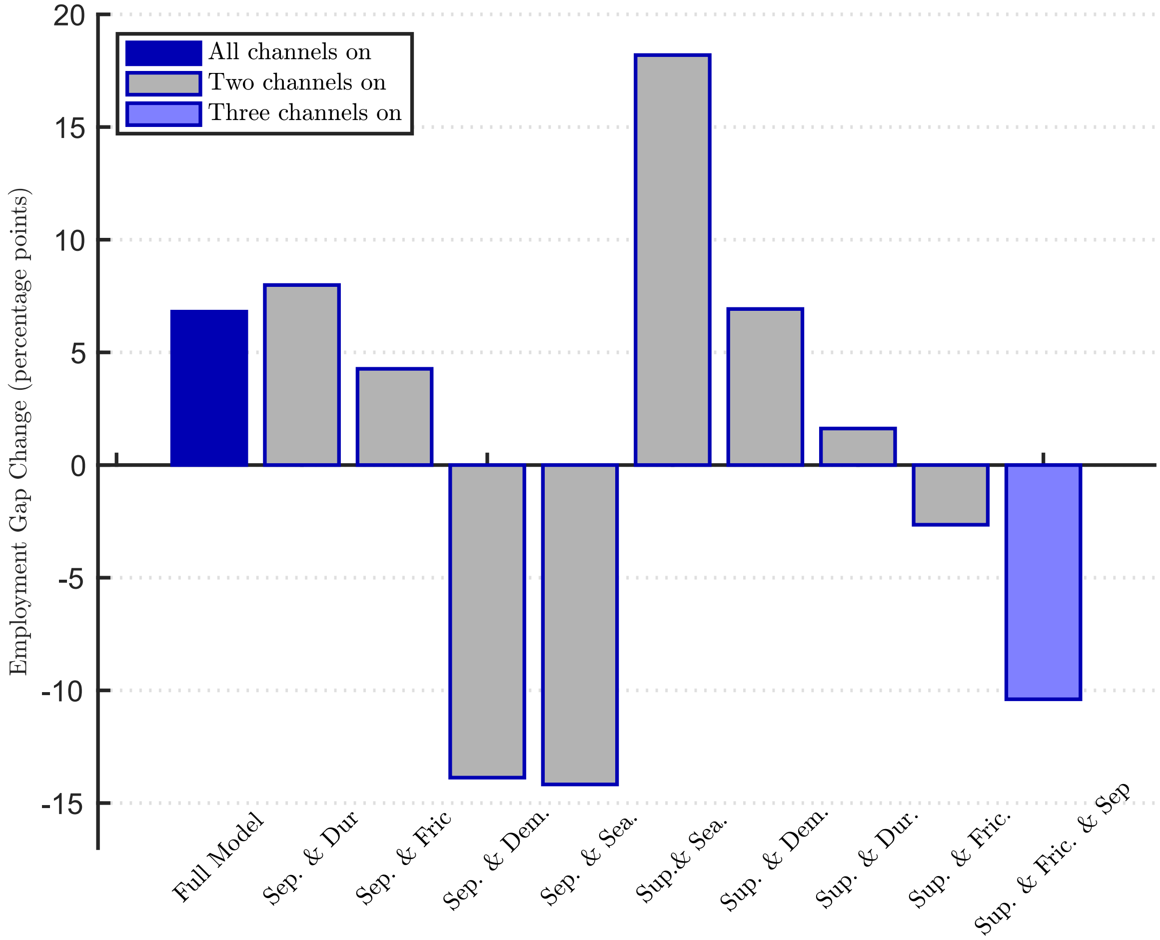

Figure 2.5: Employment Inequality—Two or Three Channels: Counterfactual Experiments

Notes: The red bar shows the rise in the gap in employment rates (in percentage points) between high school dropouts and graduates between 1990 and 2019 in the data. The dark blue bar is the corresponding gap in the model when all channels are allowed to change between the two periods, whereas the remaining light blue bars capture the results from the counterfactual experiments where two or more channels are allowed to change between 1990 and 2019. The abbreviations refer to: Sep- separations, Dur - duration, Fric - labour market efficiency (search frictions), Dem - demand, Sup - supply, and Sea - search effort.

accompanied by an improvement in matching efficiency is supposed to lower the gap. Perhaps a more efficient labour market when separation rates are high cannot significantly help keep the employment rates of low-educated men from falling. The results for job search are as expected and consistent with the findings above. If separation rates rise, but individuals concurrently exert more effort in job search, employment inequality can fall if the effect of the latter is strong enough (5th bar). However, "Sup & Sea" bar shows that a rise in search effort or advancement in job search technology when the high school dropouts are less motivated to work can worsen the problem of inequality. For example, an increase in the value of leisure implies firms have to increase wages to still attract workers. Instead, firms react by posting fewer vacancies, which can discourage search, further compounding the problem. Results in Section 2.7 show the importance of search effort in the model.

The last bar on the right indicates that shifts in frictions alone would have been enough to cancel the rise in inequality due to supply-side and job separation rates. Interestingly the 3rd bar from the left suggests that improvements in search effort alone would not have been enough to cancel the effect of separation rates. I interpret these results as indicating the vital role of shifts in labour supply in the model. For example, the "Sup & Sea" bar indicates that the interaction of other supply factors and job search alone would have led to a rise in inequality by 18 percentage points. If high school dropouts are increasingly unwilling to work, firms will first try to offer higher wages to attract workers (as seen in the rise of real wages in Fig. A-3), after which they reduce the number of low-educated vacancies they post in the market. Sensing the limited availability of vacancies, workers may reduce their effort, further increasing inequality.

2.6.2.3 Drivers of Shifts in Labour Supply

The results in the previous section indicate that labour supply is the primary driver of the rising employment inequality. The challenge in the above analysis is that it does not identify which labour supply factors are contributing to the rise in the employment gap. The labour supply parameter is a catchall parameter that includes factors such as home production, leisure and government insurance programs (progressive earnings taxation, unemployment insurance, means-tested programs, disability insurance and social security). Table 2.1 in Section 2.3 already showed significant changes in time spent on home production, especially for graduates. The fact that time spent on home production significantly increased for graduates than dropouts indicates that this cannot be the leading cause of shifts in labour supply. The table also shows the changes in time spent on leisure. While leisure time for dropouts remains relatively constant between the two periods, the fact that graduates significantly reduced time on leisure partly explains the effect on labour supply.

Government insurance programs can also provide a safety net for workers, influencing their job search and employment choices. Fig. 2.6 indicates that the share of public transfers in household income for high-educated men decreased from 30% in 1990 to only around 12% by 2018. In comparison, the share for low-educated men initially fell from 30% in 1990 to 24% by 2006 before recovering to 30% by 2018. These diverging trends suggest that public transfers are more critical for low-educated households than high-educated ones. These trends are also consistent with the evidence in Fig. A-5 that shows an increase in the share of low-educated men living in households receiving some form of public transfers.

Another important source of income for nonemployed men is employment insurance. Fig. A-10 plots the changes in the average tax for prime-age men in Canada. Although the average tax for low-educated and high-educated men rose from the 1980s to the early 1990s, the decrease post-1992 recession is also similar. Thus, changes in the average tax level are less likely to cause shifts in the labour supply documented in the model.

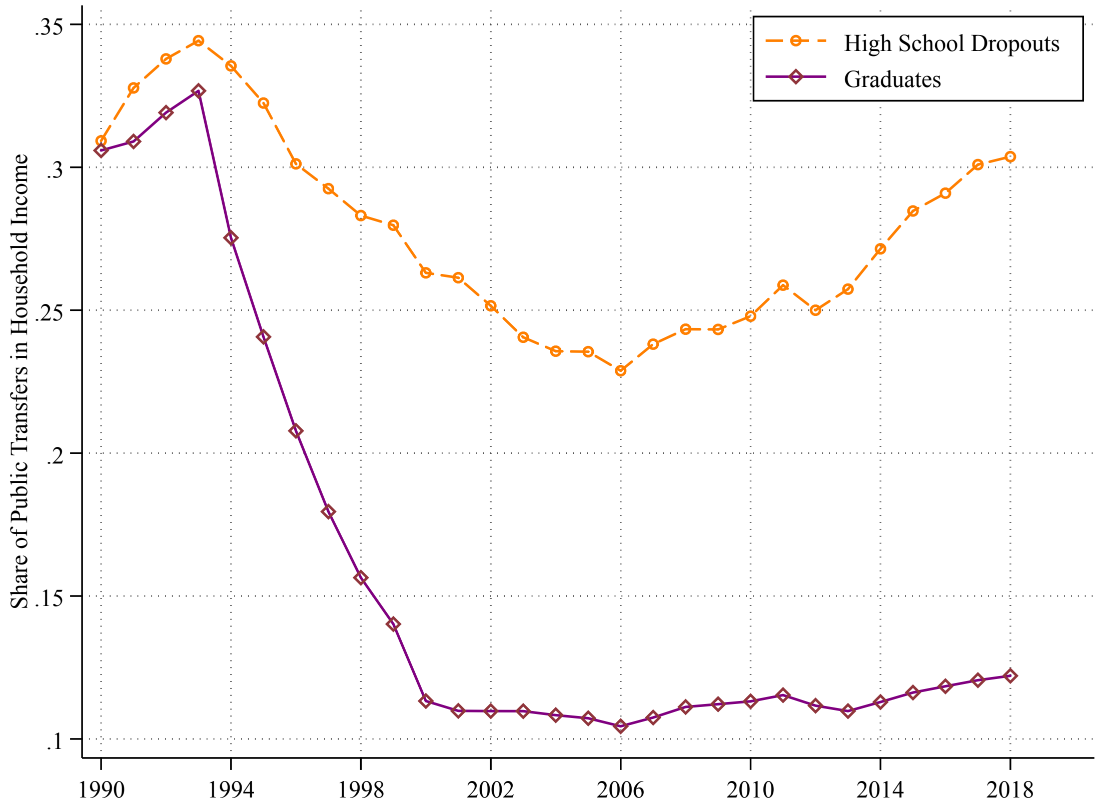

Figure 2.6: Share of Public Transfers in Household Income by Education

Notes: The figure shows the share of public transfers in household income for prime-age men in Canada. These shares are derived from three different surveys: Survey of Consumer Finances, 1980-1996, Survey of Labour and Income Dynamics,1997-2011 and Consumer Income Surveys, 2012-2018.

The employment rates of men can also fall if women are increasingly deciding to find work than they did in the past, creating an additional household income and reducing the burden for men. For instance, Keldenich and Knabe (2022) show that women tend to respond to their partner's unemployment status. Employment inequality will then rise if the spouses of one group change their labour supply choices by a relatively more considerable margin. For example, if wives of low-educated men choose to join the labour market more than those of high-educated men, the added-worker effect will substantially impact the behaviour of

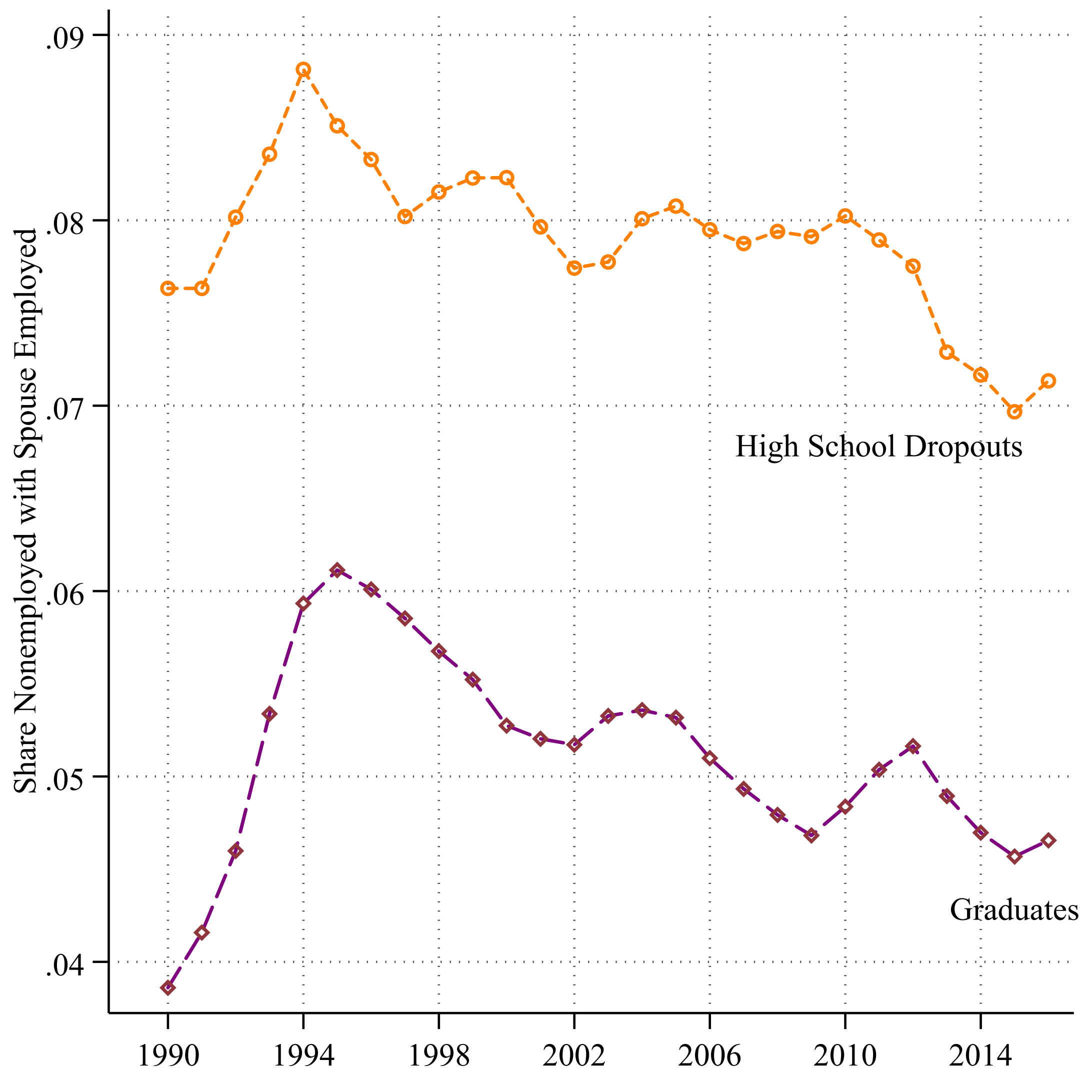

Figure 2.7: Share of Nonemployed Workers With An Employed Spouse

Notes: The figure shows the share of married prime-age men that are in nonemployment with a spouse that is currently employed. These shares are derived from Labour Force Surveys. The trends end in 2016 because the post-2016 information does not allow consistent estimation of these shares.

low-educated men. The evidence in Fig. 2.7 suggests that this channel is less likely to be the cause because the changes in the share of men who are nonemployed with employed partners are relatively similar for both groups.

Fig. A-11 shows significant shifts in the composition of nonemployed workers. The share of nonemployed workers actively looking for work (unemployed) decreased from 45% and 50% to 27% and 45% for less-educated and high-educated men, respectively. While the share of high-educated men fell by only 5%, that of less-educated men declined by 18%. These trends indicate that most nonemployed men are out of the labour force. Whether the reluctance to actively search is due to being discouraged after failing to find work or is driven by a conscious decision not to work is less clear. The LFS data do not allow checking whether these shifts are due to pull or push factors. If it is the latter, more aggressive measures would be needed to address these shifts because an inability to find a job when there is a desire to work can lead to other social ills. However, the overall evidence suggests that the shifts are driven by a conscious decision not to work, as the evidence in shifts in preference for leisure and generosity of government insurance programs above shows.

2.7 Robustness Checks

The results in Sections 2.6.2.1-2.6.2.2 are from a model that includes job search and duration dependence of job-finding rates. The conclusion in that section may differ from Wolcott (2021), who finds labour demand as the main driver of the rise in employment inequality, because of these new two additions to the model. To confirm this is not the case, I re-calibrated the model without variable search effort and duration dependence of job-finding rates. The calibrated model and counterfactual results are in Fig. 2.8. The calibrated model matches the employment gap in the data well, and job separations and labour supply are still the main drivers of the rise in employment inequality, as in Fig. 2.4. However, there is a difference in the role of labour market efficiency. Compared to Fig. 2.4,

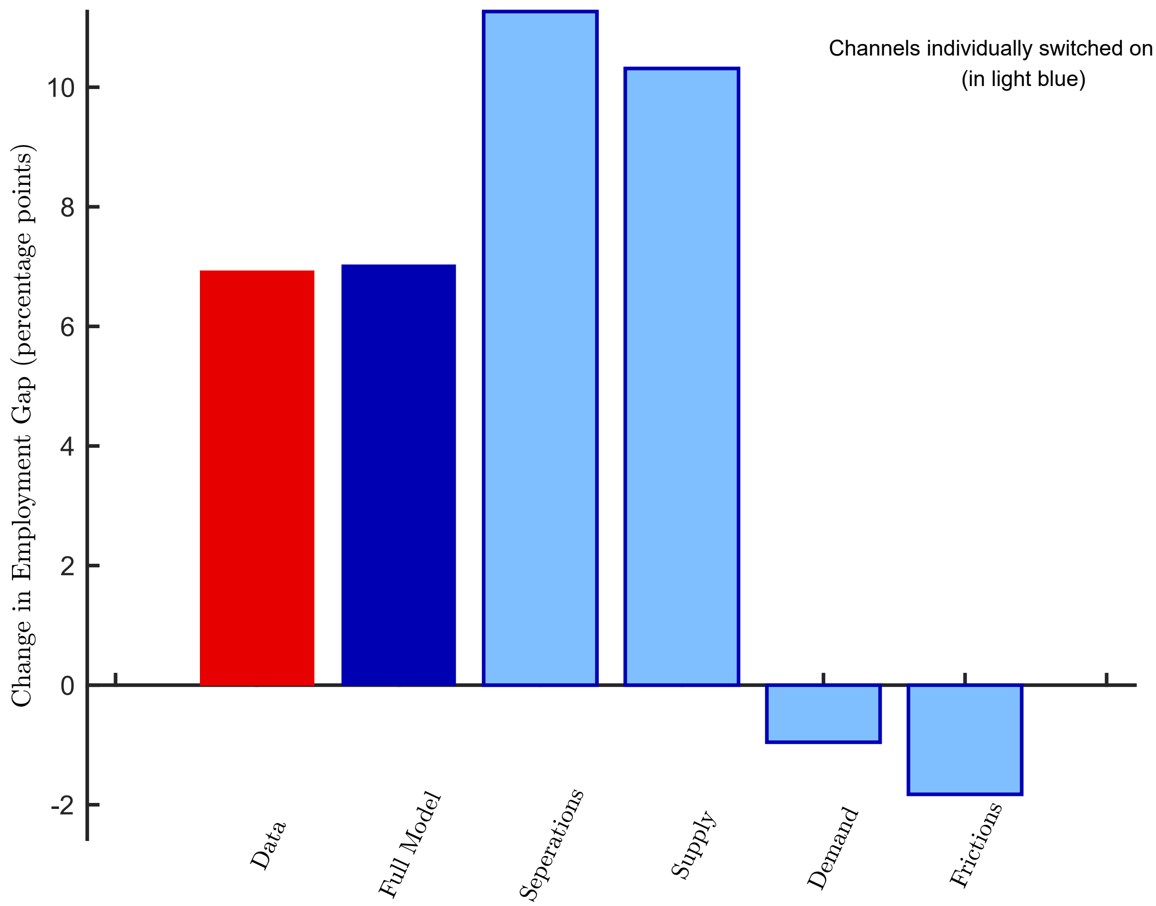

Figure 2.8: Employment Inequality: Model without Job Search and Duration Dependence

Notes: The red bar shows the rise in the gap of employment rates (in percentage points) between high school dropouts and graduates between 1990 and 2019 in the data. The dark blue bar is the corresponding gap in the model when all channels are allowed to change between the two periods, whereas the remaining light blue bars captures the results from the counterfactual experiments where only one channel is allowed to change between 1990 and 2019.

the model underestimates the role of improvements in search frictions in moderating the rise in the employment gap.

Adding search effort to the model introduces the ability of workers to influence their job-finding rates by choosing how much effort to exert in a job search. For instance, if workers decide to search for work more intensely, it increases the rate at which firms match with workers. This increased rate translates into an overall improvement in matching efficiency, equivalent to reducing job search frictions in the model. The fact that job-finding rates adjust depending on the composition of nonemployment duration also helps to amplify the role of matching efficiency in the model. These results further support the inclusion of job search effort and duration dependence of job-finding rates in the model.

2.8 Conclusion

The employment rates of prime-age men have been falling over time in Canada. Although this decline occurred for all men, the fall is particularly steep for low-educated men (less than a high school diploma). The difference in the rate of decline entails a widening gap in employment rates over time, which I define as a rise in employment inequality. This chapter investigates the potential reasons for this rise in employment inequality. I build a search and matching model with search frictions that integrate the role of five endogenous channels: labour demand, labour supply, matching efficiency, nonemployment duration dependence of job-finding rates and job search effort, to examine the rise in the employment gap between 1990 and 2019. The model also includes an exogenous job separation rate channel. In the model, workers are heterogeneous in ability and decide the degree of search effort to exert when nonemployed, and firms post ability-specific vacancies to find a match with a desire to produce output for a profit. Employment inequality rises due to shifts in these channels that differ across education groups. The results indicate that shifts in labour supply (endogenous) and separation rates (exogenous) are the primary cause of the rise in employment inequality. This labour supply effect is consistent with the data that indicates a widening preference for leisure between low-educated and high-educated men. The Canadian data also shows that government insurance programs, such as employment insurance, means-tested benefits, disability insurance and the tax systems, have become increasingly more generous to less-educated men. The results also indicate that improvements in job search technology have helped to improve matching efficiency, which in turn helped to moderate the rise in employment inequality. Although by smaller margins, the other channels contributed to reducing the widening of the employment gap.

Chapter 3

A Guide to Estimating the Canonical Income Process in Quasidifferences

with Dmytro Hryshko

3.1 Introduction

Answering many questions in macro and labor economics requires modeling of the individual and household income dynamics. It is common to posit a model of log-idiosyncratic income as a sum of permanent and transitory components and a fixed effect. Due to its relative simplicity and wide applicability, such an income process is often labeled as the canonical income process; e.g., Arellano et al. (2017).

There are various approaches to estimating the income process parameters. The literature typically relies on a minimum-distance estimation and autocovariance moments of log-income levels, y_{it}, or differences, $\Delta y_{it} = y_{it} - y_{it-1}$; see, e.g., Daly et al. (2021) for a discussion. The objective of this chapter is to provide a guide to estimating the canonical income process using quasidifferences defined as $\tilde{\Delta} y_{it} = y_{it} - \rho y_{it-1}$, where ρ is the persistence of shocks to the permanent component.

Estimations in levels, differences and quasidifferences have their advantages and disadvantages. Estimation in levels can recover the variance of fixed effects but requires a stand on the distribution of fixed effects and initial conditions and can be cumbersome to implement if the shocks' variances change with time and age. Estimation in differences is preferable when the permanent component is a random walk since fixed effects and initial conditions get differenced out and do not affect the estimated parameters. However, even though there is a consensus that the shocks to the permanent component are persistent, the literature is far from the agreement that these shocks are fully permanent. If one departs from assuming that the permanent component is a random walk, estimation in differences loses its advantages relative to estimation in levels. Estimation in quasidifferences inherits the advantages of estimation in levels and differences. Like estimation in differences when the permanent component is a random walk, it is not dependent on the distribution of initial conditions and easy in implementation. In quasidifferences, this is true even if the shocks to the persistent component are not fully permanent. Like estimation in levels, it allows for estimation of the variance of fixed effects. The challenging aspect of estimation in quasidifferences, however, is the requirement of an estimate of the persistence. In principle, estimation could proceed in two stages: first, estimation of the persistence, and second, estimation of the other parameters of the income process. Persistence could be estimated by GMM, which, like estimation in levels, requires modeling of the initial conditions.[1] To overcome this challenge, Blundell et al. (2015) proposed recently to jointly estimate the persistence and the other income process parameters using quasidifferences. The procedure they suggest involves evaluation of the distance between the model and the data autocovariance moments for a predefined grid in persistence and choosing the set of the income process parameters that minimizes the distance as the model estimates. Although this procedure is appealing in its simplicity and its merits relative to the estimations in levels and differences, nothing is known about its performance in realistic settings. Using Monte Carlo simulations, the chapter is the first to conduct such an analysis and examine the biases in the estimated parameters using quasidifferences for various true values of the persistence, N, T, initial conditions and weighting schemes.

Following a large GMM literature,[2] our primary focus is on evaluating the biases in the estimates of persistence, but we also catalog the

[1] Moreover, estimates of the persistence could be biased in small samples; e.g., Chen et al. (2019).

[2] See, e.g., Bun and Sarafidis (2015) for a review.

results for the other parameters in the appendix. Precise estimation of the persistence is important for the design of the optimal tax policy as highlighted in Farhi and Werning (2012) and for the interpretation of transmission coefficients for permanent income shocks to household consumption as emphasized in Blundell (2014), Hryshko and Manovskii (2022) and Bryukhanov and Hryshko (2020).[3] Blundell et al. (2015) measure the persistence of permanent shocks to male earnings, male and family disposable income to evaluate the insurance role of the tax and transfer system and family labor supply in moderating the persistence of longer-lasting shocks to male earnings. We find that equally-weighted estimation often results in downward-biased estimates of the persistence in small and bigger samples regardless of the magnitude of true persistence.

The bias is bigger, *ceteris paribus*, when the variance of fixed effects is smaller, the variance of permanent shocks is smaller, or the variance of transitory shocks is higher. Only when the variance of permanent shocks is bigger than the variance of transitory shocks, equally-weighted estimation is reliable in recovering the true values of the persistence. This is the setting of Blundell et al. (2015) who used equally-weighted estimation and very big samples from administrative Norwegian data on earnings and incomes. Optimal and diagonal weighting produce nontrivial upward biases in the estimated persistence when the true persistence is low and the variance of incomes in the first sample year is nonnegligible—this may happen either when the average age of individuals in the first sample year is sufficiently large (due to accumulation of persistent shocks) or when the vari-

[3] Transmission coefficients for permanent and transitory income shocks to household consumption as measured by the methodology of Blundell et al. (2008) have been used recently as a metric of the performance of quantitative macro models; see, e.g., De Nardi et al. (2019) and Daly et al. (2021). These coefficients and the income process parameters can also be estimated using quasidifferences when the income process allows for persistent rather than fully permanent shocks. See Section 3.5.4 for implementation.

ance of initial conditions is nonnegligible. Similar to the equal weighting, optimal and diagonal weighting produce unbiased estimates of the persistence when the variance of permanent shocks is bigger than the variance of transitory shocks but only when the variance of the permanent component in the first sample year is small (e.g., when following a cohort of individuals from the start of their working careers).

Other applications require precise estimates of the variance of permanent and transitory shocks and fixed effects. The variance of permanent shocks, e.g., is an important driver of wealth accumulation in a buffer-stock model of saving (Carroll, 1997), whereas the variance of fixed effects determines if inequality is important for the determination of the aggregate demand (Auclert and Rognlie, 2018). Using quasidifferences, we find that the variance of fixed effects is severely upward-biased when the true persistence is high, but the biases in its estimation are small when the true persistence is low. The variances of permanent and transitory shocks are estimated with small biases, especially so when the number of sample individuals is large.

In an empirical application, we use administrative data on earnings for a large sample of Danish males born in 1952 observed during the 1984–2008 period to estimate the canonical income process using quasidifferences. We find that the variance of transitory shocks is about twice as large as the variance of permanent shocks and that optimal weighting of the moments yields a relatively high estimate of the persistence followed by smaller estimates when using a diagonal and equal weighting of the moments. We next estimate the income process using small bootstrap samples out of the original sample of Danish males. Optimal weighting produces a relatively stable estimate of the persistence in big and small samples, whereas diagonal weighting results in a substantially higher estimate of the persistence in small samples, and equal weighting yields a smaller estimate in small samples. These results agree with our results for big and small samples from simulated data generated for high persistence of permanent shocks when the variance of transitory shocks is higher than the variance of permanent shocks and the variance of the permanent component in the initial sample year is nonnegligible (due to a history of accumulation of persistent shocks or nonzero initial conditions).

Estimation in quasidifferences is amenable to the methodology of Blundell et al. (2008) that aims to recover, besides the income process parameters, the transmission of permanent and transitory shocks to household consumption, that is, the fraction of permanent and transitory shocks that are absorbed by consumption. It generalizes the original methodology due to its potential to recover the persistence of permanent shocks that may differ from unity and the variance of fixed effects in household incomes. To explore the effectiveness of estimation in quasidifferences using consumption and income data, we calibrate a standard incomplete-markets model for low and high values of persistence and simulate the data replicating the age structure of data from the Panel Study of Income Dynamics (PSID) used in Hryshko and Manovskii (2022). We find significant biases in the estimated persistence and the transmission coefficients for permanent and transitory shocks for high and low values of persistence, different weighting schemes and the number of households in the simulated data.

The results in this chapter warn against the routine use of estimation in quasidifferences despite its attractive features. However, there is one case when estimation in quasidifferences reliably recovers the persistence of permanent shocks. When the variance of permanent shocks is higher than the variance of transitory shocks, equally-weighted estimation recovers the true persistence of permanent shocks for different T, N and initial conditions. Since biases in the variances of the shocks using estimation in quasidifferencs are small, our advice is to estimate the income process using equal weighting of the moments and use the results with confidence if the estimated variance of permanent shocks is higher than the variance of transitory shocks.

The rest of the chapter is structured as follows. In Section 3.2, we present theoretical moments in levels, differences, and quasidifferences for the canonical income process and discuss identification and estimation of the income-process parameters. In Section 3.3, we present the details of our Monte Carlo simulations, and in Section 3.4 we present the results of estimations in quasidifferences using the simulated data. Section 3.5 analyzes the bias in the estimated persistence using simulated income data, presents the results from an empirical application using administrative data on earnings from Denmark, and analyzes the bias in the income- and consumption-process parameters using data from a calibrated lifecycle model of consumption. Section 3.6 discusses implications of our results for practical use, and Section 3.7 concludes.

3.2 The Canonical Income Process

We consider the canonical decomposition of idiosyncratic (log-)income for individual i at time t, y_{it}, into a sum of an autoregressive permanent component z_{it}, fixed effect α_i and transitory shock ϵ_{it}:

$$
\begin{aligned}
y_{it} &= \alpha_i + z_{it} + \epsilon_{it} \\
z_{it} &= \rho z_{it-1} + \eta_{it} \qquad (3.1)\\
\eta_{i,t} \sim iid(0, \sigma^2_\eta) \quad & \epsilon_{i,t} \sim iid(0, \sigma^2_\epsilon) \quad \alpha_i \sim iid(0, \sigma^2_\alpha).
\end{aligned}
$$

Shocks η_{it} have persistence ρ (we will refer to them below as persistent or permanent interchangeably);[4] all shocks and the fixed effect are i.i.d. and are orthogonal to each other. In quantitative macro, it is common to assume that the initial condition z_{i0} takes the value of zero for everyone, whereas in the panel-data literature, the initial condition is commonly assumed to be drawn from some distribution. We will consider both of these assumptions in our Monte Carlo simulations.

The parameter ρ captures the degree of persistence in the autoregressive component. Under this structure α_i captures the individual fixed effects that ex-ante differentiates individuals before they join the labour market. The persistent shocks $\eta_{i,t}$ are independent and identically distributed, and orthogonal to α_i and $\epsilon_{i,t}$. Given the canonical income process above, the focus is to identify a set of parameters $\Phi = \{\rho, \sigma^2_\eta, \sigma^2_\alpha, \sigma^2_\epsilon\}$. To identify these parameters, we match data and theoretical moments of income growth computed from quasidifferences $\tilde{\Delta} y_{i,t} = y_{i,t} - \rho y_{i,t-1}$.

3.2.1 Theoretical Moments

Assuming that initial conditions z_{i0} are i.i.d. and orthogonal to the shocks and fixed effects, the autocovariance function of incomes in levels, differences, and quasidifferences for a cohort of individuals who start their work-

[4] When $\rho < 1$, the shocks are permanent in the sense that they do not fully die out, but, obviously, they do not affect current and future incomes with the same magnitude as is the case when $\rho = 1$.

ing life at time 0 can be characterized as follows.

$$
\begin{aligned}
&\texttt{Levels:}\\
&E[y_{it}y_{it}] = \sigma_\alpha^2 + \rho^{2t}\mathbf{var}(z_{i0}) + \frac{1-\rho^{2t}}{1-\rho^2}\sigma_\eta^2 + \sigma_\epsilon^2, \quad t = 1, 2, \ldots, T \qquad (3.2)\\
&E[y_{it+k}y_{it}] = \sigma_\alpha^2 + \rho^k\left[\rho^{2t}\mathbf{var}(z_{i0}) + \frac{1-\rho^{2t}}{1-\rho^2}\sigma_\eta^2\right], \quad t = 1, 2, \ldots, T-k;\\
&\qquad 1 \le k \le T-t. \qquad (3.3)
\end{aligned}
$$

$$
\begin{aligned}
&\texttt{Differences:}\\
&E[\Delta y_{it}\Delta y_{it}] = \rho^{2(t-1)}(1-\rho)^2\mathbf{var}(z_{i0}) + \left[1 + \frac{1-\rho}{1+\rho}\left(1-\rho^{2(t-1)}\right)\right]\sigma_\eta^2 + 2\sigma_\epsilon^2,\\
&\qquad t = 2, \ldots, T \qquad (3.4)\\
&E[\Delta y_{it+1}\Delta y_{it}] = \rho^{2t-1}(1-\rho)^2\mathbf{var}(z_{i0}) + \left[-(1-\rho) + \rho\frac{1-\rho}{1+\rho}\left(1-\rho^{2(t-2)}\right)\right]\sigma_\eta^2 - \sigma_\epsilon^2,\\
&\qquad t = 2, \ldots, T-1 \qquad (3.5)\\
&E[\Delta y_{it+k}\Delta y_{it}] = \rho^{k-1}\left[\rho^{2t-1}(1-\rho)^2\mathbf{var}(z_{i0}) - (1-\rho) + \rho\frac{1-\rho}{1+\rho}\left(1-\rho^{2(t-2)}\right)\right]\sigma_\eta^2,\\
&\qquad t = 2, \ldots, T-k; \qquad 2 \le k \le T-t. \qquad (3.6)
\end{aligned}
$$

$$
\begin{aligned}
&\texttt{Quasidifferences:}\\
&E[\Delta\tilde{y}_{it}\Delta\tilde{y}_{it}] = (1-\rho)^2\sigma_\alpha^2 + \sigma_\eta^2 + (1-\rho)^2\sigma_\epsilon^2, \quad t = 2, \ldots, T \qquad (3.7)\\
&E[\Delta\tilde{y}_{it+1}\Delta\tilde{y}_{it}] = (1-\rho)^2\sigma_\alpha^2 - \rho\sigma_\epsilon^2, \qquad t = 2, \ldots, T-1 \qquad (3.8)\\
&E[\Delta\tilde{y}_{it+k}\Delta\tilde{y}_{it}] = (1-\rho)^2\sigma_\alpha^2, \qquad t = 2, \ldots, T-k; \quad 2 \le k \le T-t. \qquad (3.9)
\end{aligned}
$$

3.2.2 Identification and Estimation

Identification of the income process parameters in Eq. (3.1) is typically established by using the moments involving combinations of autocovariances of income in levels or differences, whereas estimation is commonly performed relying on the minimum-distance method that involves matching all available autocovariance moments of incomes in levels or differences to their theoretical counterparts; see, e.g., Meghir and Pistaferri (2004), Heathcote et al. (2010) and Daly et al. (2021).

There are advantages and disadvantages of estimations relying on the moments in levels or differences. Estimation in differences is the most obvious choice if the permanent component is a random walk, $\rho = 1$, as it is valid for various distributions of the initial conditions for the permanent component and fixed effects. However, it does not allow for the identification of the variance of fixed effects. Estimation in levels, in turn, requires making a stand on the distribution of initial conditions for the permanent component even if the persistence of permanent shocks is equal to one. Our focus is on recovering the model parameters using quasidifferences defined as $\tilde{\Delta} y_{it} = y_{it} - \rho y_{it-1}$. This method inherits the simplicity of estimation in differences when the permanent component is a random walk, allows for estimation of the variance of fixed effects, but requires an estimate of the persistence, ρ. Eqs. (3.7)–(3.9) list three unique moments and four unknown parameters—fixing ρ allows identification of the variance of fixed effects and the variances of permanent and transitory shocks.

How would one obtain a reliable estimate of the persistence? One possibility is to use GMM. Similar to the minimum-distance estimation in levels, GMM requires modeling initial conditions. It can also be substantially biased in small samples when ρ is close to one; see, e.g., Gouriéroux et al. (2010). Another possibility is to use indirect inference, but, like GMM, it does not allow for identification of the model parameters apart from the persistence; Gouriéroux et al. (2010). These challenges could be potentially overcome if ρ is estimated simultaneously with the other parameters when using quasidifferences as was recently suggested by Blundell et al. (2015). This approach involves minimizing a distance between the full set of data and model autocovariance moments for $\tilde{\Delta} y_{it}$ for a prespecified grid of values of ρ. One then chooses $\hat{\rho}$ and the corresponding set of the model parameters yielding the lowest distance across all grid values of ρ as estimates of the income process parameters for a given sample.

3.3 Monte Carlo Setup

In this section, we will analyze how well estimation in quasidifferences recovers the true persistence in various Monte Carlo settings.

3.3.1 Time Dimension

To examine the influence of the time dimension on the identification of the income process parameters, we consider earnings histories simulated for fifteen and thirty periods. Survey data with a shorter time span of fifteen periods and less was used, e.g., in Abowd and Card (1989) and more recently in Blundell et al. (2008). Recent literature uses administrative data that allows for longer samples spanning more than twenty years of individuals' life cycles; see, e.g., Blundell et al. (2015), Daly et al. (2021) and Guvenen et al. (2021).

3.3.2 Sample Size

We consider two sample sizes, a small and a bigger one, with the number of panel individuals equal to 1,000 and 10,000, respectively. Small sample sizes are not uncommon in research relying on survey data, whereas bigger sample sizes are encountered in the literature utilizing administrative data. The low number for simulated panel units is motivated by a recent paper of Hryshko and Manovskii (2022) who considered samples of around 1,000 households and lower. We settled on 10,000 for our bigger sample size for computational reasons. This number is in the ballpark of sample sizes from

German administrative data used by Daly et al. (2021) and in our empirical application in Section 3.5.3.

3.3.3 Initial Conditions

To examine the influence of initial conditions, we proceed as follows. In the first experiment, we make a departure from the assumption of zero permanent component for every individual in period zero. In the second experiment, we maintain that assumption, simulate the income process for thirty periods, and estimate it using simulated data for the last fifteen periods. In both of the experiments, initial periods are characterized by the variances of the permanent component significantly different from zero.

3.3.4 Parameter Values

We focus on two values for the persistence of permanent shocks, $\rho = 0.90$ and $\rho = 0.995$. The higher value characterizes the permanent component as a near-unit root process, which is frequently studied in the literature. The lower value was found to characterize the persistence of permanent shocks for about half of the PSID families formed after 1968, the year of the dataset's inception; see Hryshko and Manovskii (2022).[5] Our benchmark values for the variance of fixed effects, variance of permanent shocks, and variance of transitory shocks are standard and equal to 0.10, 0.01, and 0.04, respectively; see, e.g., Meghir and Pistaferri (2004), Storesletten et al. (2004), Hryshko (2012), and Guvenen et al. (2021). In separate experiments, holding the other parameters fixed at their benchmark values, we consider a higher variance of permanent shocks at 0.04, a lower variance of transitory shocks at 0.01, and equally low and high variances of permanent

[5] We also considered an intermediate value $\rho = 0.95$ but chose not to tabulate the results to avoid cluttering.

and transitory shocks, simultaneously set to 0.01 and 0.04, respectively. We have also analyzed a case with a lower variance of fixed effects equal to 0.05, as estimated in Guvenen (2009). We assume that all shocks and fixed effects are drawn from normal distributions.[6]

3.3.5 Implementation

As mentioned above, one requires a prespecified grid for persistence, ρ, to implement estimation in quasidifferences. The grid should be wide, with many points, while allowing for a manageable estimation time. After some experimentation, we settled on a piecewise linear grid with 100 points ranging from 0.3 to 1.4.[7]. The grid is split into four parts corresponding to ranges [0.3, 0.5], (0.5,0.7], (0.7,1.1], and (1.1,1.4] with 5, 10, 70 and 15 grid points (equidistant within each range), respectively. For the two values of true persistence and a given set of the other parameters, we simulate one hundred samples. For each of those samples, we perform one hundred minimum-distance estimations (for each grid value) and set the grid value yielding the minimum objective value as an estimate of the persistence. Objective values are calculated as a weighted distance between the data and theoretical autocovariance moments. We use three popular weighting schemes in estimations—equal weighting of the moments; diagonal weighting, with the diagonal elements of the weight matrix containing the diagonal of the inverse of the variance-covariance matrix of the data moments and zero off-diagonal elements; and optimal weighting, with the weight matrix being the inverse of the variance-covariance matrix of the

[6]We have also experimented with heavy-tailed distributions, but the results were similar.

[7]Although we do not consider income processes with a true persistence of permanent shocks above one, it is important to allow for an estimated persistence to exceed one; see, e.g., Andrews (1993) and Gustavsson and Österholm (2014). If this was not allowed, the estimated persistence will be surely biased downward when the true persistence is very close or equal to one.

data moments.

3.4 Results

In this section, we present our results from Monte Carlo simulations for different weighting schemes, N, T and initial conditions for the permanent component.

3.4.1 Equal Weighting

3.4.1.1 Low Persistence

Table 3.1 contains the results for low persistence and no variation in the permanent component at time zero. Panel A contains the results for samples with a small number of individuals, while Panel B contains the results for larger samples.

Row (1) focuses on the benchmark parameters for the variances of fixed effects, permanent and transitory shocks. Column (1) shows the results for estimations utilizing the first thirty periods of simulated data. There is a sizeable downward bias in the estimated persistence that equals about 0.10. Row (2) keeps the variance of fixed effects and transitory shocks fixed but increases the variance of permanent shocks to 0.04. The estimated persistence is, on average, 0.867, featuring a smaller downward bias relative to that in row (1). Row (3) conducts a comparative statics exercise of lowering the variance of transitory shocks from 0.04 in row (1) to 0.01. The estimated persistence is now 0.898, so that the bias nearly vanishes. Row (4) conducts a comparative statics exercise of increasing the variance of permanent shocks from 0.01 in row (3) to 0.04 while keeping the variance of transitory shocks at a low value of 0.01. The estimated persistence is somewhat higher than in row (3) but close to the true value.

Summing up, it appears that increasing the variance of transitory shocks raises a downward bias in the estimated persistence, more so when the variance of permanent shocks is low. The downward bias becomes smaller with an increase in the variance of permanent shocks. Both of these effects indicate that a higher ratio of transitory to permanent variances results in a bigger downward bias in the estimated persistence. Row (5) conducts a comparative statics exercise of lowering the variance of fixed effects to 0.05 from 0.10 in row (1) and shows that the downward bias becomes bigger when the variance of fixed effects is smaller. In column (2) we use the first fifteen years of the simulated data whereas in column (3) we use its last fifteen years. The results are qualitatively similar—with generally bigger downward biases than in column (1)—and, as expected, are less precise.

In Panel B, rows (1)–(5), we conduct the same analysis but now for samples that are ten times bigger cross-sectionally. A downward bias in the estimated persistence remains substantial when the ratio of the variance of transitory to permanent shocks is high, rows (1) and (5), but nearly vanishes for the other experiments in rows (2) to (4).

3.4.1.2 High Persistence

Table 3.2 repeats the analysis of Table 3.1 for the persistence of permanent shocks equal to 0.995. The results are similar overall, although, in addition, a substantive downward bias remains even when $N = 10,000$ and the variances of permanent and transitory shocks are both low, Panel B row (3).

Table 3.1: Estimated persistence. True persistence $\rho = 0.9$. Zero Initial Conditions

Weighting	Equal			Optimal			Diagonal		
$t =$	[1-30]	[1-15]	[16-30]	[1-30]	[1-15]	[16-30]	[1-30]	[1-15]	[16-30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Panel A: $N = 1000$								
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.7971	0.7733	0.7692	0.9071	0.9180	0.9953	0.7352	0.7382	1.0236
	(0.0230)	(0.0494)	(0.0486)	(0.0457)	(0.0886)	(0.1120)	(0.0190)	(0.1223)	(0.3415)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.8671	0.8463	0.8540	0.9020	0.8906	1.0076	0.8523	0.8319	1.0382
	(0.0113)	(0.0261)	(0.0209)	(0.0120)	(0.0267)	(0.1113)	(0.0190)	(0.1223)	(0.3415)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.8983	0.8919	0.8966	0.8998	0.9059	1.0222	0.8815	0.8813	1.0207
	(0.0146)	(0.0348)	(0.0330)	(0.0131)	(0.0370)	(0.1062)	(0.0127)	(0.0496)	(0.1469)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.8987	0.8936	0.8988	0.9003	0.8999	1.0319	0.8954	0.8897	1.0201
	(0.0078)	(0.0209)	(0.0178)	(0.0106)	(0.0222)	(0.1084)	(0.0074)	(0.0125)	(0.1229)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.7686	0.7391	0.7353	0.8985	0.8895	0.9955	0.7217	0.6956	0.9571
	(0.0170)	(0.0433)	(0.0391)	(0.0158)	(0.0442)	(0.1387)	(0.0661)	(0.0747)	(0.3592)
	Panel B: $N = 10000$								
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.8602	0.8361	0.8400	0.9002	0.9021	1.0013	0.8277	0.8091	1.0142
	(0.0078)	(0.0194)	(0.0196)	(0.0046)	(0.0191)	(0.1040)	(0.0078)	(0.0171)	(0.2218)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.8938	0.8887	0.8880	0.8999	0.9009	0.9904	0.8872	0.8753	1.0092
	(0.0035)	(0.0108)	(0.0109)	(0.0032)	(0.0099)	(0.1050)	(0.0033)	(0.0067)	(0.1367)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.8999	0.9011	0.9002	0.8997	0.9018	0.9882	0.8957	0.8894	1.0007
	(0.0049)	(0.0156)	(0.0136)	(0.0029)	(0.0121)	(0.1069)	(0.0041)	(0.0122)	(0.1222)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.9005	0.9006	0.9004	0.9002	0.8994	0.9809	0.8990	0.8969	0.9787
	(0.0026)	(0.0075)	(0.0075)	(0.0032)	(0.0077)	(0.1035)	(0.0027)	(0.0072)	(0.1070)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.8533	0.8332	0.8305	0.8997	0.9024	1.0008	0.8213	0.7981	1.0605
	(0.0070)	(0.0161)	(0.0191)	(0.0040)	(0.0179)	(0.1066)	(0.0059)	(0.0144)	(0.2354)

Notes: The table shows estimated persistence for various T, N and weighting schemes from simulated data for $\rho = 0.9$ and various values for the variances of fixed effects and shocks. σ_α^2, σ_η^2 and σ_ϵ^2 are the variances of fixed effects, persistent and transitory shocks, respectively. Standard errors are in parentheses.

3.4.2 Optimal Weighting

3.4.2.1 Low Persistence

Using an optimal weighting matrix and the first thirty periods of the simulated data produces nearly unbiased estimates of the persistence; see column (4) of Table 3.1. Using the first fifteen periods makes the estimates somewhat noisier, but the results are qualitatively similar. This is true both for samples with small and big N, panels A and B, respectively. Using the last fifteen periods produces a substantial upward bias in the estimated persistence. The key difference between the samples used in columns (4) and (5) versus column (6) is that, in the latter case, the variance of incomes in the first panel year is substantially higher due to the accumulation of persistent shocks. We will show below that the high variability of the persistent component in the first sample year is the driver of the divergent results for the first fifteen vs. the last fifteen periods.

3.4.2.2 High Persistence

For the case of high persistence, using the first fifteen periods in estimation produces substantively downward-biased estimates of the persistence even if the number of sample individuals is large when the ratio of the variance of transitory to permanent shocks is large or when both permanent and transitory variances are small—rows (1) and (5), and row (3) in Table 3.2, respectively. Similar to the case of low persistence, optimal weighting typically produces upward-biased estimates of the persistence when the number of sample individuals is small, and one uses the last fifteen sample periods. However, the bias is small in magnitude.

Table 3.2: Estimated persistence. True persistence $\rho = 0.995$. Zero Initial Conditions

Weighting	Equal			Optimal			Diagonal		
$t =$	[1-30]	[1-15]	[16-30]	[1-30]	[1-15]	[16-30]	[1-30]	[1-15]	[16-30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Panel A: $N = 1000$								
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.9170	0.8739	0.8951	0.9836	0.9899	0.9996	0.9222	0.9317	1.0203
	(0.0095)	(0.0237)	(0.0343)	(0.0142)	(0.0482)	(0.0535)	(0.0838)	(0.1622)	(0.1698)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.9862	0.9760	0.9875	0.9949	0.9958	1.0148	0.9784	0.9499	1.0819
	(0.0040)	(0.0245)	(0.0065)	(0.0030)	(0.0056)	(0.0283)	(0.0838)	(0.1622)	(0.1698)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.9659	0.9516	0.9578	0.9827	0.9839	0.9950	0.9580	0.9403	0.9969
	(0.0079)	(0.0212)	(0.0184)	(0.0127)	(0.0344)	(0.0434)	(0.0057)	(0.0365)	(0.0744)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.9935	0.9932	0.9935	0.9945	0.9942	1.0106	0.9923	0.9913	1.0206
	(0.0030)	(0.0051)	(0.0041)	(0.0030)	(0.0105)	(0.0264)	(0.0032)	(0.0059)	(0.0347)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.9127	0.8690	0.8887	0.9838	0.9750	1.0037	0.8892	0.8811	1.0211
	(0.0089)	(0.0188)	(0.0280)	(0.0134)	(0.0274)	(0.0486)	(0.0267)	(0.1275)	(0.1796)
	Panel B: $N = 10000$								
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.9536	0.9293	0.9332	0.9903	0.9832	0.9906	0.9392	0.9070	0.9173
	(0.0040)	(0.0115)	(0.0123)	(0.0051)	(0.0134)	(0.0266)	(0.0037)	(0.0076)	(0.0466)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.9929	0.9931	0.9934	0.9951	0.9947	0.9965	0.9914	0.9917	1.0077
	(0.0025)	(0.0026)	(0.0029)	(0.0028)	(0.0028)	(0.0093)	(0.0000)	(0.0021)	(0.0287)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.9783	0.9693	0.9710	0.9899	0.9853	0.9896	0.9735	0.9578	0.9638
	(0.0038)	(0.0086)	(0.0124)	(0.0049)	(0.0105)	(0.0189)	(0.0021)	(0.0058)	(0.0264)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.9953	0.9943	0.9942	0.9953	0.9949	0.9947	0.9946	0.9946	0.9953
	(0.0027)	(0.0029)	(0.0029)	(0.0027)	(0.0028)	(0.0028)	(0.0029)	(0.0029)	(0.0074)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.9534	0.9292	0.9289	0.9898	0.9833	0.9900	0.9384	0.9050	0.9106
	(0.0037)	(0.0104)	(0.0119)	(0.0047)	(0.0131)	(0.0256)	(0.0033)	(0.0068)	(0.0390)

Notes: The table shows estimated persistence for various T, N and weighting schemes from simulated data for $\rho = 0.995$ and various values for the variances of fixed effects and shocks. σ_α^2, σ_η^2 and σ_ϵ^2 are the variances of fixed effects, persistent and transitory shocks, respectively. Standard errors are in parentheses.

3.4.3 Diagonal Weighting

Diagonal weighting produces the results similar to equal weighting when i) the true persistence is high and the number of individuals is large; ii) when the true persistence is high, the number of individuals is low and one uses the first fifteen or the first thirty sample periods; or iii) when the true persistence is low and one uses the first fifteen or the first thirty sample periods, regardless of the number of sample individuals. Diagonal weighting produces results similar to optimal weighting when i) the true persistence is low and one uses the last fifteen sample periods, or ii) when the true persistence is high and one uses the last fifteen sample periods when the number of sample individuals is small.

3.4.4 Nonzero Initial Conditions

In period 16, the variance of the permanent component equals $\sigma_\eta^2(1-\rho^{32})/(1-\rho^2)$ when the initial variance of the permanent component is zero—see Eq. (3.2). To verify our conjecture above that it is the variance of the permanent component in the first sample year that makes our diagonal and optimal weighting results for periods $t = 1, \ldots, 15$ and $t = 16, \ldots, 30$ diverge, we conduct separate experiments where we assume that var$(z_{i0}) = \sigma_\eta^2(1-\rho^{30})/(1-\rho^2)$. We then simulate data for fifteen periods and compare the results with those in Tables 3.1–3.2 that are based on zero initial conditions and the last fifteen periods. By design, income variances in the first sample year in these experiments are identical.[8] For convenience, columns (4)–(6) and (10)–(12) of Table 3.3 reproduce our results from Tables 3.1–3.2 for low and high persistence, respectively. Reassuringly, our

[8]This is so because the variance of the permanent component in the first year for the case with nonzero initial conditions equals $\rho^2\text{var}(z_{i0}) + \sigma_\eta^2 = \rho^2\sigma_\eta^2(1-\rho^{30})/(1-\rho^2) + \sigma_\eta^2 = \sigma_\eta^2(1-\rho^{32})/(1-\rho^2)$.

new results in columns (1)–(3) and (7)–(9) of Table 3.3 based on nonzero initial conditions are nearly identical to the corresponding results in the reproduced columns, both in small and large samples. For example, using the first fifteen periods of the samples with nonzero initial conditions and diagonally-weighted estimation produces a substantial upward bias in the persistence when the true persistence is low—column (3) Panels A and B of Table 3.3. Those results are nearly the same in all of our experiments when we use the last fifteen periods of the data with zero initial conditions and diagonally-weighted estimation—column (6) of Panels A and B.

An important takeaway is that even when using big samples that mix various cohorts of individuals at different stages of their life cycles, with optimal and diagonal weighting, one should expect an upward-biased estimate of persistence.

3.4.5 The Mechanics of an Upward Bias in Optimally- and Diagonally-weighted Estimations

As we documented above in Tables 3.1–3.2, there is a stark difference in the estimated persistence for equal versus optimal and diagonal weighting when the true persistence is low and one relies on the first fifteen (or first thirty periods) as opposed to the last fifteen periods. Since the number of individuals and time periods is held constant across those two scenarios, the key difference between them is the variance of permanent incomes in the initial period of the data—no variance of the permanent component for the samples utilizing the first fifteen periods in estimation, as we assume that initial conditions are zero for everyone, and a relatively high variance of the permanent component for the samples utilizing the last fifteen periods, due to a history of accumulation of persistent shocks.

To further explore the importance of this mechanism for the results,

Table 3.3: Estimated persistence. Nonzero vs. Zero Initial Conditions

Persistence:	$\rho = 0.9$						$\rho = 0.995$					
Initial Conditions	Nonzero			Zero			Nonzero			Zero		
$t =$	[1-15]			[16-30]			[1-15]			[16-30]		
Weighting	E	O	D	E	O	D	E	O	D	E	O	D
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Panel A: $N = 1000$											
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.7650	1.0180	1.0298	0.7692	0.9953	1.0236	0.8933	0.9928	0.9959	0.8951	0.9996	1.0203
	(0.0551)	(0.1244)	(0.3432)	(0.0486)	(0.1120)	(0.3415)	(0.0292)	(0.0523)	(0.1662)	(0.0343)	(0.0535)	(0.1698)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.8545	1.0094	1.0027	0.8540	1.0076	1.0382	0.9863	1.0165	1.0764	0.9875	1.0148	1.0819
	(0.0265)	(0.1135)	(0.1916)	(0.0209)	(0.1113)	(0.1959)	(0.0071)	(0.0288)	(0.0429)	(0.0065)	(0.0283)	(0.0371)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.8986	1.0006	1.0187	0.8966	1.0222	1.0207	0.9559	0.9915	0.9847	0.9578	0.9950	0.9969
	(0.0302)	(0.1041)	(0.1481)	(0.0330)	(0.1062)	(0.1469)	(0.0213)	(0.0423)	(0.0734)	(0.0184)	(0.0434)	(0.0744)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.9014	1.0104	1.0353	0.8988	1.0319	1.0201	0.9942	1.0108	1.0195	0.9935	1.0106	1.0206
	(0.0172)	(0.1061)	(0.1192)	(0.0178)	(0.1084)	(0.1229)	(0.0067)	(0.0270)	(0.0339)	(0.0041)	(0.0264)	(0.0347)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.7344	1.0130	0.9513	0.7353	0.9955	0.9571	0.8851	1.0048	1.0273	0.8887	1.0037	1.0211
	(0.0434)	(0.1221)	(0.3514)	(0.0158)	(0.1387)	(0.3592)	(0.0331)	(0.0558)	(0.1824)	(0.0486)	(0.0486)	(0.1796)
	Panel B: $N = 10000$											
(1) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.8365	1.0083	1.0216	0.8400	1.0013	1.0142	0.9315	0.9902	0.9241	0.9332	0.9906	0.9173
	(0.0223)	(0.1057)	(0.2243)	(0.0196)	(0.1040)	(0.2218)	(0.0130)	(0.0252)	(0.0587)	(0.0123)	(0.0266)	(0.0466)
(2) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.04$	0.8869	0.9874	0.9929	0.8880	0.9904	1.0092	0.9933	0.9954	1.0111	0.9934	0.9965	1.0077
	(0.0088)	(0.1037)	(0.1350)	(0.0109)	(0.1050)	(0.1367)	(0.0029)	(0.0061)	(0.0301)	(0.0029)	(0.0093)	(0.0287)
(3) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.01$	0.9015	0.9766	0.9839	0.9002	0.9882	1.0007	0.9713	0.9885	0.9599	0.9710	0.9896	0.9638
	(0.0176)	(0.1003)	(0.1181)	(0.0136)	(0.1069)	(0.1222)	(0.0120)	(0.0183)	(0.0201)	(0.0124)	(0.0189)	(0.0264)
(4) $\sigma_\alpha^2 = 0.10, \sigma_\epsilon^2 = 0.01, \sigma_\eta^2 = 0.04$	0.8999	0.9975	0.9907	0.9004	0.9809	0.9787	0.9946	0.9956	0.9951	0.9942	0.9947	0.9953
	(0.0077)	(0.1064)	(0.1098)	(0.0075)	(0.1035)	(0.1070)	(0.0029)	(0.0060)	(0.0061)	(0.0029)	(0.0028)	(0.0074)
(5) $\sigma_\alpha^2 = 0.05, \sigma_\epsilon^2 = 0.04, \sigma_\eta^2 = 0.01$	0.8278	1.0007	1.0132	0.8305	1.0008	1.0605	0.9280	0.9898	0.9201	0.9289	0.9900	0.9106
	(0.0186)	(0.1048)	(0.2372)	(0.0191)	(0.1066)	(0.2354)	(0.0126)	(0.0234)	(0.0591)	(0.0119)	(0.0256)	(0.0390)

Notes: The table shows estimated persistence for various N, initial conditions and weighting schemes from simulated data for $\rho = 0.995$ and $\rho = 0.90$ and various values for the variances of fixed effects and shocks. σ_α^2, σ_η^2 and σ_ϵ^2 are the variances of fixed effects, persistent and transitory shocks, respectively. Standard errors are in parentheses. "E" ["O"] ("D") stands for equally [optimal] (diagonally) weighted estimations. The variance of the permanent component in period 0, var(z_{i0}), equals 0 or $\sigma_\eta^2(1-\rho^{30})/(1-\rho^2)$ for zero and nonzero initial conditions, respectively.

we simulate data for a very large number of individuals—$N = 50,000$—using our benchmark parameters and the persistence of permanent shocks of 0.90 and plot the minimized objective function values at each value of the persistence grid. Panels (a)–(b) of Figure 3.1 (A) plot the minimized objective values for the equally-weighted estimations using the first fifteen and the last fifteen periods for the case of zero initial conditions, whereas Panel (c) does the same but for the estimations based on the first fifteen periods and nonzero initial conditions for the permanent component, with the variance equal to $\sigma_\eta^2(1-\rho^{30})/(1-\rho^2)$. The plots have two minima, one well below the true value of 0.90 and another one above 1 but below 1.1.

All estimations capture well the global minimum that generates the downward bias in the estimated persistence we have documented in Tables 3.1 and 3.3. Panels (d)–(f) and (g)–(i) of Figure 3.1 (A) contain the same plots but for the optimally and diagonally-weighted estimations, respectively. When the first fifteen periods are used in estimation and initial conditions for the permanent component are zero, both optimal and diagonal weighting estimations capture well the global minimum visible in Panels (d) and (g). When the variance of permanent incomes in the initial year of the sample is high—due to nonzero initial conditions or when using the last fifteen periods for the case of zero initial conditions—there are two local minima, one around 0.90 and another one above 1 that yield close values of the objective function; see panels (e)–(f) for optimal and (h)–(i) for diagonal estimations, respectively. Both of these minima are frequently chosen in the estimations resulting in an upward-biased estimate of the persistence when the variance of incomes in the first year of the sample is high and either optimally- or diagonally-weighted estimations are utilized. These estimations also result in a relatively higher variability in the

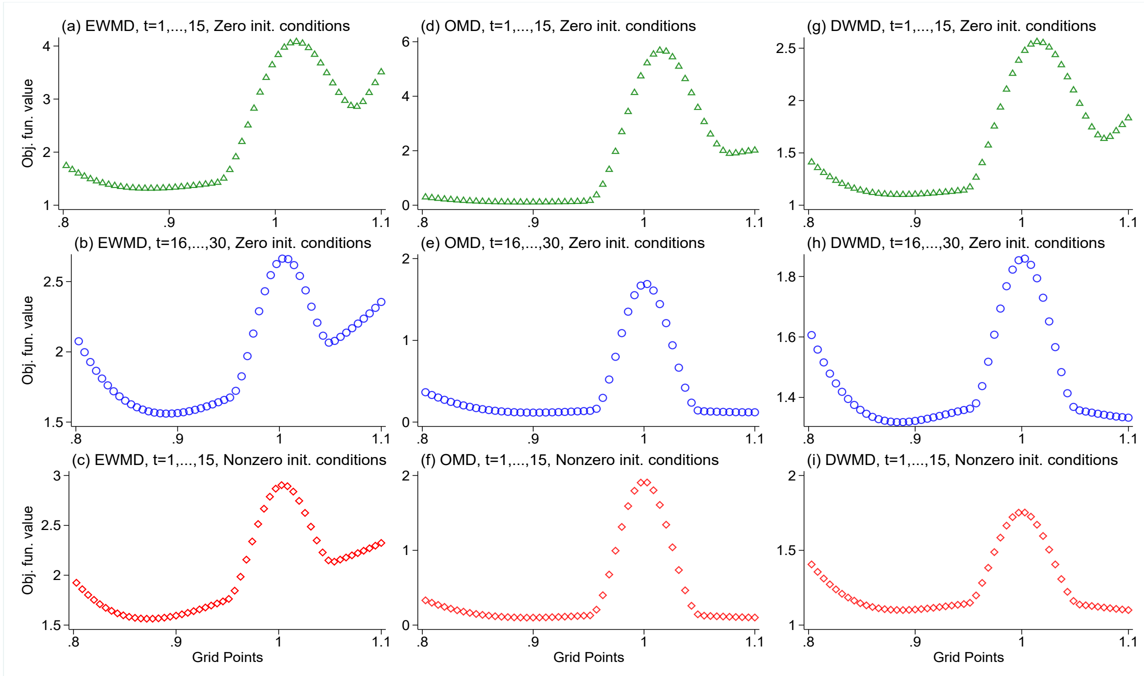

(A) Low persistence: $\rho = 0.90$

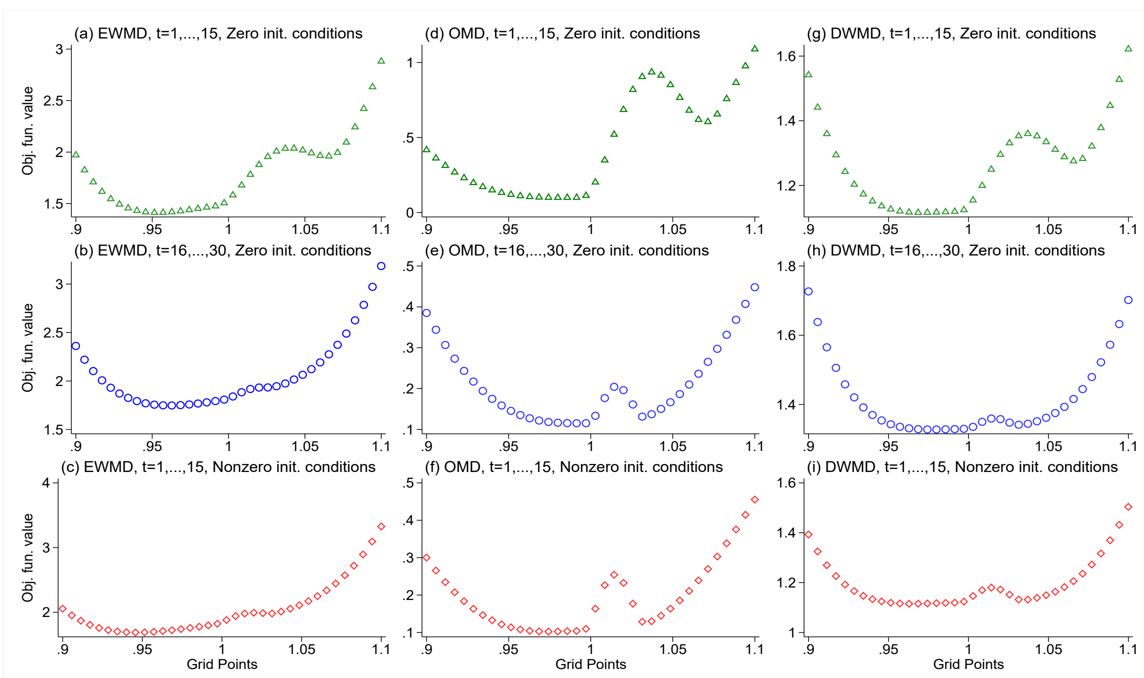

(B) High persistence: $\rho = 0.995$

Figure 3.1: Objective Function Value at Grid Points for the Persistence. Simulated data. Large N

Notes: Each panel shows objective function values at various grid points for the persistence for equally, optimally, and diagonally-weighted minimum distance estimation (EWMD, OMD and DWMD). OMD and DWMD objective function values are divided by 1000 and 100, respectively.

estimated persistence.

When the true persistence is high, equally-weighted estimations using the first fifteen or the last fifteen observations produce global minima, which are well below the true value of persistence—panels (a)–(c) of Figure 3.1 (B). Similar to the case of low persistence, estimations capture these global minima well. For diagonal and optimal weighting, some estimations feature local minima, below and above the true value, that are not far apart—panels (e), (f), (h) and (i). As with the case of low persistence, these estimations result in a relatively higher variability of the estimated persistence.

3.5 Biases in the Estimated Parameters: Quantitative Evaluation

In this section, we, first, examine biases in the persistence and the other parameters using the regression analysis; second, perform an empirical application of estimation in quasidifferences to Danish administrative data on male earnings; and, third, analyze biases in the income and consumption process parameters from a minimum-distance estimation based on consumption data and income data in quasidifferences from a calibrated standard incomplete markets model.

3.5.1 Biases in Persistence

In Table 3.4 we report the results from regressions of a bias in the estimated persistence, $\hat{\rho} - \rho$, on the true values for the variances of fixed effects, permanent shocks and transitory shocks. In odd columns, we combine the samples that use the first fifteen or the first thirty observations, while in even columns, we use the simulated data for the last fifteen observations from Tables 3.1–3.2. Our qualitative analysis above emphasized that those samples yield drastically different results for the estimations relying on the optimal and diagonal weighting. The dependent variable is divided by 100 and the independent variables are standardized so that they all have mean zero and a standard deviation of one.

When the true persistence is low, equally-weighted estimates have an average downward bias of 0.06 to 0.07 in small samples; the bias drops to 0.02–0.03 in samples with 10,000 individuals—see the estimated constants in columns (1)–(2) of Panel A and B, respectively. In small samples, an increase in the variance of transitory shocks by one standard deviation,[9] ceteris paribus, raises the downward bias by additional 0.04 points, resulting in an estimate of the persistence of 0.80. A reduction in the variance of permanent shocks by one standard deviation raises the downward bias by additional 0.02 points, whereas a reduction in the variance of fixed effects by one standard deviation raises the downward bias by about 0.02 points.[10] All these effects are significant at the 1 percent level. They are quantitatively smaller when N is large, although they are still statistically significant—panel B, columns (1)–(2). Thus, even if samples are large, a particular configuration of model parameters could result in biased estimates of the persistence.

When the true persistence is high, the average bias is similar, the effects of a variation in the variance of fixed effects and transitory shocks are somewhat smaller, whereas the effect of a change in the variance of permanent shocks is somewhat larger—columns (7)–(8).

Optimally-weighted estimates have small biases when the first fif-

[9] In the estimation sample of column (1), mean and standard deviations for the variance of transitory shocks equal about 0.03 and 0.015, respectively.

[10] Mean and standard deviation for the variance of permanent shocks and fixed effects in this sample are 0.02 and 0.09, and 0.015 and 0.02, respectively.

Table 3.4: Bias in the estimated persistence. Regression analysis

	$\rho = 0.90$						$\rho = 0.995$					
Weighting	Equal		Optimal		Diagonal		Equal		Optimal		Diagonal	
$t =$	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Panel A: $N = 1000$											
σ^2_α	1.961***	2.184***	0.607***	0.021	2.228***	2.813	0.693***	0.821***	0.270***	-0.171	1.636***	0.350
	(0.142)	(0.230)	(0.179)	(0.676)	(0.268)	(1.786)	(0.102)	(0.161)	(0.100)	(0.259)	(0.374)	(0.886)
σ^2_η	1.777***	2.135***	-0.466***	0.538	2.852***	0.344	2.946***	3.144***	0.481***	0.756***	1.972***	2.090***
	(0.114)	(0.189)	(0.140)	(0.536)	(0.202)	(1.073)	(0.079)	(0.120)	(0.078)	(0.193)	(0.232)	(0.473)
σ^2_ϵ	-3.662***	-4.222***	0.144	-1.257**	-4.783***	0.516	-1.851***	-1.684***	0.109	0.216	-1.269***	2.079***
	(0.114)	(0.189)	(0.140)	(0.536)	(0.202)	(1.073)	(0.079)	(0.120)	(0.078)	(0.193)	(0.232)	(0.473)
Const.	-6.261***	-6.923***	0.117	11.051***	-8.844***	11.194***	-5.110***	-5.046***	-0.717***	0.973***	-5.212***	3.314***
	(0.106)	(0.173)	(0.124)	(0.518)	(0.182)	(1.132)	(0.075)	(0.112)	(0.071)	(0.185)	(0.228)	(0.527)
	Panel B: $N = 10000$											
σ^2_α	0.629***	0.861***	0.007	-0.014	0.929***	-1.467	0.312***	0.543***	0.014	0.033	0.452***	0.823***
	(0.067)	(0.107)	(0.048)	(0.559)	(0.068)	(1.172)	(0.055)	(0.070)	(0.038)	(0.131)	(0.066)	(0.218)
σ^2_η	1.057***	1.181***	-0.040	-0.449	1.676***	-1.196	1.779***	2.044***	0.382***	0.269***	2.390***	3.068***
	(0.057)	(0.089)	(0.033)	(0.514)	(0.062)	(0.751)	(0.040)	(0.063)	(0.025)	(0.084)	(0.048)	(0.160)
σ^2_ϵ	-1.510***	-1.781***	0.025	0.556	-2.214***	0.547	-0.837***	-0.946***	-0.025	0.070	-1.110***	-0.759***
	(0.057)	(0.089)	(0.033)	(0.514)	(0.062)	(0.751)	(0.040)	(0.063)	(0.025)	(0.084)	(0.048)	(0.160)
Const.	-2.326***	-2.818***	0.063*	9.233***	-3.993***	11.701***	-2.614***	-3.085***	-0.483***	-0.271***	-3.581***	-3.671***
	(0.051)	(0.082)	(0.032)	(0.470)	(0.055)	(0.772)	(0.039)	(0.057)	(0.025)	(0.085)	(0.046)	(0.152)
No. obs.	1000	500	1000	500	1000	500	1000	500	1000	500	1000	500

Notes: The table contains the results from a regression of the bias in the persistence scaled by 100, $100 \cdot (\hat{\rho} - \rho)$, on the standardized variances of fixed effects, σ^2_α, permanent shocks, σ^2_η and transitory shocks, σ^2_ϵ. Columns $t =$[1–15], [1–30] ($t =$[16–30]) utilize estimation data based on the first fifteen or thirty (last fifteen) observations from Tables 3.1–3.2. Standard errors are in parentheses. *** (**) [*] significant at the 1% (5%) [10%] level.

teen or thirty observations are used and little variation of these biases when the model parameters change—columns (3) and (9). An upward bias is substantial when the last fifteen observations are used and the true persistence is low, and the bias varies little with the model parameters, especially so when N is large—column (4).

Diagonally-weighted results are quantitatively similar to the equally-weighted results when the first fifteen or thirty observations are used in estimation and similar to the optimally-weighted results when the last fifteen observations are used in estimation.

3.5.2 Biases in the Other Parameters

Appendix Tables B-1–B-3 describe the results from analogous regressions for the biases in the variance of fixed effects, permanent and transitory shocks, respectively. Briefly, the variance of fixed effects is severely upward-biased when the true persistence is high regardless of the weighting matrix, sample size, or the number of periods. The variances of permanent and transitory shocks typically have small biases, especially so when samples are big cross-sectionally. See Appendix for the full details.

3.5.3 Empirical Application

To illustrate our results empirically, we use administrative earnings records from Denmark for the cohort of males born in 1952 observed during 1981–2006. We use a balanced sample and the sample selection criteria of Daly et al. (2021) that are standard in the literature. Briefly, we dropped individuals who were ever self-employed, dropped records for males working less than 10 percent of the year as full-time employees, and earnings histories with growth outliers defined as an increase of earnings in adjacent periods by more than 500 percent or a drop by less than –80 percent. Earnings are expressed in 1981 Danish kroner. We remove predictable variation in earnings by running, for each year, cross-sectional regressions of log earnings on educational dummies, a third-order polynomial in age, and the interaction of the educational dummies with the age polynomial. Our final sample comprises 13,543 individuals with 26 earnings observations.

As in Daly et al. (2021), we allow for an MA(1) transitory component but also allow for the variances of permanent and transitory shocks to vary over time.[11] Using an optimal weighting matrix and estimation in quasidifferences, similar to Daly et al. (2021), we find that the variance of transitory shocks is about twice as large as the variance of permanent shocks. The persistence of permanent shocks is estimated at 0.975, close to an estimate in Daly et al. (2021) for their balanced sample. Using diagonal weighting, we estimate persistence at 0.935, whereas equal weighting of the moments yields an estimate of 0.895. The most fitting reference to our empirical setting are the simulation results in rows (1) and (5) and columns (3), (6) and (9) of Table 3.2 Panel B. Those results are based on the experiments where the variance of transitory shocks is relatively higher than the variance of permanent shocks, samples are cross-sectionally large, and individuals have already accumulated persistent shocks for a number of years when they are first observed in the data. Similar to our empirical results, optimal weighting yields a higher estimate of the persistence than equal and diagonal weighting of the moments. Using 100 bootstrap samples of 1,000 individuals from our Danish data, we obtain the average values of persistence of 0.984, 1.00 and 0.867 for optimal, diagonal and equal weighting, respectively. This corresponds to our simulation results where the estimates of the persistence are comparable for big and small samples

[11] Our results in Tables 3.1–3.3 are similar if we allow for an MA(1) transitory component instead of a purely transitory shock.

for optimal weighting (column (6) row (1) Panels A and B, or column (6) row (5) Panels A and B), are substantially higher in small samples when diagonal weighting is used (column (9) row (1) Panels A and B, or column (9) row (5), both panels), and are lower in small samples when equal weighting is used (column (3) row (1), or column (3) row (5) in both panels).

3.5.4 Biases When Using Consumption and Income Data

Blundell et al. (2008) developed a methodology for estimating the income-process and consumption insurance parameters using auto- and cross-covariances of consumption and income growth rates. They applied it to longitudinal data from the U.S., assuming that the permanent component is a random walk. Specifically, they estimated the variances of permanent and transitory income shocks in Eq. (3.1) along with the transmission of those shocks to household consumption, evaluating the following equation:

$$\Delta c_{it} = \phi \eta_{it} + \psi \epsilon_{it} + \zeta_{it} + \Delta u_{it},$$

where c_{it} is household i's log consumption at time t, ϕ and ψ are the transmission coefficients for permanent and transitory income shocks to household consumption, respectively, u_{it} is measurement error in consumption and ζ_{it} is a permanent shock to consumption that is orthogonal to the income shocks.

Estimation using quasidifferencing can be easily embedded into the methodology of Blundell et al. (2008) when the permanent component is, instead, an auroregressive process with finite persistence.[12] Although we showed above that estimation in quasidifferences is often unreliable when

[12]See Kaplan and Violante (2010) for the income-consumption moments, in addition to those listed in Eq. (3.7)–(3.9), needed for identification of income and consumption processes.

using income data, consumption data could provide valuable information for identifying the persistence of the shocks and, used jointly with income, might potentially help correct the biases outlined above. To explore this possibility, we simulate data from a calibration of the standard incomplete-markets model of consumption.

In the model, households value consumption, C, using CRRA utility function $u(C) = \frac{C^{1-\gamma}}{1-\gamma}$, face zero borrowing constraints, and accumulate savings for retirement and precautionary reasons to insulate their consumption from permanent and transitory shocks to income and the predictable income drop at retirement. The real interest rate on saving is set to 4 percent. Households discount future utility exponentially using the time discount factor β and face mortality risk after retirement. They start their working life at age 26, retire at 65 and die with certainty at age 90. Income is exogenous and is subject to the persistent and transitory risk before retirement, as in Eq. (3.1). After retirement, income equals 70 percent of its permanent component at age 65.

We focus on the two persistence values as we have done above. Those values were found to characterize the income dynamics of the families formed by sons and daughters of the original PSID households; see Hryshko and Manovskii (2022). The variance of permanent and transitory shocks and the variance of fixed effects are set to the benchmark values of row (1) in Table 3.1. We assume that all households start with zero permanent component of incomes at age 26, and all shocks and fixed effects are normally distributed. The predictable component of income and mortality risk are taken from Hryshko and Manovskii (2022). The time discount factor is calibrated by matching the wealth-to-income ratio in each economy characterized by different persistence of permanent income shocks to the

Table 3.5: Estimates of the Income Process Parameters. Data from a Calibrated Lifecycle Model

Weighting	$\rho = 0.90$			$\rho = 0.995$		
	Equal	Optimal	Diagonal	Equal	Optimal	Diagonal
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: $N = 1000$					
Persistence, ρ	0.7522	1.0267	0.9469	0.8394	1.0108	1.1364
	(0.0570)	(0.1756)	(0.3537)	(0.0525)	(0.1877)	(0.2815)
Variance fixed effects, σ^2_α	0.1215	0.0696	0.1224	0.2562	0.2967	0.1938
	(0.0086)	(0.0517)	(0.0079)	(0.0285)	(0.8675)	(0.0369)
Variance perm. shocks, σ^2_η	0.0115	0.0119	0.0156	0.0099	0.0096	0.0149
	(0.0570)	(0.0031)	(0.0048)	(0.0020)	(0.0027)	(0.0049)
Variance trans. shocks, σ^2_ϵ	0.0388	0.0383	0.0381	0.0396	0.0385	0.0390
	(0.0015	(0.0012)	(0.0015)	(0.0014)	(0.0011)	(0.0015)
	Panel B: $N = 10000$					
Persistence, ρ	0.8058	0.9889	1.0467	0.8598	0.9889	0.9561
	(0.0248)	(0.1205)	(0.2994)	(0.0378)	(0.0675)	(0.2104)
Variance fixed effects, σ^2_α	0.1208	0.0840	0.1209	0.2648	0.2282	0.2190
	(0.0032)	(0.0335)	(0.0036)	(0.0213)	(0.1793)	(0.0410)
Variance perm. shocks, σ^2_η	0.0103	0.0108	0.0147	(0.0099	0.0100	0.0117
	(0.0006)	(0.0016)	(0.0042)	(0.0005)	(0.0007)	(0.0029)
Variance trans. shocks, σ^2_ϵ	0.0396	0.0399	0.0392	0.0396	0.0399	0.0394
	(0.0005)	(0.0004)	(0.0005)	(0.0004)	(0.0003)	(0.0005)

Notes: The table shows results from estimations in quasidifferences using simulated income data from the calibrated standard incomplete-markets model for different values of persistence, N and weighting schemes. See Section 3.5.4 for the details on the model and simulations. Age distribution in the simulated data is as in the PSID. True values for the variance of fixed effects, permanent, and transitory shocks are 0.10, 0.01, and 0.04, respectively. Standard errors are in parentheses.

value of 3. The calibrated values are 0.9529 and 0.9555 in the economies with the persistence of permanent shocks equal to 0.995 and 0.90, respectively. Table 3.6 summarizes details of the calibration.

We next simulate income and consumption data for a large number of households and randomly select either 1,000 or 10,000 out of the simulated dataset. We assume that consumption is measured with error in the data and set the variance of measurement error to 0.04. We replicate the PSID design of the data in Hryshko and Manovskii (2022) by creating

Table 3.6: Calibration

Parameter	Value	Internally calibrated	Data source
Income age profile, ages 26–65	various	No	Kaplan and Violante (2010)
Survival probabilities, ages 66–90	various	No	Hryshko (2014)
CRRA, γ	2	No	Kaplan and Violante (2010)
Interest rate	4%	No	Carroll (2009)
Variance of permanent income shocks, σ_η^2	0.01	No	benchmark value
Variance of transitory income shocks, σ_ϵ^2	0.04	No	benchmark value
Variance of fixed effects, σ_α^2	0.10	No	benchmark value
Persistence of permanent shocks, ρ	0.90 [0.995]	No	Hryshko and Manovskii (2022)
Replacement rate	0.70	No	Hryshko and Manovskii (2022)
Time discount factor, β	0.9555 [0.9529]	Yes	

Notes: The table shows various inputs for calibration of the standard incomplete-markets economy with zero borrowing constraints. The time discount factor is calibrated internally by matching the wealth-to-income ratio in the economies characterized by different values of the persistence of permanent shocks to the value of 3.

fifteen years of data on consumption and income and matching the age distribution in the PSID during the 1978–1992 period analyzed in Blundell et al. (2008) and Hryshko and Manovskii (2022). Specifically, we observe the distribution of birth years for households formed by sons and daughters of the original PSID households, and we use the income and consumption data when those households are 30 to 65 years of age within the 1978–1992 period.

This is a different data structure relative to the ones we analyzed above due to the presence of various birth-year cohorts in each year. In the simulated data, similar to the PSID, households are on average 36.5 years of age in 1978 (first year of the simulated dataset) and 42.5 years of age in 1992 (last year of the dataset). Since many households have accumulated a substantial history of shocks prior to the first year in the simulated data, we expect our results to resemble the above results for the estimations relying on periods 16 to 30 of the simulated income data. Table 3.5 presents the results from estimations in quasidifferences using simulated income data only and verifies this conjecture. Both for small and large samples, equal weighting yields substantial downward biases in the estimated persistence when the true persistence is low or high. Diagonal and optimal weighting yield substantial upward biases when the true persistence is low for small and large samples, and some downward biases when the true persistence is high and samples are large. The estimated variances of permanent and transitory shocks are not very far from their true values, whereas the variance of fixed effects is biased, more so for the estimations on the data with high true persistence.

Table 3.7 shows the results for the income-process and consumption insurance parameters from estimations using consumption data and incomes in quasidifferences. Equal weighting yields an estimate of persistence close to its true value when the true persistence is low and the sample is large, whereas optimal weighting yields slightly upward-biased estimates of the persistence when the true persistence is high, both in small and large samples. The variance of permanent shocks is substantially overestimated when the diagonal weighting matrix is used. The true transmission coefficients for permanent and transitory shocks in our low-persistence economy are 0.59 and 0.25, respectively. Both transmission coefficients are close to their true values when optimal weighting is used, both for small and large samples. However, in this setting, one would conclude that consumers are excessively insured against permanent income shocks since the estimated persistence is well above the true value of 0.90. Diagonal weighting produces noisy and biased estimates of the persistence and the transmission coefficient for permanent shocks in small samples and is very unreliable

Table 3.7: Estimates of Income and Consumption Insurance Parameters. Data from a Calibrated Lifecycle Model

	$\rho = 0.90$			$\rho = 0.995$		
Weighting	Equal	Optimal	Diagonal	Equal	Optimal	Diagonal
Parameters	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: $N = 1000$					
Persistence, ρ	0.7775	1.0188	0.9366	0.8763	1.0096	1.2313
	(0.0400)	(0.0821)	(0.3576)	(0.0294)	(0.0535)	0.1802)
Variance fixed effects, σ_α^2	0.1177	0.0347	0.1197	0.2494	0.0193	0.1685
	(0.0094)	(0.0465)	(0.0082)	(0.0244)	(0.0467)	(0.0215)
Variance perm. shocks, σ_η^2	0.0127	0.0100	0.0196	0.0130	0.0081	0.0178
	(0.0013)	(0.0021)	(0.0066)	(0.0012)	(0.0009)	(0.0028)
Variance trans. shocks, σ_ϵ^2	0.0383	0.0358	0.0369	0.0383	0.0351	0.0385
	(0.0013)	(0.0011)	(0.0013)	(0.0011)	(0.0011)	(0.0013)
Variance meas. error, σ_u^2	0.0417	0.0367	0.0421	0.0410	0.0367	0.0416
	(0.0009)	(0.0010)	(0.0010)	(0.0010)	(0.0011)	(0.0011)
Transm. of perm. shocks, ϕ	0.7075	0.6554	0.5503	0.9169	0.9165	0.7115
	(0.0681)	(0.0776)	(0.1020)	(0.0720)	(0.0597)	(0.1025)
Transm. of trans. shocks, ψ	0.2704	0.1943	0.2828	0.2277	0.1679	0.2551
	(0.0322)	(0.0287)	(0.0450)	(0.0367)	(0.0303)	(0.0390)
	Panel B: $N = 10000$					
Persistence, ρ	0.8845	1.0489	1.3996	0.9579	1.0161	1.1780
	(0.0196)	(0.0024)	(0.0028)	(0.0102)	(0.0256)	(0.0141)
Variance fixed effects, σ_α^2	0.0913	0.0006	0.1158	0.2238	0.0000	0.1393
	(0.0168)	(0.0016)	(0.0019)	(0.0410)	(0.0000)	(0.0069)
Variance perm. shocks, σ_η^2	0.0120	0.0102	0.0273	0.0132	0.0107	0.0161
	(0.0004)	(0.0002)	(0.0007)	(0.0004)	(0.0010)	(0.0007)
Variance trans. shocks, σ_ϵ^2	0.0391	0.0387	0.0375	0.0387	0.0378	0.0389
	(0.0004)	(0.0003)	(0.0003)	(0.0004)	(0.0008)	(0.0004)
Variance meas. error, σ_u^2	0.0414	0.0398	0.0421	0.0411	0.0394	0.0416
	(0.0003)	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0004)
Transm. of perm. shocks, ϕ	0.7539	0.6431	0.4419	0.9023	0.9183	0.7610
	(0.0211)	(0.0120)	(0.0121)	(0.0222)	(0.0239)	(0.0299)
Transm. of trans. shocks, ψ	0.2818	0.2101	0.3167	0.2404	0.1691	0.2504
	(0.0106)	(0.0047)	(0.0095)	(0.0097)	(0.0269)	(0.0087)

Notes: The table shows results from estimations using simulated consumption data and income data in quasidifferences from the calibrated standard incomplete-markets model for different values of persistence, N and weighting schemes. See Section 3.5.4 for the details on the model and simulations. Age distribution in the simulated data is as in the PSID. True values for the variance of fixed effects, permanent, and transitory shocks are 0.10, 0.01, and 0.04, respectively. Standard errors are in parentheses.

in large samples.

The true transmission coefficients for permanent and transitory shocks in our high persistence economy are 0.90 and 0.20, respectively. Diagonal weighting substantially underestimates the transmission of permanent shocks to consumption. Equal weighting yields transmission coefficients that are close to their true values in small and big samples yet significantly underestimates the persistence, producing an erroneous impression of underinsurance of permanent income shocks. Diagonal weighting produces slightly biased transmission coefficients but an upward-biased estimate of the variance of permanent shocks.

To sum up, using data from a calibrated lifecycle model of consumption, we showed, for the benchmark income parameters, that the estimated income and consumption process parameters are biased regardless of the weighting scheme and the number of sample households.

3.6 Implications

Despite its virtues in implementation, estimation in quasidifferences is biased in the predominant number of cases that we considered. What are the implications for researchers who wish to rely on quasidifferences to recover the income-process parameters? There are a number of positive results from our analysis. First, looking closely at the results in Tables 3.1–3.3 reveals that, regardless of the sample size, initial conditions, or the true persistence, the estimated persistence is virtually unbiased if one uses equally-weighted estimation when the variance of permanent shocks is significantly bigger than the variance of transitory shocks—columns (1)–(3) and row (4) of Tables 3.1 and 3.2, and columns (1) and (7), row (3) in Table 3.3.[13] Al-

[13]Optimal and diagonal weighting also produce unbiased estimates of the persistence when the variance of permanent shocks is bigger than the variance of transitory shocks but only when the variance of the permanent component in the first sample year is small (e.g., when following a cohort of individuals from the start of their working careers).

though this is not characteristic of U.S. survey and administrative data,[14] it is the case for administrative income data from Norway as was found in Blundell et al. (2015). This result is based on the knowledge of the relative variances of permanent and transitory shocks, not available a priori. However, our second positive result is that the variances of permanent and transitory shocks are not significantly biased for the cases we considered. Our advice, therefore, is to apply equally-weighted estimation of the income process in quasidifferences and use the results with confidence if the variance of permanent shocks is found to be larger than the variance of transitory shocks. It would be ideal to have a robust method for correcting the biases in all other circumstances, but unfortunately, after some experimentation, we couldn't come up with such a method.[15]

3.7 Conclusion

The income process parameters are key in quantitative models featuring incomplete insurance markets. To recover the parameters, one typically matches autocovariance moments for incomes in levels or growth rates. Blundell et al. (2015) recently applied estimation in quasidifferences to administrative data on incomes from Norway assuming a variant of the canonical income process featuring fixed effects, permanent and transitory components. Estimation in quasidifferences combines the features of estimations in levels and differences but relies on more concise moments

[14]See, e.g., Guvenen et al. (2021) and Hryshko and Manovskii (2022) for recent studies using U.S. administrative and survey data, respectively.

[15]We tried a debiased estimator of Chen et al. (2019) that relies on splitting the sample in half on the cross-sectional or time dimension to correct biases in the persistence but, as can be seen from Tables 3.1 and 3.2, splitting on the time dimension will not work (e.g., the estimated persistence in columns (1), (2), (3) relying on fifteen or thirty periods is very similar for the equally-weighted case), and the estimated persistence does not drastically change when splitting the sample, big or small cross-sectionally, in half (we do not report these results to avoid cluttering).

when the true persistence of permanent shocks is different from one. It also requires searching through a predefined grid for the persistence of permanent shocks as the persistence is not jointly identified with the other parameters. In contrast to a voluminous GMM literature, nothing is known about how well the persistence is recovered using estimation in quasidifferences. In this chapter, we provide a guide to estimating the canonical income process using quasidifferences by conducting Monte Carlo simulations and cataloging biases for various N, T, initial conditions and weighting schemes.

We find that equally-weighted estimations result in downward-biased estimates of the persistence when the variance of transitory shocks is higher than the variance of permanent shocks, whereas optimally and diagonally weighted estimations result in upward-biased estimates of the persistence when the variance of the permanent component in the first sample year is nonnegligible, which is a typical feature of the data containing individuals of different ages in the first sample year of the data. The variance of fixed effects is substantially biased upward when the true persistence is high. The biases in the estimated variances of permanent and transitory shocks are, however, small. Our estimations based on Danish administrative earnings data yield divergent estimates of the persistence for different weighting schemes, which conform to the simulated results for the high persistence and nonzero initial conditions. Using data from a calibrated lifecycle model of consumption, we also show, for the benchmark income parameters, that the estimated income and consumption process parameters are biased regardless of the weighting scheme and the number of sample households.

The results in the chapter provide a warning against the routine use of estimation in quasidifferences despite its attractive features. However, there is one case when estimation in quasidifferences recovers the true persistence of permanent shocks. When the variance of permanent shocks is higher than the variance of transitory shocks, equally-weighted estimation is reliable for different T, N and initial conditions. Since biases in the variances of the shocks using estimation in quasidifferencs are small, our advice is to estimate the income process using equal weighting of the moments and use the results with confidence if the estimated variance of permanent shocks is higher than the variance of transitory shocks. This is, e.g., the case for Norwegian administrative data as shown in Blundell et al. (2015).

Chapter 4

The Role of Private and Public Transfers for Consumption Insurance in South Africa

4.1 Introduction

South Africa has been struggling with high levels of unemployment for several decades, which is worsened by continued sluggish economic growth.[1] High unemployment makes it harder for many families to earn labour income. In the absence of labour earnings, households often resort to reducing consumption expenditure when they lack access to other sources of income, such as private and public transfers, personal savings and borrowing. Families that reduce consumption, especially when it is already close to subsistence, can experience long-term negative effects.[2] Consumption can also vary due to other household income shocks, e.g., occupational disability, property theft, death, sickness and poor harvest, and these shocks make access to different sources of income for smoothing consumption critical. This chapter estimates the ability of households in South Africa to insure consumption from income shocks during the period 2008-2017 and explores the role of different income sources in providing insurance.

The inefficiencies of the labour market in South Africa partly reflect the effects of Apartheid (1948-1994) policies that disenfranchised some groups, primarily based on race. The disenfranchised groups were not allowed to own assets such as fertile land or formal housing, likely creating long-run disparities in households' ability to accumulate wealth.[3] Upon gaining independence in 1994, the South African government expanded access to social welfare programs and enacted policies to create opportunities for formerly disadvantaged groups.[4] The abolishment of restrictive laws

[1] See Levinsohn (2008), Alvaredo and Atkinson (2010) and Festus et al. (2016).

[2] See Jacoby and Skoufias (1997), Janzen and Carter (2018) and De Magalhães et al. (2019).

[3] Although black people constituted more than 75 percent of the population, the 1913 Land Act set aside only 7.5 percent of the land in South Africa for black people, and owning land was restricted to tribal areas.

[4] Public transfers include public works programs, unemployment insurance, medical

led to rural-to-urban migration to seek better employment opportunities.

Family members who migrate to other regions or cities are a significant source of income for some rural households through private transfers. These members often send transfers home to supplement their income. Government social welfare programs are also becoming an essential source of income for many households. Fig. 4.1 shows that these two income sources constitute more than 85 percent of household income for families with no labour earnings; the proportion of individuals living in households with no labour income earners increased from 18 percent in 2008 to 25 percent by 2019; see Fig. C-1. The reliance on public transfers is significantly higher for less-educated households; see Fig. C-2. This chapter estimates the role of these two income sources as insurance devices.

Households are likely to adjust their consumption in response to fluctuations in income when they have insufficient resources to cushion consumption from movements in income. The availability of insurance devices can help attenuate the transmission of shocks to consumption, and the devices include personal savings, borrowing from credit markets (formal or informal), family labour supply, welfare benefits, transfer payments and the tax system. There is also growing literature studying the use of informal channels to smooth consumption.[5] Understanding the magnitude

provisions, housing subsidies and cash. The cash welfare program includes child support grants, foster child grants, old-age pensions, disability grants, care dependency grants and war veterans grants. Section 9(2) of the South African Constitution allows the enactment of measures to protect and advance the interests of previously disadvantaged groups.

[5]See Townsend (1994) for India Attanasio and Ríos-Rull (2000) for Mexico and Kinnan (2022) for Thailand.

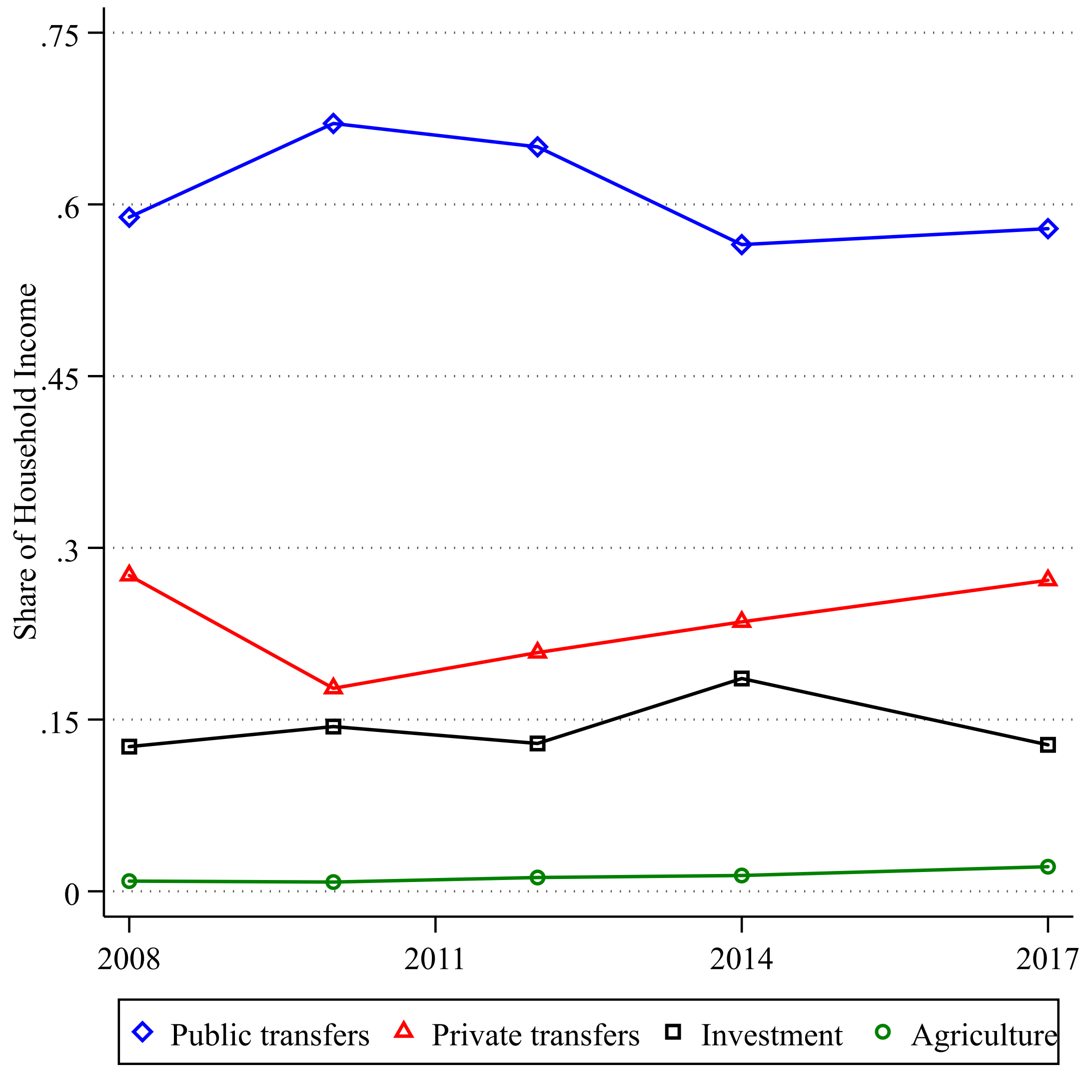

Figure 4.1: Household Income Shares for Households with Zero Labour Income

Notes: The figure depicts the proportions of public transfers (social grants), private transfers from family and friends, investment, and agriculture farming in household income for households without labour income. The shares are computed from amounts in real Rand terms, deflated using a monthly CPI corresponding to the month and year of the interview.

of consumption insurance in developing countries can help shape public policies (e.g., tax systems, food relief programs, agriculture input schemes and welfare transfers) that are better at assisting households in mitigating the impact of adverse shocks.[6]

[6]See Baez (2006) for a discussion on income volatility and risk coping mechanisms in

This chapter uses the methodology of Blundell et al. (2008) to estimate the proportion of movements in income passing through to consumption. Under the methodology, income is a sum of permanent and transitory components and shocks to these components cause consumption to vary. The transmission parameters of the shocks are identified from the comovement of income and consumption growth by minimizing the distance between data and model moments. The proportion of shocks that does not transmit to consumption indicates the degree of consumption insurance. This estimation approach differs from other South African studies that focus on the role of specific programs.[7] Rather, it provides an estimate of the overall ability of households to cushion shocks while abstracting from the source of the risk.

The results indicate that households insure about 31 percent of permanent income shocks. In comparison, households have full insurance against transitory income shocks—estimates of the transmission of transitory shocks to consumption are small (close to zero) and mostly insignificant. There is significant heterogeneity in the level of insurance across education, race, region of residence and age. Households headed by educated, white, or older heads or those living in urban areas have relatively more insurance than the average household. As for insurance devices, private transfers seem more effective in helping households smooth shocks than public transfers. Because private transfers flow from multiple sources, are more responsive to the occurrence of a shock and require no means-testing

---

developing countries, especially for rural families. The chapter also reviews several studies that estimate the degree of insurance against specific shocks (e.g., crop failure, illness) and potential coping devices.

[7] See Case and Deaton (1998), Bertrand et al. (2003), Duflo (2003), Dutronc-Postel and Tondini (forthcoming), Ardington et al. (2009) on social pensions. Case et al. (2005) and Edmonds (2006) on child support grants.

like public transfers, it makes them more effective at providing insurance.[8]

4.2 Related Literature

Studies on the comovement of income and consumption date back many decades ago, e.g., Modigliani and Brumberg (1954), Friedman (1957). Since then, different theories have emerged that theorize the relationship between consumption and income. The complete-markets hypothesis predicts that households can fully insure all idiosyncratic income shocks. On the other hand, incomplete-markets models such as the life cycle hypothesis of Modigliani and Brumberg (1954) and permanent income hypothesis of Friedman (1957) assert that households use personal savings to smooth transitory income shocks but do not have insurance against permanent shocks. Many studies (Campbell and Deaton, 1989; Cochrane, 1991; Townsend, 1994; Attanasio and Davis, 1996) test these competing hypotheses and find mixed results. Blundell et al. (2008) propose a flexible partial insurance model that accommodates these two competing hypotheses by allowing consumption to respond to both transitory and permanent income shocks. They then use the model to estimate consumption insurance in the US using Panel Study of Income Dynamics (PSID) data. They find full insurance against transitory shocks and 36 percent partial insurance of permanent shocks.

Many studies have since employed various versions of the Blundell et al. (2008) methodology. De Nardi et al. (2019) use a richer income process that includes a persistent rather than permanent component and find a similar degree of partial insurance (43 percent) of permanent shocks. Kaplan and Violante (2010) calibrate a standard lifecycle incomplete-markets

[8] See Fig. C-7 in the Appendix on the multiple sources of private transfers.

model and report a lower (22 percent) partial insurance against permanent shocks in the model. The degree of insurance in the model was also sensitive to the assumptions about the ability of households to borrow. In other developed countries, Casado (2011) uses the Spanish Household Continuous Survey and the Blundell et al. (2008) methodology to estimate the degree of insurance in Spain. The study reports insurance against permanent income shocks of 52 percent, with the degree of insurance increasing with education and household wealth. Ahn et al. (2021) and Kubota (2021) report similar magnitudes for South Korea and Japan, respectively. Gathergood and Wylie (2018) show significant evidence of households borrowing from credit markets and social networks to smooth consumption in the UK. They find that the level of insurance increases with the degree of financial literacy of the household head.

Similar research for developing countries is limited due to the lack of panel data. Townsend (1994) tests and statistically rejects the presence of full insurance in rural India but documents significant risk-sharing across families. Santaeulàlia-Llopis and Zheng (2018) build a panel data set running from 1989 to 2009 in China to analyze the degree of consumption insurance. They examine how insurance changed over time as China went through decades of rapid economic growth. The main results show that the ability to cushion consumption from permanent income shocks deteriorated with economic growth.

In Africa, Dercon et al. (2005) examine the ability of households to smooth health and drought shocks in Ethiopia. They show that female-headed, low-educated and low-wealth households are the most vulnerable groups to drought shocks. De Magalhães et al. (2019) study consumption smoothing over the life cycle in Malawi and show that rural households rely significantly on self-farmed food rather than formal insurance channels. They also report that households are likely to withdraw their children from school when faced with income shocks. Dafuleya and Tregenna (2021) show that households in Bulawayo, a city in the south of Zimbabwe, use both formal and informal channels to insure the family from funeral shocks. They find that funeral insurance or availability of multiple assets helps households partially insure food consumption from death shocks. Riley (2018) shows that rural households in Tanzania significantly rely on remittances after experiencing floods or drought shocks. The probability of receiving the transfers is higher for families with access to mobile money services because it reduces the cost of sending transfers. Bellemare et al. (2021) show that contract farming is an effective insurance mechanism for households with limited access to formal insurance markets in Madagascar.

Unlike many African countries, South Africa has a more advanced social welfare program. These government programs include a series of grants: child support grants, child foster grants, disability grants and old-age pension payments, unemployment insurance and workers' compensations. Many studies analyze the role of these government programs on food consumption (Case and Deaton, 1998), labour supply (Bertrand et al., 2003; Dutronc-Postel and Tondini, forthcoming), school choices (Case et al., 2005; Edmonds, 2006).

Case and Deaton (1998) studies the relationship between income and food consumption. They focus on the role of old-age pensions in South Africa and find that income from pensions significantly affects food consumption. Bertrand et al. (2003) observe a decrease in employment when a household has a member who qualifies to receive an old-age pension. Edmonds (2006) documents a significant relationship between the decision to send children to school and the timing of social pension payments. These results support findings in Case et al. (2005) who find that children that receive child support grants are more likely to be enrolled in school. Bellon et al. (2020) examine the evolution of household income shocks and consumption in South Africa and Tanzania. However, they only focus on the impact of job loss. They report that government transfers and access to credit play a significant role in cushioning households from job loss.

These South African studies either focus on the impact of specific shocks, e.g., drought, death of a family member, job loss, or on the role of specific income sources such as old-age pension and child support grants. Instead, this chapter remains agnostic on the source of income shocks and only classifies shocks as either permanent or transitory. It then examines the role of different income sources in shielding these two broad classes of shocks.

4.3 Data

The chapter uses the public files of the National Income Dynamics Survey (NIDS) data from South Africa. NIDS commenced in 2008 by collecting information on a nationally representative sample of 7,300 households. Since then, the survey has been conducted biennially following the original households and their split-offs, that is, children of the original sample. There are currently five waves available.[9] Due to attrition problems, especially among the White and Asian households and high-income earners, a top-up sample of 2,775 households was added in Wave 5. The top-up sam-

[9]The 5th wave was mainly conducted in 2017 rather than 2016. I adjust for this time difference in the estimation of the results similar to Santaeulàlia-Llopis and Zheng (2018) who also make adjustments for the Chinese data.

ple aims to maintain the national representativeness of the panel. Thus, an individual may appear at least once and a maximum of five waves. The response rate for Waves 1 to 5 among the original sample is 78 percent or above. The survey collects socio-demographic and socioeconomic characteristics, including income and disaggregated expenditure categories such as food, health and transportation. Details on household wealth were only collected in Waves 2, 4 and 5. The information is gathered via a recall method where respondents are asked the level of expenditure or income in the last month. Some variables, such as employment hours, are recorded as weekly estimates.

The final sample is restricted to household heads between the ages of 25 and 65.[10] The study excludes households from the 2017 top-up subsample because the estimation of insurance parameters requires at least two consecutive observations per household. Also, the chapter winsorizes the bottom and top 1 percent income observations to minimize the problem of outliers. Households whose income grew by more than 500 percent are also excluded from the sample. The focus is on nondurable consumption, including food, alcohol, tobacco, utilities, public and private transport, personal care, education, donations, entertainment and clothing. Section V provides an estimate of consumption insurance for aggregate consumption that includes both nondurable and durable expenditure and explores the effect of excluding some categories that can be classified as investment in human capital (e.g., education and health expenditure). All the income and expenditure variables are in real terms, deflated using the monthly

[10] The cut-off at the age of 65 is to ensure that the final sample excludes retired household heads. The trade-off is that this sample restriction will likely minimize the role of social assistance in providing household insurance because many family heads who receive welfare transfers are old, for instance, old-aged and child support grants for some retirees who look after their grandchildren.

consumer price index corresponding to the month of data collection. The final sample consists of 28,160 observations.

The Southern Africa Labour Dynamics Research Unit (SALDRU), the organization conducting the surveys, imputes some of the components of income.[11] Following Blundell et al. (2008), the baseline results exclude financial income, but Section V explores the contribution of financial income as a source of consumption insurance.

Table 4.1 presents descriptive statistics from the baseline sample across households and household head characteristics. The demographic characteristics are relatively stable across the study period, with few exceptions. First, the average age of the household heads is about 44 years, with the fraction of those who are more educated (high school and college) increasing at the expense of those with less than high school. Given that the attrition was mainly among the groups with higher education (high-income earners, White individuals and Asian individuals), the rise in education is likely to be higher. Second, there is a downward trend in the share of White households in the sample, which is partly why SALDRU introduced an additional subsample in Wave 5. About half of the households are employed, and of those in employment, 13 percent work in unionized jobs. Close to 60 percent of the households live in urban areas with an average household size of four people and two children.

[11] According to the NIDS Survey User Guides, *Item non-response occurs when the respondent refuses to answer a particular question in the survey or states that they "Don't Know" the answer. In these circumstances, imputation can be performed on the individual variables affected. This is conducted only once a few qualifying conditions are satisfied. Single imputation regressions are run only when there are: a) 100 or more "valid" responses for a variable and b) the percent of missings does not exceed 40 percent... rule-based imputation process is followed for the State Old Age Pension, Child Support Grant, Disability Grant, Care Dependency Grant and Foster Care Grant. Respondents acknowledging receipt of one of these grants, but failing to provide an amount, are assigned the maximum value of the grant for the month in which the interview took place. This is because individuals receiving one of the state grants rarely receive less than the full amount.*

Table 4.1: Cross-sectional Descriptive Statistics, 2008-2017

	Variable Means				
Year	2008	2010	2012	2014	2017
	Panel A: Household Head Characteristics				
Age	45.11	45.46	44.17	44.14	43.98
Gender (male = 1)	0.573	0.443	0.388	0.444	0.427
Education (Percentage %):					
Less than High School	0.698	0.714	0.673	0.656	0.617
High School	0.149	0.137	0.161	0.174	0.203
College	0.153	0.148	0.166	0.170	0.180
Race (%):					
Black	0.754	0.802	0.808	0.828	0.838
Coloured	0.153	0.142	0.144	0.135	0.134
White	0.094	0.055	0.049	0.037	0.0274
Employed	0.550	0.468	0.520	0.582	0.571
Union	0.132	0.114	0.126	0.140	0.141
	Panel B: Household Characteristics				
Married	0.934	0.899	0.813	0.875	0.832
Household size	3.835	4.385	4.184	4.035	3.930
Number of children	2.285	2.369	2.277	2.328	2.228
Urban	0.575	0.543	0.566	0.579	0.580
Province (%):					
Western Cape	0.151	0.127	0.131	0.122	0.117
Eastern Cape	0.121	0.121	0.124	0.121	0.117
Northern Cape	0.083	0.079	0.077	0.079	0.074
Free State	0.066	0.063	0.067	0.064	0.062
KwaZulu-Natal	0.223	0.252	0.247	0.251	0.258
North West	0.069	0.065	0.066	0.065	0.063
Gauteng	0.146	0.131	0.137	0.149	0.155
Mpumalanga	0.076	0.083	0.074	0.074	0.075
Limpopo	0.067	0.080	0.076	0.075	0.079
Expenditure and Income (R 000s):					
Food consumption	14.062	13.594	12.681	13.235	12.583
Nondurable consumption	43.936	36.830	29.874	32.551	30.854
Total consumption	48.474	39.637	33.343	36.071	34.154
Income before tax	65.415	67.135	66.351	74.373	67.053
Disposable income	64.076	66.045	65.137	73.732	66.114
Observations	5,107	4,162	5,453	6,670	6,768

Notes: The table presents summary statistics across individual and household characteristics. The white percentage under race includes White and Asian headed households. *Data source:* South Africa National Income Dynamics Survey (NIDS), 2008-2017.

The share of households across provinces is relatively stable, with about 65 percent of households living in the Western Cape, East Cape, KwaZulu Natal and Gauteng provinces. The average expenditure on household food and nondurable consumption decreased during the period. Income slightly rose until 2014, after which it started declining. This decrease can be due to two potential reasons. As explained above, the attrition of White-headed households and high-income earners who usually have high income and expenditure such that excluding them pushes the mean, in real terms, down. Another potential reason is that the decrease reflects the actual trends in the economy in which household income is falling. As Fig. C-5 in the appendix shows, the GDP per capita has been falling in South Africa since 2013.

4.4 Cross-Sectional Micro Facts

4.4.1 A. Household Income Structure

Fig. 4.2 presents the dynamics in household income structure during the study period. The plot shows that the contribution of different income sources to household income depends on demographics. Income from earnings and public transfers contributes above 90 percent of the household income for rural households; earnings alone almost match this level for urban households. Public and private transfers are a significant source of income for rural households, while investment income matters more than private transfers for urban households. Private transfers almost contribute 10 percent of household income for rural households by 2017; thus, they are likely to play an important role in helping households smooth income shocks. Fig. C-3 in the appendix shows the share of income sources by household percentiles. Regardless of whether it is urban or rural, public

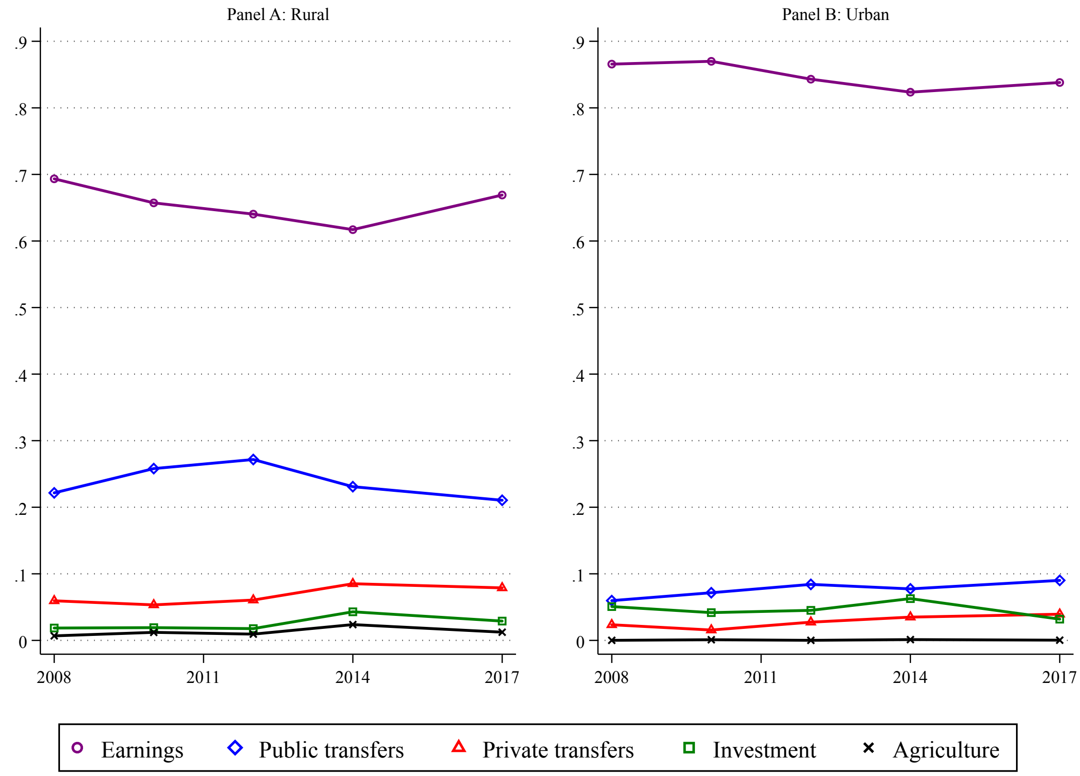

Figure 4.2: Household Income Shares by Source

Notes: Panel A shows trends in the share of earnings, investment, receipts from social assistance, agriculture and private transfers from family and friends in household income for rural households. Panel B depicts the same trends for urban households. The household income aggregates resources from the five income sources. The shares are computed from amounts in real Rand terms, deflated using a monthly CPI corresponding to the month and year of the interview.

transfers are the main source of income for income-poor households. For rural households below the 75th percentile, the contribution of public transfers has been slowly rising. Fig C-4 presents income shares by race. The results show that public and private transfers are more important for Black and Coloured households.

The share of households receiving zero labour earnings, with no one working, has been rising over time (see Fig. C-1). When labour income is elusive other sources of household income are likely to be increasingly important. One source of income when earnings fall to zero is borrowing from formal or informal channels. Fig. C-8 in the appendix shows that the proportion of households indicating that they borrowed in the last year has risen over time. The rise is even starker given that the share of individuals with a bank account or credit card slightly declined during the period. This trend can only be rationalized by the fact that those who borrow have increasingly relied on informal borrowing channels. Besides borrowing, households can also rely on private transfers from either household members living outside the household or family and friends.

Fig. C-7 in the appendix depicts source and destination of remittances by province of residence and the relationship between the remitter and the household. The data indicates that an individual household can receive transfers from up to eight individuals. The figure only shows information for the first three main remitters. Most remitters live in Western Cape, KwaZulu Natal and Gauteng, which is consistent with the fact that these provinces have relatively more job opportunities. Interestingly, KwaZulu Natal is a major remitter as well as a receiver of transfers. This trend might be because the province has both a significant rural and urban population. The figure indicates that, besides spouses living outside the household, the biggest senders of transfers are children and friends.

The significant role of children aligns with rural-urban migration, where children migrate in search of better opportunities and have to send money home once they start working. Fig. C-6 depicts a positive covariance between labour earnings and private transfers, implying that even households with employed individuals still receive remittances. However, there is a negative covariance between labour earnings and public transfers. This inverse relationship points to social assistance being well-targeted at low-income households (as expected because most transfers are means-tested).

The results echo a conclusion in Leibbrandt et al. (2011) who show that South Africa has more redistributive policies than many developing countries, including most Latin American countries. The figure also shows that the household can simultaneously receive income from private and public transfers. Many studies (Leibbrandt et al., 2011; Maboshe and Woolard, 2018; Leibbrandt et al., 2018) find social assistance programs to be mostly progressive even though their effects seem to be waning over time. Similarly, Case and Deaton (1998) show that most of the public transfers in South Africa are received by low-income households.

Although all these other sources are essential sources of income, labour income remains the primary source of income for many households. However, Fig. C-1 show that even when employed, a significant fraction (4 percent) of individuals earn less than their skill due to underemployment problems. The increasing share of zero labour earnings households and high unemployment is partially reflected in the aggregate average real income that has barely increased in South Africa since independence. Fig. C-5 in the appendix shows that increase in GDP per capita has been lagging behind that of other emerging economies in the BRICS (Brazil, Russia, India, China and South Africa) economic block even though the average rate of population growth is comparable for most of the countries.[12]

4.4.2 Labour Market Trends

Besides being already high, the level of unemployment (narrowly defined) in South Africa has been trending upwards since independence from the Apartheid government in 1994 (except in the early 2000s). Burger and von Fintel (2009) provide a detailed discussion of these trends, and Levinsohn (2008) discusses the role of Apartheid policies in shaping South Africa's

[12] Income per capita in South Africa has been falling since 2013.

current labour market. Panel A in Fig. 4.3 sheds insight on why unemployment remains persistently high regardless of the structural and affirmative action policies implemented since independence. The figure shows that labour-intensive industries (mining, manufacturing and construction) are contributing increasingly less towards total employment. On the other hand, Panel B indicates that the share of low-skilled jobs (significant share of individuals in the labour force are low-skilled) has fallen by 13 percent from 37 percent in 1994 to 24 percent by 2018. In comparison, employment in high-skilled industries rose from 35 percent to 45 percent during the same period.[13]

The low-educated and poor households are likely to bear the burden of these structural changes. Due to the high unemployment rate, the returns to labour market experience and tertiary education continue to rise in post-independence South Africa (Leibbrandt et al., 2011, 2018). In a market with excess labour, firms can be selective by preferring to employ only high-educated or more experienced individuals. These firms' decisions burden poor households, low-educated individuals and the youths (youth unemployment is above 35 percent). Another notable trend is the decrease in employment share for industries (mining, manufacturing and construction) with relatively high labour union intensity. [14]

When the labour market fails to provide individuals with employment opportunities, many households can not use labour supply via the

[13]The level of skill is defined according to the average level of education in these sectors. While this classification misses the high-educated individuals in manufacturing, mining and construction, e.g., middle to upper-level managers, such grouping provides an insight into the general trends in the labour market. It is also important to note that education is not an accurate measure of skills in the labour market.

[14]This observation does not imply causation as many other factors can also play a role, such as the transformation of the country into more of a tertiary sector-driven economy. Perhaps it can be argued that South Africa has struggled to utilize her abundant resource - human capital.

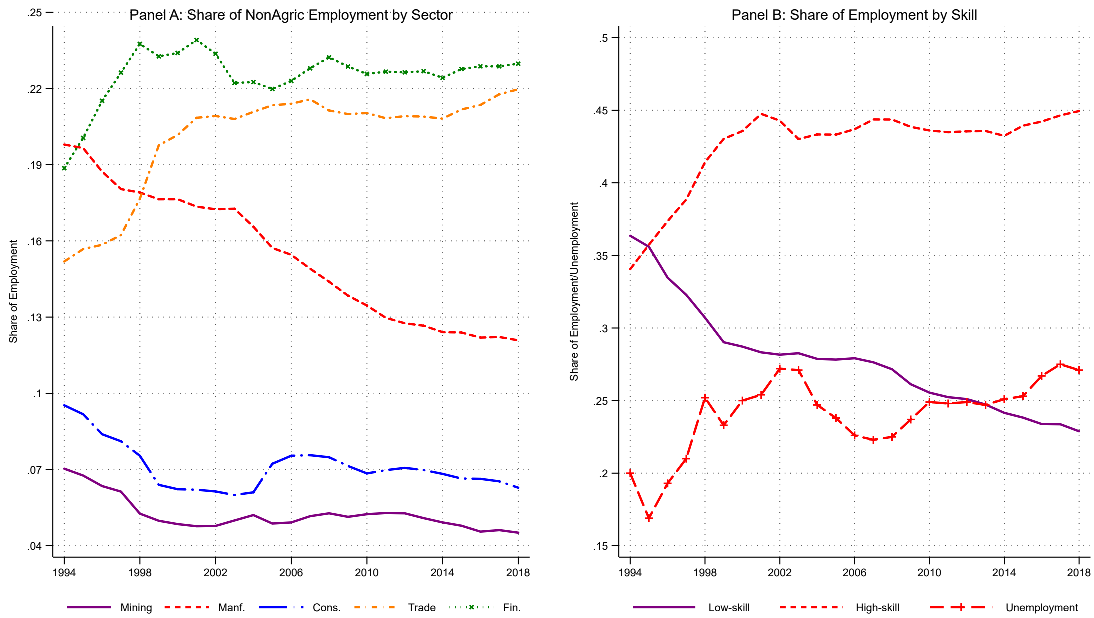

Figure 4.3: Employment Shares Across Sectors, Skill and Unemployment Rate

Notes: The non-agricultural employment data is from the South African Reserve Bank (SARB) historical data query portal published as series KBP7002J-KBP7009J. The data series are available as indices accompanied by 2010 employment numbers, which can be used to compute the level of employment. In Panel A, the trends show the contribution of each sector to the total level of non-agricultural employment. The two trends on top (in 1994) in Panel B depict the share of employment based on the average level of skills (education) required to work in the sector. The low-skill employment line traces the share of employment in mining, manufacturing and construction, while high skill is for employment in the trade and financial sectors. The bottom line depicts the trend in the official unemployment rate (narrow definition).

added-worker effect or changes in the intensive margin as a device to smooth consumption. Because family labour supply is one of the most effective sources of consumption insurance, as noted in Blundell et al. (2016), lack of job opportunities means these households have to rely on other insurance mechanisms.[15] With the level of financial inclusion and possession of assets to use as collateral still low, it means the financial sector as a source of insurance is still inaccessible to many people. Given that the financial sector and labour market may be out of reach for many, domes-

[15]The authors show that family labour supply is the primary source of insurance against income shocks in the U.S.

tic private transfers and private transfers ought to be critical in helping families smooth shocks.

4.4.3 Residual Income and Consumption Inequality

This section discusses the evolution of residual inequality in log adult-equivalized disposable income and nondurable consumption. The income residuals are from a regression of log adult-equivalized disposable income on education, race, sex, number of children, household size, province of residence, type of household settlement (formal/informal), area of residence (rural/urban) and age. Consumption residuals are from a regression of log equivalized nondurable consumption on the same demographics. This section focuses on variances of the residuals and percentile ratios. Panels A-D in Fig. 4.4 show trends in variances, 50^{th}-10^{th} (P50-P10), and 90^{th}-50^{th} (P90-P50) percentile ratios. Panels A and B depict trends in consumption, and Panels C and D show the same trends for disposable income. The variances and ratios are normalized to zero in 2008 to clearly illustrate the inequality changes between the two periods.

On average, consumption inequality increased between 2008 and 2017 (Panel A). Residual and raw consumption inequality increased by 0.2 and 0.1 log points, respectively. In comparison, panel C indicates that residual and raw income inequality fell by 0.2 and 0.5 log points, respectively. This observation is different from findings in Hundenborn et al. (2018) and Leibbrandt et al. (2018), who report that the level of income inequality remains relatively stable and high in South Africa. Leibbrandt et al. (2018) credit the unresponsiveness of income inequality to the labour market that continuously favours highly educated individuals and the declining effectiveness of social assistance. The significant difference between

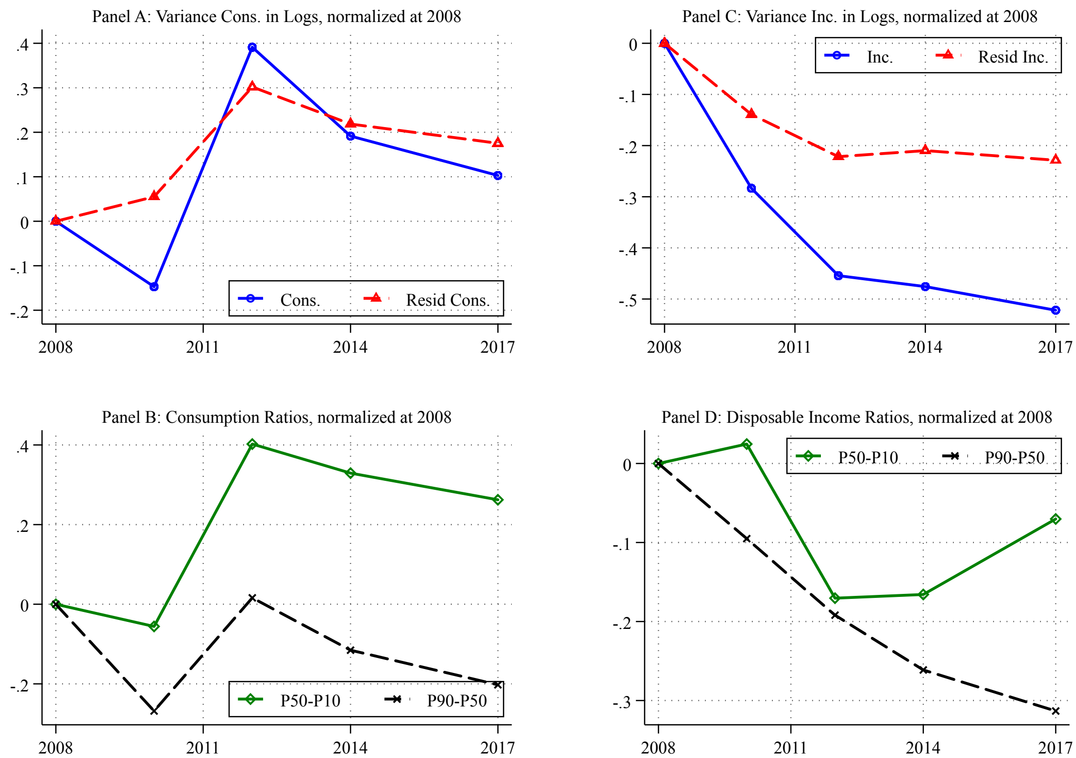

Figure 4.4: Residual Inequality in Equivalized Income and Consumption

Notes: The figure depicts variance of residuals (after controlling for household head and household observable characteristics) for equivalized disposable income and consumption using OECD scales (1 + 0.7 × (# adults - 1) + 0.5× # children). Panel A and C show trends in inequality using the variance of logs of nondurable consumption and disposable income, respectively, while Panel B and D depict inequality across the distribution (50-10 and 90-50 percentile ratios). Panels A and C show both inequality in raw variables and residuals.

raw and residual inequality reinforces the importance of controlling for observable characteristics.

Panel B and D provide insights into the distributional shifts in inequality over the period. The trends indicate more mobility at the top than at the bottom of the distribution. This difference is consistent with Hundenborn et al. (2018) who note that the labour market continues to benefit the highly skilled and more educated groups. The results show that while consumption inequality fell at the top of the distribution, it rose at the bottom. These results differ from the evidence from other countries such as

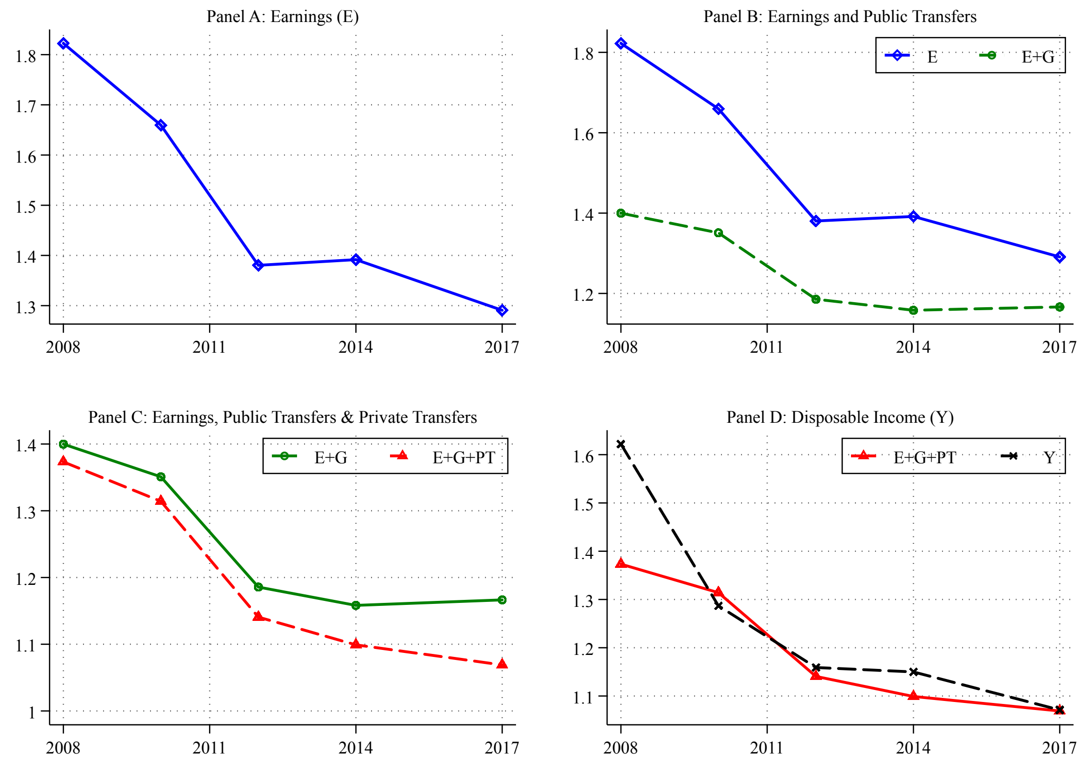

Figure 4.5: Income Inequality at Different Aggregation Levels

Notes: The plots depict income dynamics at different aggregation levels with the labels referring to E- Household earnings, PT- private transfers, and G - social assistance payments. +X means the previous income component plus the subcomponent X. Panel D includes invest and agriculture income received by the household. The figure presents income components at household level using NIDS data. The panels depict variance of income logs.

the United States (Krueger and Perri, 2006; Blundell et al., 2008; Aguiar and Bils, 2015), China (Santaeulàlia-Llopis and Zheng, 2018) who report a rise in both consumption and income inequality. However, it is essential to note that these countries still have levels of inequality far below that of South Africa.

Instead of focusing on total household income, it can be informative to consider the different components and their contribution to overall inequality, as shown in Fig. 4.5. Both private and public transfers (social assistance) are critical sources of income that reduce income inequality.

However, public transfers are more effective at reducing income inequality than private transfers. For instance, comparing Panels B and C, private transfers tend to lower inequality by, on average, 0.02 log points while private transfers decrease by 0.4 log points. The less effectiveness of private transfers in reducing inequality is likely because any household can receive funds (including the well-off). At the same time, public transfers target families at the bottom of the income distribution. Although public transfers seem more effective at reducing inequality, it does not necessarily imply effectiveness at insuring households against shocks. The following section estimates the degree of insurance of different income sources, including private and public transfers.

4.5 Joint Evolution of Consumption and Income

To measure the magnitude of shocks, econometricians often decompose income shocks into persistent and transitory components. This decomposition is under the assumption that the household income process is a combination of a persistent component (sometimes assumed to be a random walk) and a transitory component. The chapter follows Blundell et al. (2008) in assuming that the persistent component follows a random walk. By following this approach, the study abstracts from directly considering labour supply or other heterogeneities (as in Heathcote et al., 2014; Blundell et al., 2016). The measure of income and consumption is the residuals (unexplained component) from regressing the log of equivalized income and consumption on education, race, province of residence, employment status, number of children, cohort and year dummies, and interactions of education, race, family size, employment status and number of children with year. The residuals from these regressions are the variables of interest in the estimation below. In the model, income is a combination of permanent and transitory components that can be represented as:

$$\log Y_{i,t} = \mathbf{X}'_{i,t}\beta_t + z_{i,t} + \epsilon_{i,t} \tag{4.1}$$

where $\mathbf{X}$ controls for observable income characteristics for household i in period t, $z_{i,t}$ is the permanent component and $\epsilon_{i,t}$ is the transitory component. The characteristics $\mathbf{X}$ include the factors mentioned above. The permanent component follows a random walk:

$$z_{i,t} = z_{i,t-1} + \eta_{i,t} \tag{4.2}$$

where $\eta_{i,t}$ is serially uncorrelated idiosyncratic shocks that are orthogonal to the transitory component $\epsilon_{i,t}$. Thus, both shocks are independent and identically distributed across households and waves. Blundell et al. (2008) assume that the transitory risks follow a moving average of order one MA(1). Because the NIDS data is biennial, it is not possible to identify the MA(1) parameter.[16] The only difference between the definition of the permanent component in Eqn. 4.2 and the one in Chapter 3 above is that I assume that these shocks are genuinely permanent; thus, $\rho = 1$.

Using the residuals, the growth in the unexplained component of income becomes

$$\Delta^2 y_{i,t} = \eta_{i,t} + \Delta^2 \epsilon_{i,t} \tag{4.3}$$

where $y_{i,t} = \log Y_{i,t} - \mathbf{X}'_{i,t}\beta_t$ is the log of income net of predictable observable characteristics and Δ^2 means two year growth (biennial data).

Similarly, Blundell et al. (2008) show that the growth of residual consumption can be specified as:

$$\Delta^2 c_{i,t} = \phi\eta_{i,t} + \psi\epsilon_{i,t} + \xi_{i,t} + \Delta^2 \zeta_{it} \tag{4.4}$$

[16]The transitory component can still follow an MA(1), but the parameter cannot be recovered with biennial data. Implicit in the specification is the assumption that households' transitory component does not follow a process higher than MA(1).

where $\Delta^2 c_{i,t}$ is the growth of consumption between period $t-2$ and t, $\eta_{i,t}$ and ϵ_{it} are permanent and transitory shocks to household income, respectively, $\xi_{i,t}$ captures shocks to consumption that are orthogonal to those in disposable income and $\zeta_{i,t}$ captures any measurement error in consumption. The parameters ϕ and ψ capture the transmission of permanent and transitory shocks to consumption, respectively. In this model, these two parameters can also be used to derive the level of consumption insurance against permanent $(1-\phi)$ and transitory $(1-\psi)$ income shocks. The magnitude of insurance gives an indication of the extent to which households encounter fluctuations in their income without having to adjust their consumption.

Given $\Delta^2 c_{i,t}$ and $\Delta^2 y_{i,t}$, the insurance parameters (ϕ, ψ), variances of permanent and transitory shocks $(\sigma^2_\eta, \sigma^2_\epsilon)$, variances of consumption growth and measurement error $(\sigma^2_\xi, \sigma^2_\zeta)$, are recovered by matching the autocovariances and cross-covariances of income and consumption against the theoretical moments. The optimization process is via diagonally weighted minimum distance (DWMD), that is, minimizing the weighted distance between data moments and the moments from the model.

4.6 Results

4.6.1 Partial Insurance Heterogeneity

This section explores the level of insurance across different cohorts to show that there is heterogeneity in terms of households' preparedness to smooth consumption. The square dots in Panel A of Fig. 4.6 depict the transmission of permanent shocks (ϕ), with $1-\phi$ as the magnitude of insurance and the bars around each dot represent 95 percent confidence intervals. Similarly, Panel B presents the transmission of transitory shocks (ψ). The results are across five groups; the baseline sample, education, region of residence,

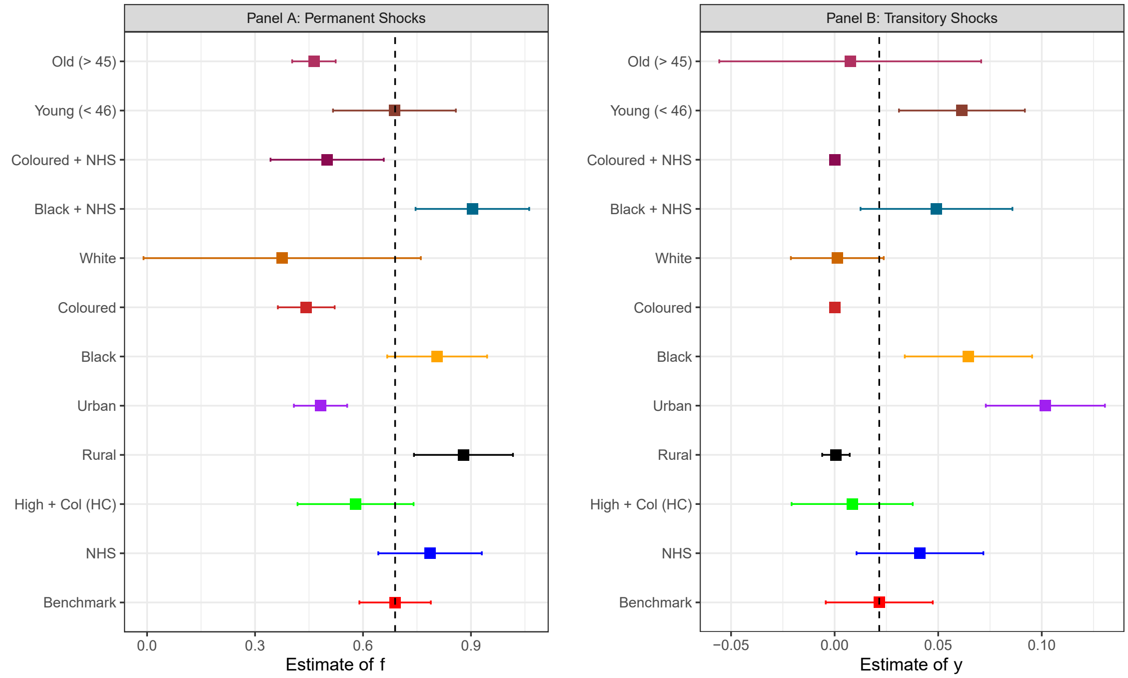

Figure 4.6: Minimum Distance Estimates: Transmission of Permanent and Transitory Shocks

Notes: The panels depict diagonally-weighted minimum distance (DWMD) estimates for different sub-samples, and the dots show the pass-through of permanent and transitory income shocks. Tables with these estimates and standard errors are in the Appendix. The benchmark estimates represent the transmission of shocks from disposable income to nondurable consumption for the full sample (continuously married or single) and the other estimates are for different cohorts of education, region of residence, race and age. NHS represents the sub-sample of individuals with less than high school and Black + NHS means an individual who is Black and has less than high school education (same for Coloured + NHS). The bars around coefficients represents 95 percent confidence intervals. To overcome the problem of some subsamples being too small, the standard errors and confidence intervals are generated via bootstrapping with 500 replications.

race and age of household head.

Except for the cohort of White-headed families, all the permanent estimates are statistically significant (at least at 5 percent) while most of the transitory estimates are insignificant even at 10 percent. First, the benchmark estimate for ϕ is 0.69, which means for every 10 percent movement in household disposable income, nondurable consumption changes by 6.9 percent. Alternatively, the estimate indicates that, on average, households smooth about 31 percent of permanent income shocks. The estimate is close to Blundell et al. (2008) who report the level of insurance in the U.S. at 36 percent. However, there are two factors that differentiate this study from Blundell et al. (2008). The analysis includes both married and single households who are between the age of 25 and 65 years, whereas Blundell et al. (2008) restrict their sample to only continuously married male heads between the ages of 30 and 65 years. Nevertheless, this level of insurance points to partial insurance against permanent income shocks in South Africa. However, the magnitude is far below that of Chinese households of 70 percent (reported in Santaeulàlia-Llopis and Zheng (2018)). The estimates for transitory shocks are in the range of 0 to 0.10, which implies insurance above 90 percent with full insurance for some groups (White, Coloured, Older cohorts). However, these estimates are mostly statistically insignificant even at 10 percent, which means the hypothesis of full insurance against transitory shocks cannot be rejected (except for Black-headed and urban households). [17]

The results suggest that the level of insurance against permanent shocks is positively correlated with the household head's education. Because the group of individuals with a college degree is small in the sample, the estimates combine those with high school and above against individuals with no schooling or high school dropouts. While only 57 percent of permanent shocks passthrough to consumption for households whose head has at least high school attainment, the magnitude is 19 percent higher (76 percent) for those with less than high school. These results confirm the assertion in Leibbrandt et al. (2018) that there are higher returns to being educated in South Africa, especially with the persistent lack of equal

[17] Although a negative estimate, as suggested by some of the confidence intervals in Fig. 4.6, makes no theoretical sense. A test of whether they are significantly different from zero fails to reject.

opportunities. This disparity is most likely due to the high level of unemployment such that individuals with tertiary education stand a better chance of securing employment. Such results raise a gloomy picture for individuals with low education, especially after Blundell et al. (2016) provide evidence that family labour supply is the most effective tool that most households use to smooth consumption in the U.S. Given that the South African government is currently evaluating the merits of offering free college education, these results might be encouraging in that college holders become better at smoothing shocks.[18]

As the conjecture in Section 4.4.2 alludes, the burden of South Africa's sluggish growth falls mostly on those who struggle to access job opportunities in the labour market. The sub-sample of low-educated (NHS) individuals who are Black has only 10 percent insurance against persistent risks. This level of insurance compares to about 50 percent for the same group, but from Coloured families.[19] If the sample is further restricted to Blacks who are younger household heads (less than 45 years old) with less than high school, they have zero insurance against these shocks (estimate not reported in the figure). With the non-college youth (15-24 years old) unemployment rate hovering above 50 percent, the slow growth and lack of opportunities are exerting a larger burden on the low-educated.

In terms of region of residence, Fig. 4.6 shows that about 48 per-

[18]The country experienced a wave of protests across universities in 2015 and 2016, in which students were demanding the introduction of free tertiary education. The South African Student Aid Scheme (NASFAS) already provides university education loans to low-income households. However, the main grievance of the students' protests was that the coverage of NASFAS was inadequate in that many middle-income families who are struggling to pay do not qualify for the loans. See http://www.nsfas.org.za/content/mission.html. Also, the assertion that more education ensures households can smooth shocks is speculative in this case as the current analysis cannot assess the dynamic effects of everyone in the country becoming increasingly more educated.

[19]The coefficients for the White-headed households with less than high school is not reported because the sample is too small to get precise estimates.

cent of permanent income shocks pass-through to consumption in urban areas and the magnitude is almost double for rural households (88 percent). These magnitudes imply that urban and rural households have about 52 percent and 12 percent insurance against permanent income shocks, respectively. These findings differ from Notten and de Crombrugghe (2012) and De Magalhães et al. (2019) who report higher insurance for rural households in Russia and Malawi, respectively.[20] Although relatively more private transfers go to rural areas (55 percent), the results seem to suggest that such transfers are not enough to equalize the ability of rural households to smooth consumption as those in the big cities. Another possible reason for the differences is that rural households (36 percent) in South Africa are not as intensively reliant on farming as a source of income or insurance as those in Malawi and Russia.

The results also indicate the potential persistence of Apartheid policies that still contributes to skewed labour market opportunities decades after they were abolished. In Black-headed households, 81 percent of the permanent shocks affect consumption. The level of transmission is lower for Coloured-headed households (44 percent) and white-headed households (37 percent).[21] Fig. C-4 in the appendix shows the difference in wealth across the race groups. While financial income, which points to a household holding wealth that generates the income, constitutes about 20 percent of household income for White-headed households, the share is almost zero for Black-headed households with Coloured-headed families somewhere in between. Thus, Black families face a double jeopardy of lacking

[20] Gerry and Li (2010) find the opposite results for Russia in that urban households had more insurance.

[21] Coloured refers to individuals of mixed race and this term is predominantly used in South Africa. The White group includes individuals of Indian descent to increase sample size and also due to the fact that most of these families are in the top income and wealth quantiles in South Africa.

wealth to draw from in times of income shocks as well as facing higher unemployment in the labour market.[22]

Lastly, the findings indicate that young households have less insurance (31 percent) in comparison to older households (54 percent). These results are consistent with the theory that older families would have accumulated more assets, thus giving them more resources to smooth risks than younger households. Such life cycle trends are well documented in the literature (See Kaplan and Violante, 2010; Karahan and Ozkan, 2013).

4.6.2 The Role of Private Transfers, Public Transfers and Taxes

This section discusses the role of taxes, private and public transfers and the progressivity of the tax system in South Africa. First, to assess the progressivity of the tax system, the chapter follows a simple functional form in Bénabou (2002), Guner et al. (2014) and Heathcote et al. (2014). Specifically, the tax system can be summarised by:

$$t(\tilde{y}) = 1 - \lambda \tilde{y}^{-\tau}$$

where $t(\tilde{y})$ is the average tax rate in the data and $\tilde{y}$ is a multiple of mean household income. As defined in Guner et al. (2014), $\tilde{y} = 2$ corresponds to the average tax rate that is paid by an individual who earns twice the income of the mean household income in the data. The parameter τ measures the degree of progressivity in the tax schedule such that $\tau = 0$ implies a proportional tax rate and $\tau > (<)0$ signifies a progressive (regressive) tax system. The parameter λ measures the level of taxes.

Using the benchmark sample, τ is positive at 0.0442 with a standard error of 0.0066, which indicates that the economy has a progressive

[22]See Gradin (2019) for a discussion of the extent of disaggregation in the South African labour market.

tax system.[23] Fig. C-9 in the appendix depicts parameters estimated by year. All the estimates of τ are positive during the period. A question remains whether this progressive tax system plays some role in helping households to insure against income shocks and such a question is an empirical pursuit. Therefore, Fig. 4.7 provides evidence of the role of the tax system, financial income and private and public transfers. The benchmark estimates are the same as in Fig. 4.6, but are included as a point of comparison. The estimates for permanent shocks are all statistically significant at 1 percent, while all the coefficients of the transmission of transitory shocks are insignificant (except for income excluding private transfers).

The results indicate that the tax system indeed plays some role in helping households to mitigate the impact of permanent shocks (0.69 versus 0.58). When household disposable income is inclusive of financial income, the transmission of shocks falls from 69 percent to 62 percent. Although financial income is a source of insurance, the evidence shows that the tax system contributes relatively more to consumption insurance (decrease from 69 percent to 59 percent). Perhaps, the reason for this low protection from financial income is that only a few households have wealth that generates the income as discussed above.

Interestingly, the results suggest that private transfers are more critical in helping households to reduce the impact of permanent shocks (a drop of 15 percent).[24] In the absence of private transfers, families struggle even to insure against transitory shocks with 7 percent of these risks

[23] The parameters are estimated using ordinary least squares and $\lambda = 0.9953$ with a standard error of 0.0006.

[24] To see that private transfers play a role, note that consumption still includes insurance from private transfers. However, private transfers have been excluded in household income; thus, the effect is that there should be less transmission of shocks to consumption. The decrease gives insights into the extent of insurance being contributed by private transfers. The same applies to the other categories in the figure.

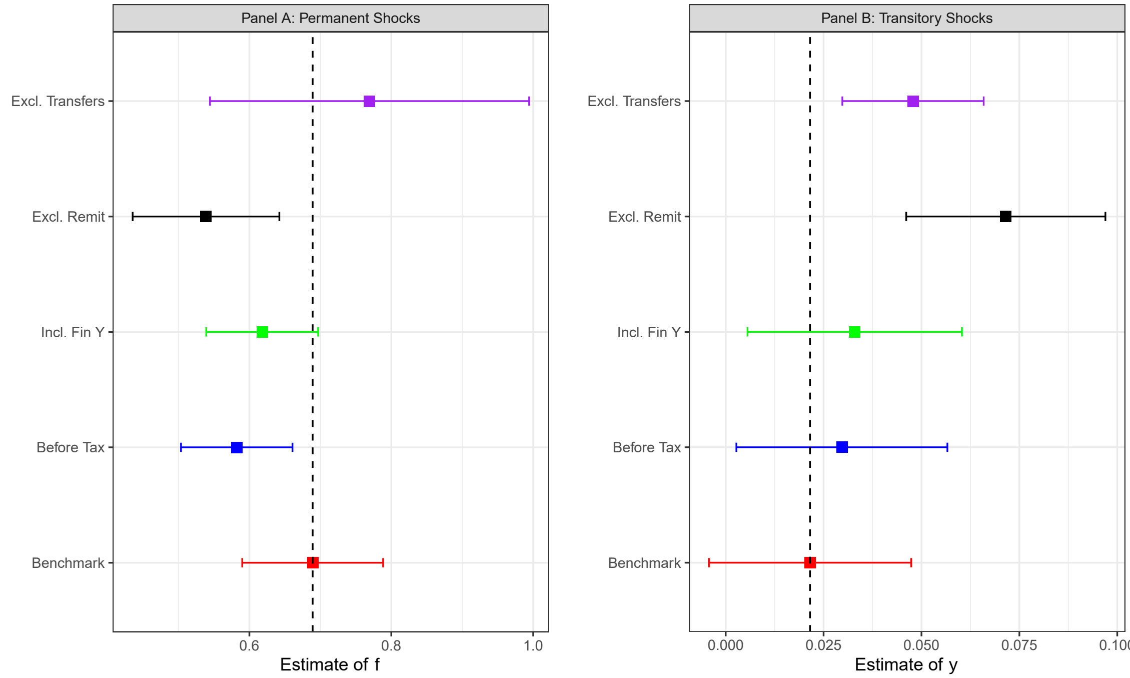

Figure 4.7: Minimum Distance Estimates: Transmission of Permanent and Transitory Shocks

Notes: The panels report the transmission (point estimates) of permanent and transitory shocks computed using diagonally weighted minimum distance (DWMD). Tables with standard errors are in the Appendix. The benchmark estimates represent the transmission of shocks from disposable income to nondurable consumption for the full sample (continuously married or single). Excl. transfers means disposable income net of financial income and social assistance, excl. remit means disposable income net of financial income and private transfers. The bars around coefficients represents 95 percent confidence intervals. Standard errors and confidence intervals are generated via bootstrapping with 500 replications.

affecting consumption. Given that a significant share of private transfers are often from family members (Fig. C-7), these results differ to findings in Altonji et al. (1996) who report no evidence of within-family insurance in the U.S.[25] On the other hand, there is no evidence that public transfers help in consumption smoothing.[26] The results from Canada (Brzozowski

[25]Schulhofer-Wohl (2011) argue that the earlier studies that failed to find evidence of risk-sharing in the U.S. ignored the possibility that individuals might self-select into different occupations depending on their level of risk tolerance. Using data from the Health and Retirement Survey, Schulhofer-Wohl (2011) finds the evidence of this self-selection and show that once this correlation is controlled for, there is a considerable level of within-family risk-sharing in the U.S. (consumption growth was less sensitive to shocks).

[26]Public transfers in South Africa include child support grants, old aged grants, disability support grants and foster care grants.

et al., 2010) and Norway (Blundell et al., 2015) show that the tax system plays an important role in minimizing the pass-through of income shocks to consumption in those countries. However, Brzozowski et al. (2010) report that public transfers were also effective, which differs from the results in Fig. 4.7.

The smaller role of social assistance is somewhat surprising. In fact, the magnitude of the coefficient relative to the benchmark estimate is at odds with theoretical expectations.[27] Because social assistance is disbursed on a means-test-basis such that they target low-income families that need help, it should play a more prominent role in insurance than private transfers, which is not the case. Perhaps the fact that the transfers are only paid at the end of the month make them less effective at mitigating impacts that occur during the month, especially for hand-to-mouth households, while private transfers can be received at any time of the month. An important observation from these results and the discussion in Section 4.4.3 is that private transfers are more effective at providing households with insurance while public transfers is relatively more effective at reducing the level of income inequality. Therefore, these two sources of income are important for households.

4.6.3 Different Components of Consumption

The analysis above focuses on nondurable consumption and it might be informative to consider consumption at different levels of aggregation. Because food expenditure is a necessity, the prediction is that households should be able to smooth food more than other broader consumption categories. If not, that is a worrisome *status quo*. The lack of consumption

[27] As explained in footnote 24, the insurance parameter after excluding public transfers (77 percent) is supposed to be less than the benchmark (69 percent). The wrong direction of this coefficient is puzzling and no immediate explanation comes to mind.

insurance against food would imply that some families are forced to cut food expenditure when they experience shocks. Reducing food expenditure for households that are already at subsistence has human capital development implications, especially for children.

The literature (Jamison, 1986; Themane et al., 2003) has shown that children that come from low-income households and are malnourished tend to relatively perform less at school. Another reason for studying pass-through at different levels is that health and education expenditure should be excluded in computing nondurable expenditure because they rather capture the cost of investment in human capital. Alternatively, the exclusion is justified on the account that the utility benefit from these two categories varies widely across the life cycle (see Attanasio et al., 2005; Aguiar and Hurst, 2013). Therefore, Fig. 4.8 depicts the magnitude of the transmission of income shocks at different levels of consumption, that is, for food, nondurable excluding health and education expenditure, necessities and total consumption (including durable expenditure).[28]

The estimates indicate that families, on average, have more insurance for food (59 percent) or the broad category of necessities (49 percent). In comparison to nondurable consumption, they almost have double the level of insurance for food. Blundell et al. (2008) and Casado (2011) find similar results for the United States and Spain, respectively. For nondurable consumption net of health and education, households have slightly more insurance (37 percent) than the benchmark level of consumption. This result seems to suggest that health and education expenditure helps households to smooth more shocks. The potential reason for this finding in South Africa is that primary health care and high school education can

[28]Necessities include expenditure on food, utilities and rent.

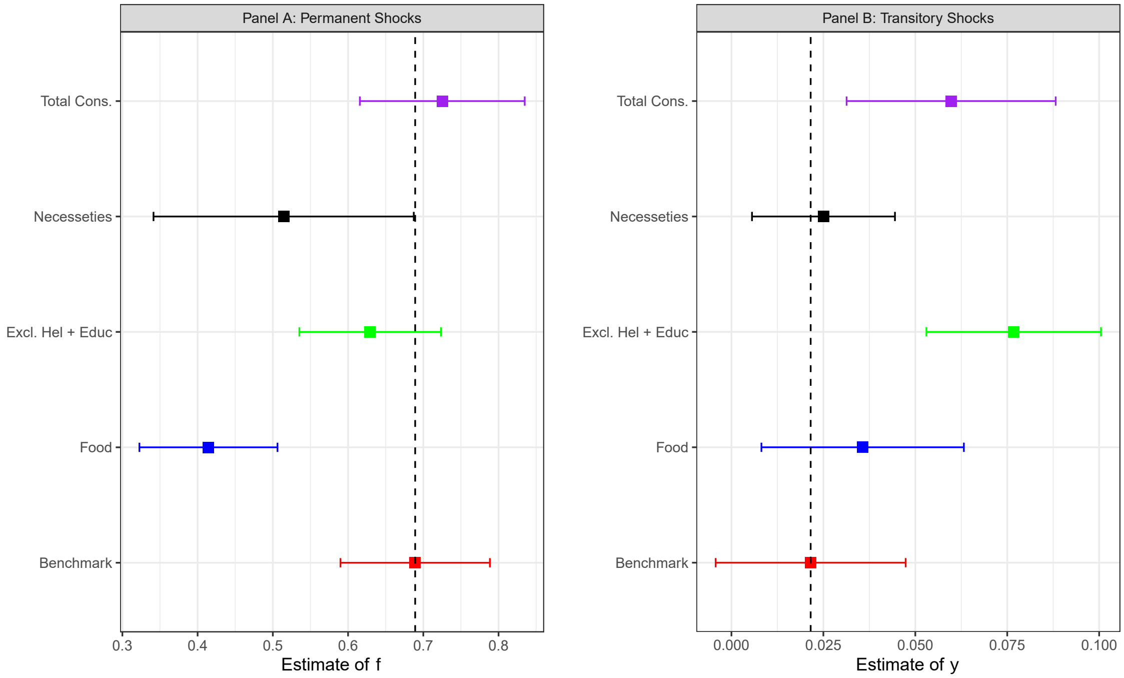

Figure 4.8: Minimum Distance Estimates: Transmission of Permanent and Transitory Shocks

Notes: The panels report the pass-through (point estimates) of permanent and transitory shocks computed using diagonally weighted minimum distance (DWMD). Tables with standard errors are in the Appendix. The benchmark estimates represent the transmission of shocks from disposable income to nondurable consumption for the full sample (continuously married or single). Necessities include expenditure on food, utilities and rent expenditure, while total cons. is the sum of nondurable and durable expenditure. Standard errors and confidence intervals are generated via bootstrapping with 500 replications.

be accessed for free if households choose to visit public clinics, hospitals and schools. More so, medication for some chronic illnesses such as TB and HIV are accessible free of charge. The fact that households can access healthcare and their children attend school even without income means that it partly helps them to avoid adjusting consumption than they were likely to do. For instance, a health shock in the United States to a family with no health insurance can be catastrophic for a household by drenching them into debt, or they intentionally avoid visiting a hospital to avoid the income risk. Given that South Africa is in the process of introducing a National Health Insurance Scheme modelled on a structure similar to that of the United Kingdom is somewhat encouraging in that the results point to the benefits of households having access to state-funded healthcare.

To summarize, the results indicate that South African households can partially insure against permanent shocks. These results are consistent with the evidence in Blundell et al. (2008), Kaplan and Violante (2010), Karahan and Ozkan (2013), Heathcote et al. (2014), De Nardi et al. (2019) for the U.S., Casado (2011) for Spain and Santaeulàlia-Llopis and Zheng (2018) for China. Also, the findings point to heterogeneity in the extent of insurance across the region of residence, education, race, education and age. In general, the estimates indicate that there is no enough evidence to reject the presence of full insurance against transitory income shocks.

4.7 Robustness Checks

This section explores the sensitivity of the above results in different subsamples. Fig. 4.9 depicts the transmission parameters for married couples only, married couples who were continuously married during the sample period, balanced sample (observations in all five surveys), male heads only and those aged between 30 and 65 years. The transmission of permanent shocks for all the subsamples (except the balanced sample), is close to the benchmark estimates. The cohort of households that were in marriage has relatively less insurance. Married couples have less insurance than the benchmark (21 percent vs 39 percent), which is puzzling. The expectation is that married couples should have more insurance through the role of family labour supply (Blundell et al., 2016). There are a few potential reasons for these odd results. First, family labour supply is likely to play a minimum role in consumption smoothing when unemployment is high. In an economy with high levels of unemployent, it is difficult for the other

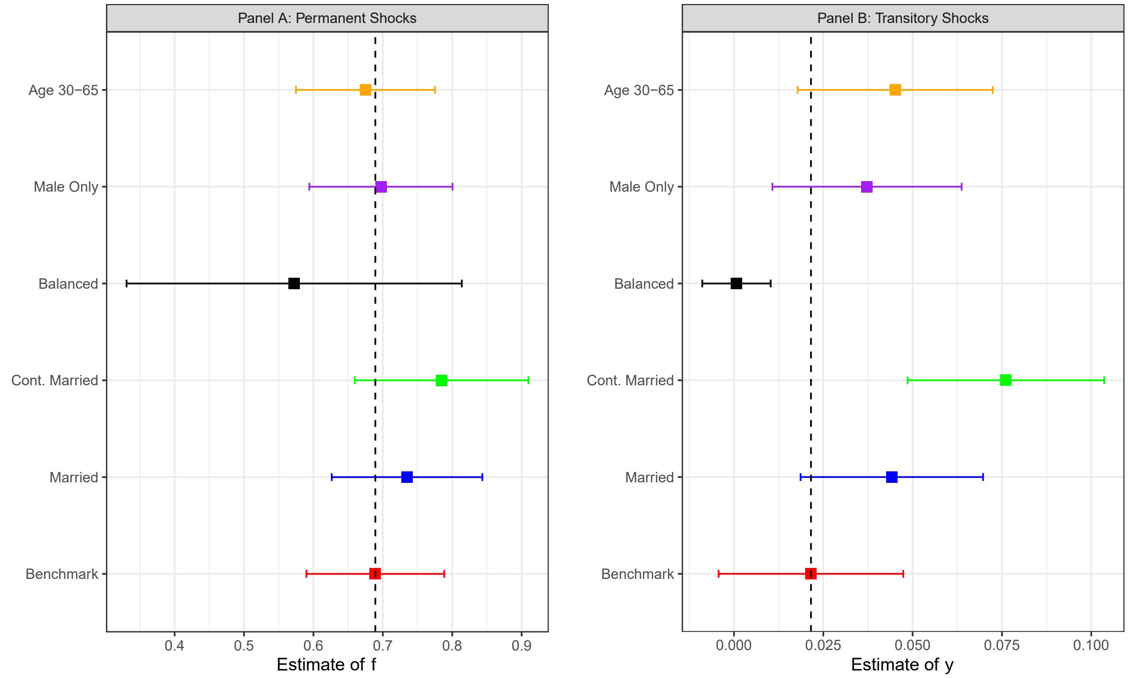

Figure 4.9: Minimum Distance Estimates: Transmission of Permanent and Transitory Shocks

Notes: The panels report the pass-through (point estimates) of permanent and transitory shocks computed using diagonally weighted minimum distance (DWMD). Tables with standard errors are in the Appendix. Cont. married means the subsample of households heads who were continuously married during the sample period. Standard errors and confidence intervals are generated via bootstrapping with 500 replications.

household members to adjust their labour supply when there are limited opportunities in the labour market.

Another reason could be a case where couples work in sectors that tend to have correlated shocks, such that they experience income changes at the same time, and that amplifies the exposure. In South Africa, a significant share of public transfers go to child support grants. The beneficiaries are predominantly single-parent households. It means these single-parent households receive more government support relative to married couples, which might explain why the latter have lower insurance. Nevertheless, the closeness of Fig. 4.9 estimates to the previous results discussed above gives an assurance that they are robust to sample selection.

4.8 Conclusion

As people age across the lifecycle, they face shocks that can force them to reduce consumption below subsistence. However, households can mitigate the impact of permanent or transitory shocks by using insurance devices (savings, borrowing, welfare payments). These insurance devices can facilitate the reduction of consumption inequality by shielding households from some shocks. This chapter investigates the degree of consumption insurance against income risk in South Africa using panel data from 2008-2017. When resources to smooth shocks are lacking, consumption volatility closely tracks the variations in income over time. Therefore, the chapter also assesses the role of different insurance mechanisms such as welfare payments, private transfers, and the tax system in smoothing income shocks.

The baseline estimate indicates that 69 percent of the permanent shocks lead to changes in consumption. In other words, households have 31 percent insurance against permanent income shocks. This level of insurance is close to the estimate in Blundell et al. (2008) (36 percent), but less than those reported in Casado (2011) (52 percent) and Santaeulàlia-Llopis and Zheng (2018) (above 70 percent). I fail to reject the hypothesis that households have full insurance against transitory shocks. There is considerable heterogeneity regarding the degree of insurance across education, area of residence and race. On average, educated and urban households have more insurance than their less educated and rural counterparts. Black-headed households have relatively lower insurance for both shocks. While private transfers seem to play a crucial role in consumption smoothing, social programs are more effective at reducing the level of consumption inequality. The tax system is progressive and helps families shield consumption from income shocks. The evidence that public transfers are less effective at providing consumption insurance but critical for reducing income inequality requires more research. The government should also promote private firms that offer instant mobile transfers, given that the results above indicate the importance of private transfers in offering families access to consumption smoothing.

Chapter 5

Concluding Remarks

5.1 Summary of Findings

This chapter concludes the study by summarising the key research findings in Chapters 2-4. The focus of this dissertation is to investigate issues related to the labour market from a macro perspective. The thesis begins by analyzing the causes of a widening employment gap (employment inequality) between low-educated and high-educated prime-age men in Canada. The evidence shows that the employment rate of low-educated men decreased by more than 25 percent between 1975 and 2020, while that of high-educated men fell by only 9 percent. This divergence in employment rates is the focus of Chapter 2. The chapter compares employment rates between 1990 and 2019, assuming that these two periods correspond to different steady states in the labour market.

The chapter then builds a search and matching model incorporating heterogeneous ability, education choice and variable job search effort. In the model, workers are either employed or decide on the degree of search effort to exert when they are nonemployed. The model accommodates the role of five channels: labour demand, labour supply, search frictions, search effort and duration dependence of job finding rates. I calibrate the model to match the rise in employment inequality between 1990 and 2019. The results from the calibrated model indicate that shifts in labour supply have been the primary driver in propelling the increase in employment inequality among prime-age men in Canada. On the other hand, substantial improvements in labour market efficiency for low-educated men helped to reduce the employment gap. The other channels also slowed the increase in inequality.

Chapter 3 provides a catalog of biases associated with estimating canonical income process parameters using quasidifferences. The chapter assumes that the income process is a sum of permanent and transitory components and fixed effects. Most studies estimate the parameters by minimizing the distance between model and data moments using either income levels or differences. The chapter focuses on a quasidifferences method that is rarely used in the literature but is simple to implement. However, little is known about the accuracy of this method's estimated parameters. The chapter uses Monte Carlo simulations to show that the degree of persistence can be precisely recovered using equal-weighting only when the variance of permanent shocks is bigger than that of transitory shocks. The other results are that the variance of fixed effects is upward-biased when the persistence is high, while the variances of permanent and transitory shocks are estimated with small biases, especially when the number of sample individuals is large. The chapter warns against using the quasidifferences method, except in the case above.

The last chapter examines the degree of consumption insurance in South Africa and the insurance devices households use to cushion consumption from income shocks. The premise of the chapter is that the extent to which consumption varies in the event of a household income shock depends on the family's ability to cushion the income loss. The chapter focuses on the ability to insure permanent and transitory shocks as in Blundell et al. (2008).[1] The chapter also documents the degree of heterogeneity in the ability of households to insure consumption across education, race, area of residence and age of household heads.

On average, the results show that households insure about 31 per-

[1] Other studies that use the same methodology include Kaplan and Violante (2010),Casado (2011), Santaeulàlia-Llopis and Zheng (2018), Hryshko and Manovskii (2022).

cent of permanent income shocks and have full insurance against transitory shocks. High-educated and urban households have more insurance than their less-educated and rural counterparts. Black-headed households have relatively less insurance for both shocks. The evidence also shows that private transfers provide significant insurance in the event of income shocks.

5.2 Areas for Further Research

Chapter 2 presents evidence that shifts in labour supply have been the primary driver of the increase in employment inequality. The model is limited because labour supply is modelled by catch-all parameters that summarize the overall shift over time. This approach does not disentangle the different components, e.g., unemployment insurance, the value of leisure, disability insurance, home production, etc., that affect labour supply. There is a need for further research to precisely identify the main drivers of these shifts in labour supply. A more nuanced understanding of the shifts in supply improves decision-making and potential policy response.

Chapter 4 utilizes biennial panel data from South Africa for 2008-2017 to measure consumption insurance. Such analysis is crucial, especially in South Africa, where unemployment levels are exceptionally high. However, the challenge is that the analysis period is short. Most developing countries lack panel income and consumption data, and this information scarcity hampers research on labour market dynamics. One would require access to administrative data to better understand labour dynamics and their link to consumption. Given that private transfers seem more effective at providing insurance, it might also be interesting to analyze the relationship between financial inclusion and the flow of private transfers.

Bibliography

Abowd, J. and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57, 411–445.

Abraham, K. G. and M. S. Kearney (2020): "Explaining the Decline in the US Employment-to-Population Ratio: A Review of the Evidence," *Journal of Economic Literature*, 58, 585–643.

Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2016): "Import Competition and the Great US Employment Sag of the 2000s," *Journal of Labor Economics*, 34, S141–S198.

Acemoglu, D. and P. Restrepo (2020): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128, 2188–2244.

Aguiar, M. and M. Bils (2015): "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review*, 105, 2725–56.

Aguiar, M., M. Bils, K. K. Charles, and E. Hurst (2021): "Leisure Luxuries and the Labor Supply of Young Men," *Journal of Political Economy*, 129, 337–382.

Aguiar, M. and E. Hurst (2007): "Measuring Trends in Leisure: The Allocation of Time Over Five Decades," *Quarterly Journal of Economics*, 122, 969–1006.

——— (2008): "The Increase in Leisure Inequality," NBER Working Papers 13837, National Bureau of Economic Research, Inc.

——— (2013): "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121, 437–492.

Ahn, T., C. G. Chee, and S. Kim (2021): "The Evolution of Income Risk and Consumption Insurance in South Korea over the Last Two Decades*," *Oxford Bulletin of Economics and Statistics*, 83, 328–351.

Altonji, J. G., F. Hayashi, and L. Kotlikoff (1996): "The Effects of Income and Wealth on Time and Money Transfers between Parents and Children," NBER Working Paper No. 5522, National Bureau of Economic Research.

Alvaredo, F. and A. B. Atkinson (2010): "Colonial Rule, Apartheid and Natural Resources: Top Incomes in South Africa 1903-2005," OxCarre Working Papers No. 046, Oxford Centre for the Analysis of Resource Rich Economies.

Andrews, D. W. K. (1993): "Exactly Median-Unbiased Estimation of First Order Autoregressive/Unit Root Models," *Econometrica*, 61, 139–165.

Ardington, C., A. Case, and V. Hosegood (2009): "Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa," *American Economic Journal: Applied Economics*, 1, 22–48.

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85, 693–734.

Attanasio, O. and S. J. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104, 1227–1262.

Attanasio, O., H. Low, and V. Sánchez-Marcos (2005): "Female Labor Supply As Insurance Against Idiosyncratic Risk," *Journal of the European Economic Association*, 3, 755–764.

Attanasio, O. and J.-V. Ríos-Rull (2000): "Consumption smoothing in island economies: Can public insurance reduce welfare?" *European Economic Review*, 44, 1225–1258.

Auclert, A. and M. Rognlie (2018): "Inequality and Aggregate Demand," Working Paper 24280, National Bureau of Economic Research.

Autor, D. H. and D. Dorn (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–97.

Autor, D. H., D. Dorn, and G. H. Hanson (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–68.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2008): "Trends in U.S. Wage Inequality: Revising the Revisionists," *The Review of Economics and Statistics*, 90, 300–323.

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, 1279–1333.

BACHMANN, R. AND D. BAUMGARTEN (2013): "How do the unemployed search for a job? – Evidence from the EU Labour Force Survey," *IZA Journal of European Labor Studies*, 22.

BAEZ, J. E. (2006): "Income Volatility, Risk-Coping Behavior and Consumption Smoothing Mechanisms in Developing Countries: A Survey." *Desarrollo y Sociedad*, 58, 37 – 83.

BELLEMARE, M. F., Y. N. LEE, AND L. NOVAK (2021): "Contract farming as partial insurance," *World Development*, 140.

BELLON, M., C. PIZZINELLI, AND R. PERRELLI (2020): "Household Consumption Volatility and Poverty Risk: Case Studies from South Africa and Tanzania," IMF Working Papers No. WP/20/51, International Monetary Fund.

BÉNABOU, R. (2002): "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70, 481–517.

BERTRAND, M., S. MULLAINATHAN, AND D. MILLER (2003): "Public Policy and Extended Families: Evidence from Pensions in South Africa," *World Bank Economic Review*, 17, 27–50.

BINDER, A. J. AND J. BOUND (2019): "The Declining Labor Market Prospects of Less-Educated Men," *Journal of Economic Perspectives*, 33, 163–90.

BLINDER, A. AND A. KRUEGER (2013): "Alternative Measures of Offshorability: A Survey Approach." *Journal of Labor Economics*, 31.

BLOOM, N., M. DRACA, AND J. VAN REENEN (2015): "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity," *The Review of Economic Studies*, 83, 87–117.

BLUNDELL, R. (2014): "Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies," *Economic Journal*, 124, 289–318.

BLUNDELL, R., M. GRABER, AND M. MOGSTAD (2015): "Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family," *Journal of Public Economics*, 127, 58–73.

BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921.

BLUNDELL, R., L. PISTAFERRI, AND I. SAPORTA-EKSTEN (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435.

BRYUKHANOV, M. AND D. HRYSHKO (2020): "Inequality in Russia over Time and over the Life Cycle," mimeo, University of Alberta.

BRZOZOWSKI, M., M. GERVAIS, P. KLEIN, AND M. SUZUKI (2010): "Consumption, Income, and Wealth Inequality in Canada," *Review of Economic Dynamics*, 13, 52–75.

BUN, M. J. AND V. SARAFIDIS (2015): "Dynamic Panel Data Models," in *The Oxford Handbook of Panel Data*, ed. by B. H. Baltagi, Oxford: Oxford University Press, chap. 3, 76–110.

BURGER, R. AND D. VON FINTEL (2009): "Determining the Causes of the Rising South African Unemployment Rate: An Age, Period and Generational Analysis," Stellenbosch Economic Working Papers No: 24, The Bereau of Economic Research.

CAMPBELL, J. AND A. DEATON (1989): "Why is Consumption So Smooth?" *Review of Economic Studies*, 56, 357–373.

CARROLL, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55.

——— (2009): "Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income," *Journal of Monetary Economics*, 56, 780–790.

CASADO, J. M. (2011): "From Income to Consumption: Measuring Household Partial Insurance," *Empirical Economics*, 40, 471–495.

CASE, A. AND A. DEATON (1998): "Large Cash Transfers to the Elderly in South Africa," *The Economic Journal*, 108, 1330–1361.

CASE, A., V. HOSEGOOD, AND F. LUND (2005): "The reach and impact of Child Support Grants: evidence from KwaZulu-Natal," *Development Southern Africa*, 22, 467–482.

CHEN, S., V. CHERNOZHUKOV, AND I. FERNÁNDEZ-VAL (2019): "Mastering Panel Metrics: Causal Impact of Democracy on Growth," *AEA Papers and Proceedings*, 109, 77–82.

CHETTY, R. (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116, 173–234.

CHEUNG, A., J. MARCHAND, AND P. MARK (2020): "Loss of Life and Labour Productivity: The Canadian Opioid Crisis," Department of economics, Working Paper No. 2020-13, University of Alberta.

CHRISTENSEN, B. J., R. LENTZ, D. T. MORTENSEN, G. R. NEUMANN, AND A. WERWATZ (2005): "On-the-Job Search and the Wage Distribution," *Journal of Labor Economics*, 23, 31–58.

COCHRANE, J. H. (1991): "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99, 957–976.

DAFULEYA, G. AND F. TREGENNA (2021): "How effectively do households insure food consumption and assets against funeral expenses? The case of urban Zimbabwe," *Review of Economics of the Household*, 19, 987–1021.

DALY, M., D. HRYSHKO, AND I. MANOVSKII (2021): "Improving the Measurement of Earnings Dynamics," *International Economic Review (forthcoming)*.

DAUTH, W., S. FINDEISEN, J. SUEDEKUM, AND N. WOESSNER (2021): "The Adjustment of Labor Markets to Robots," *Journal of the European Economic Association*, jvab012.

DE MAGALHÃES, L., D. KOH, AND R. SANTAEULÀLIA-LLOPIS (2019): "The costs of consumption smoothing: less schooling and less nutrition," *Journal of Demographic Economics*, 85, 181–208.

DE NARDI, M., G. FELLA, AND G. PAZ-PARDO (2019): "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 00, 1–37, jvz010.

DERCON, S., J. HODDINOTT, AND T. WOLDEHANNA (2005): "Shocks and

Consumption in 15 Ethiopian Villages, 1999–2004," *Journal of African Economies*, 14, 559–585.

Diamond, P. (1982): "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90, 881–94.

Dixon, J., B. Hong, and L. Wu (2021): "The Robot Revolution: Managerial and Employment Consequences for Firms," *Management Science*, 67, 5586–5605.

Duflo, E. (2003): "Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa," *The World Bank Economic Review*, 17, 1–25.

Dutronc-Postel, P. and A. Tondini (forthcoming): "Large Means-Tested Pensions with Informal Labor Markets: Evidence from South Africa," *Economic Development and Cultural Change*.

Eberstadt, N. (2016): *Men without Work*, West Conshohocken, PA: Templeton Press.

——— (2022): *Men Without Work: Post-Pandemic Edition*, Templeton Press, second edition ed.

Edmonds, E. V. (2006): "Child labor and schooling responses to anticipated income in South Africa," *Journal of Development Economics*, 81, 386–414.

Faberman, R. J. and M. Kudlyak (2019): "The Intensity of Job Search and Search Duration," *American Economic Journal: Macroeconomics*, 11, 327–57.

Farber, H. S., C. M. Herbst, D. Silverman, and T. von Wachter (2019): "Whom Do Employers Want? The Role of Recent Employment and Unemployment Status and Age," *Journal of Labor Economics*, 37, 323–349.

Farber, H. S., D. Silverman, and T. von Wachter (2016): "Determinants of Callbacks to Job Applications: An Audit Study," *American Economic Review*, 106, 314–18.

Farhi, E. and I. Werning (2012): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 80, 596–635.

Farré, L., F. Fasani, and H. Mueller (2018): "Feeling useless: the effect of unemployment on mental health in the Great Recession," *IZA Journal of*

Labor Economics, 7.

FESTUS, L., A. KASONGO, M. MOSES, AND D. YU (2016): “The South African labour market, 1995–2015,” *Development Southern Africa*, 33, 579–599.

FRIEDMAN, M. (1957): *A Theory of the Consumption Function*, National Bureau of Economic Research, Inc.

GALARNEAU, D., R. MORISSETTE, AND J. USALCAS (2013): “What has changed for young people in Canada?” Catalogue no. 75-006-X 75-006-X, Statistics Canada.

GATHERGOOD, J. AND D. WYLIE (2018): “Why are some households so poorly insured?” *Journal of Economic Behavior & Organization*, 156, 1–12.

GEORGIEFF, A. AND A. MILANEZ (2021): “What happened to jobs at high risk of automation?” Oecd publishing, Employment and Migration Working Papers, No. 255.

GERRY, C. J. AND C. A. LI (2010): “Consumption smoothing and vulnerability in Russia.” *Applied Economics*, 42, 1995 – 2007.

GOENSCH, J., A. GULYAS, AND I. KOSPENTARIS (2021): “Unemployment Insurance Reforms in a Search Model With Endogenous Labor Force Participation,” CRC TR 224 Discussion Paper Series 264, University of Bonn and University of Mannheim, Germany.

GOMME, P. AND D. LKHAGVASUREN (2015): “Worker search effort as an amplification mechanism,” *Journal of Monetary Economics*, 75, 106–122.

GOURIÉROUX, C., P. C. PHILLIPS, AND J. YU (2010): “Indirect inference for Dynamic Panel Models,” *Journal of Econometrics*, 157, 68–77.

GRADIN, C. (2019): “Occupational segregation by race in South Africa after apartheid,” *Review of Development Economics*, 23, 553–576.

GRIGOLI, F., Z. KOCZAN, AND P. TOPALOVA (2020): “Automation and labor force participation in advanced economies: Macro and micro evidence,” *European Economic Review*, 126, 1–23.

GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2014): “Income taxation of U.S. households: Facts and parametric estimates,” *Review of Economic Dynamics*, 17, 559 – 581.

GURIEV, S. (2018): “Economic Drivers of Populism,” *AEA Papers and Pro-*

ceedings, 108, 200–203.

Gustavsson, M. and P. Österholm (2014): "Does the Labor-Income Process Contain a Unit Root? Evidence from Individual-Specific Time Series," *Journal of Economic Dynamics and Control*, 47, 152–167.

Guvenen, F. (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12, 58–79.

Guvenen, F., F. Karahan, S. Ozcan, and J. Song (2021): "What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?" *Econometrica*, 89, 2303–2339.

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006," *Review of Economic Dynamics*, 13, 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2014): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," *American Economic Review*, 104, 2075–2126.

Hryshko, D. (2012): "Labor Income Profiles Are Not Heterogeneous: Evidence From Income Growth Rates," *Quantitative Economics*, 3, 177–209.

——— (2014): "Correlated Income Shocks and Excess Smoothness of Consumption," *Journal of Economic Dynamics and Control*, 48, 41–62.

Hryshko, D. and I. Manovskii (2022): "How much consumption insurance in the U.S.?" *Journal of Monetary Economics*.

Hundenborn, J., M. Leibbrandt, and I. Woolard (2018): "Drivers of Inequality in South Africa," WIDER Working Paper 2018/162, UNU-WIDER.

Jacoby, H. G. and E. Skoufias (1997): "Risk, Financial Markets, and Human Capital in a Developing Country," *The Review of Economic Studies*, 64, 311–335.

Jamison, D. T. (1986): "Child malnutrition and school performance in China," *Journal of Development Economics*, 20, 299 – 309.

Janzen, S. A. and M. R. Carter (2018): "After the Drought: The Impact of Microinsurance on Consumption Smoothing and Asset Protection," *American Journal of Agricultural Economics*, 101, 651–671.

Jarosch, G. and L. Pilossoph (2018): "Statistical Discrimination and Duration Dependence in the Job Finding Rate," *The Review of Economic Studies*, 86, 1631–1665.

Juhn, C. (1992): "The Role of Declining Market Opportunities," *The Quarterly Journal of Economics*, 107, 79–121.

Kaplan, G. and G. L. Violante (2010): "How Much Consumption Insurance beyond Self-Insurance?" *American Economic Journal: Macroeconomics*, 2, 53–87.

Karahan, F. and S. Ozkan (2013): "On the Persistence of Income Shocks over the Life Cycle: Evidence, Theory, and Implications," *Review of Economic Dynamics*, 16, 452–476.

Keldenich, C. and A. Knabe (2022): "Women's Labor Market Responses to Their Partners' Unemployment and Low-Pay Employment," *Journal of Labor Research*, 43, 134–162.

Kinnan, C. (2022): "Distinguishing Barriers to Insurance in Thai Villages," *Journal of Human Resources*, 57, 44–78.

Kospentaris, I. (2021): "Unobserved Heterogeneity and Skill Loss in a Structural Model of Duration Dependence," *Review of Economic Dynamics*, 39, 280–303.

Kroft, K., F. Lange, M. J. Notowidigdo, and L. F. Katz (2016): "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation," *Journal of Labor Economics*, 34, 7–54.

Kroft, K., F. Lange, M. J. Notowidigdo, and M. Tudball (2019): "Long Time Out: Unemployment and Joblessness in Canada and the United States," *Journal of Labor Economics*, 37, S355–S397.

Krueger, A. B. (2017): "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate," *Brookings Papers on Economic Activity*, 47, 1–87.

Krueger, D. and F. Perri (2006): "Does Income Inequality Lead To Consumption Inequality? Evidence And Theory," *Review of Economic Studies*, 73, 163–193.

KUBOTA, K. (2021): "Partial insurance in Japan," *The Japanese Economic Review*, 72, 299–328.

KUDLYAK, M. AND A. HORNSTEIN (2017): "Estimating Matching Efficiency with Variable Search Effort," 2017 Meeting Papers 881, Society for Economic Dynamics.

KUHN, P. AND A. L. ROBB (1998): "Shifting Skill Demand and the Canada-US Unemployment Gap: Evidence from Prime-Age Men," *Canadian Public Policy*, 24.

LANDAIS, C., P. MICHAILLAT, AND E. SAEZ (2018): "A Macroeconomic Approach to Optimal Unemployment Insurance: Applications," *American Economic Journal: Economic Policy*, 10, 182–216.

LARSON, R., S. SHANNON, A. SOJOURNER, AND C. UGGEN (2021): "Felon history and change in U.S. employment rates," *Social Science Research*, 102649.

LEIBBRANDT, M., V. RANCHHOD, AND G. PIPPA (2018): "Taking stock of South African income inequality," WIDER Working Paper 2018/184, UNU-WIDER.

LEIBBRANDT, M., E. WEGNER, AND A. FINN (2011): "The Policies for Reducing Income Inequality and Poverty in South Africa," A Southern Africa Labour and Development Research Unit Working Paper 64, SALDRU.

LESTER, B., D. RIVERS, AND G. TOPA (2018): "Job Referrals and Skills," 2018 Meeting Papers 306, Society for Economic Dynamics.

LEVINSOHN, J. (2008): "Two policies to alleviate unemployment in South Africa," CID Working Paper No. 166, Harvard: Centre for International Development.

LISE, J. (2013): "On-the-Job Search and Precautionary Savings," *Review of Economic Studies*, 80, 1086–1113.

LYSHOL, A. F., P. T. NENOV, AND T. WEVELSTAD (2021): "Duration Dependence and Labor Market Experience," *LABOUR*, 35, 105–134.

MABOSHE, M. AND I. WOOLARD (2018): "Revisiting the Impact of Direct Taxes and Transfers on Poverty and Inequality in South Africa," WIDER Working Paper 2018/79, UNU-WIDER.

MALGOUYRES, C. (2017): "The Impact of Chinese Import Competition on the Local Structure of Employment and Wages: Evidence from France," *Journal of Regional Science*, 57, 411–441.

MEGHIR, C. AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72, 1–32.

MODIGLIANI, F. AND H. BRUMBERG, R (1954): *Post Keynesian Economics*, Rutgers University Press, chap. Utility analysis and the consumption function: an interpretation of cross-section data, 388–436.

MOFFITT, R. A. (2012): "The Reveral of the Employment-Population Ratio in the 2000s: Facts and Explanations," *Brookings Papers on Economic Activity*, 43, 201–264.

MORTENSEN, D. (1982): "Property Rights and Efficiency in Mating, Racing, and Related Games," *American Economic Review*, 72, 968–79.

MOSCARINI, G. (2001): "Excess Worker Reallocation," *Review of Economic Studies*, 68, 593–612.

MUKOYAMA, T., C. PATTERSON, AND A. ŞAHIN (2018): "Job Search Behavior over the Business Cycle," *American Economic Journal: Macroeconomics*, 10, 190–215.

MUSSIDA, C. AND L. ZANIN (2020): "Determinants of the Choice of Job Search Channels by the Unemployed Using a Multivariate Probit Model," *Social Indicators Research*.

NAKAJIMA, M. (2012): "A quantitative analysis of unemployment benefit extensions," *Journal of Monetary Economics*, 59, 686–702.

NICKELL, S. AND B. BELL (1995): "The Collapse in Demand for the Unskilled and Unemployment across the OECD," *Oxford Review of Economic Policy*, 11, 40–62.

NOTTEN, G. AND D. DE CROMBRUGGHE (2012): "Consumption Smoothing in Russia." *Economics of Transition*, 20, 481 – 519.

PARK, S. AND D. POWELL (2021): "Is the rise in illicit opioids affecting labor supply and disability claiming rates?" *Journal of Health Economics*, 76, 102430.

PEREZ-ARCE, F. AND M. J. PRADOS (2021): "The Decline in the U.S. Labor Participation Rate: A Literature Review," *Journal of Economic Surveys*, 35, 615–652.

Picot, G. and A. Heisz (2000): "The Performance of the 1990s Canadian Labour Market," *Canadian Public Policy*, 26, 7–25.

Pissarides, C. (1985): "Short-run Equilibrium Dynamics of Unemployment Vacancies, and Real Wages," *American Economic Review*, 75, 676–90.

Pizzinelli, C. and B. Speigner (2017): "Matching efficiency and labour market heterogeneity in the United Kingdom," Bank of England working papers 667, Bank of England.

Riley, E. (2018): "Mobile money and risk sharing against village shocks," *Journal of Development Economics*, 135, 43–58.

Santaeulàlia-Llopis, R. and Y. Zheng (2018): "The Price of Growth: Consumption Insurance in China 1989–2009," *American Economic Journal: Macroeconomics*, 10, 1–35.

Schulhofer-Wohl, S. (2011): "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, 119, 925 – 958.

Sedlacek, P. (2016): "The aggregate matching function and job search from employment and out of the labor force," *Review of Economic Dynamics*, 21, 16–28.

Setty, O. and Y. Yedid-Levi (2020): "On the Provision of Unemployment Insurance when Workers are Ex-Ante Heterogeneous," *Journal of the European Economic Association*, 19, 664–706.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

——— (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15, 127–148.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 112, 695–717.

Themane, M., K. Monyeki, M. Nthangeni, H. Kemper, and J. Twisk (2003): "The relationship between health (malnutrition) and educational achievements (Maths and English) in the rural children of South Africa," *International Journal of Educational Development*, 23, 637 – 643.

Townsend, R. M. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539–591.

Van Belle, E., V. Di Stasio, R. Caers, M. De Couck, and S. Baert (2018): "Why Are Employers Put Off by Long Spells of Unemployment?" *European Sociological Review*, 34, 694–710.

Wolcott, E. L. (2021): "Employment inequality: Why do the low-skilled work less now?" *Journal of Monetary Economics*, 118, 161–177.

Ye, C. (2021): "How Does Education Affect the Time Spending on Work and Leisure?" in *Proceedings of the 6th International Conference on Financial Innovation and Economic Development (ICFIED 2021)*, Atlantis Press, 274–280.

Álvarez, R. and S. Claro (2009): "David Versus Goliath: The Impact of Chinese Competition on Developing Countries," *World Development*, 37, 560–571.

Çenesiz, M. A. and L. Guimarães (2022): "The cyclicality of job search effort in matching models," *Oxford Economic Papers*, gpac001.

Appendix A: Employment Inequality Among Prime-Age Men in Canada

A.1 Figures

This section contains the rest of figures to support results in chapter 2.

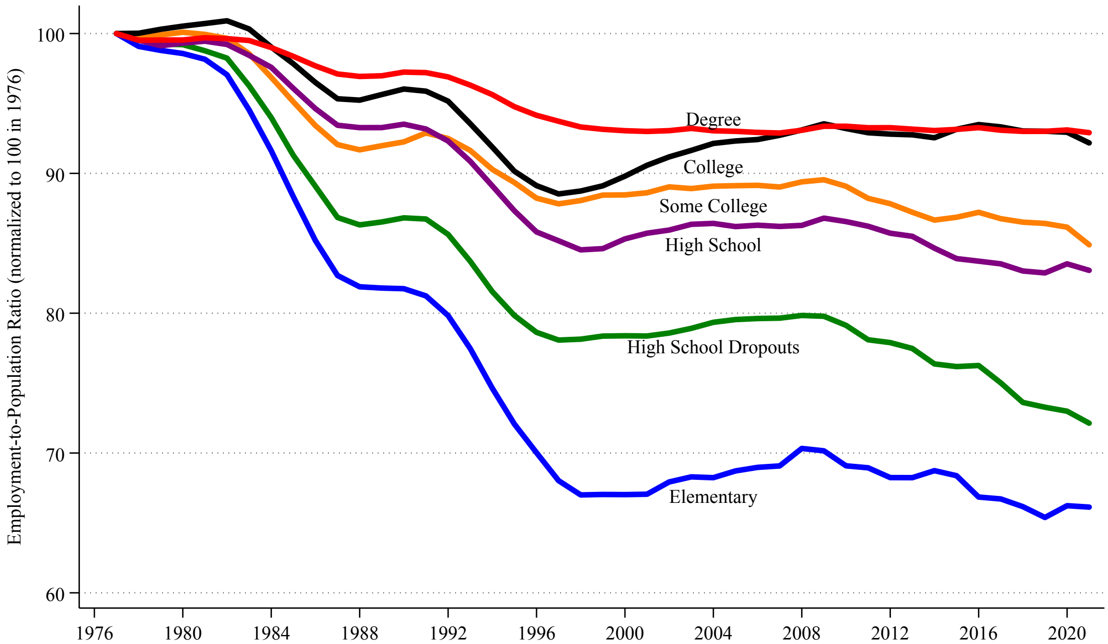

Figure A-1: Employment-to-Population Ratios of Prime-Age Men, Age 25-54

Notes: The data is annual averages of the monthly Canadian Labour Force Surveys (LFS). The plots are demographically adjusted for age. The degree group includes individuals with professional qualifications.

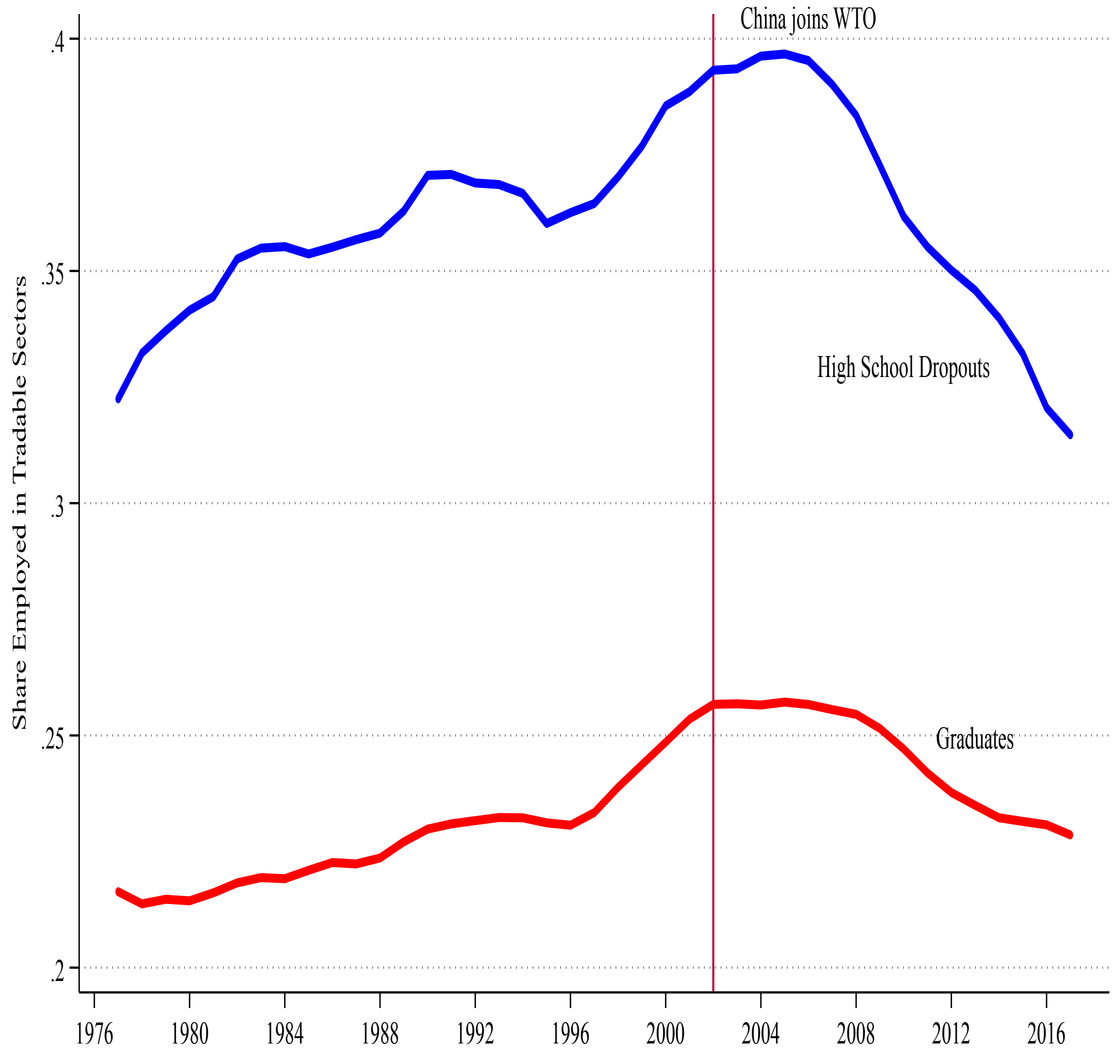

Figure A-2: Share of Men Employed in Tradable Industries

Notes: The figure shows the share of men employed in tradable sectors (directly vulnerable to trade competition). The tradable industries include Fish, Forestry, Oil & Gas, Manufacturing and Food Processing. *Data source:* Labour Force Survey, 1976-2017.

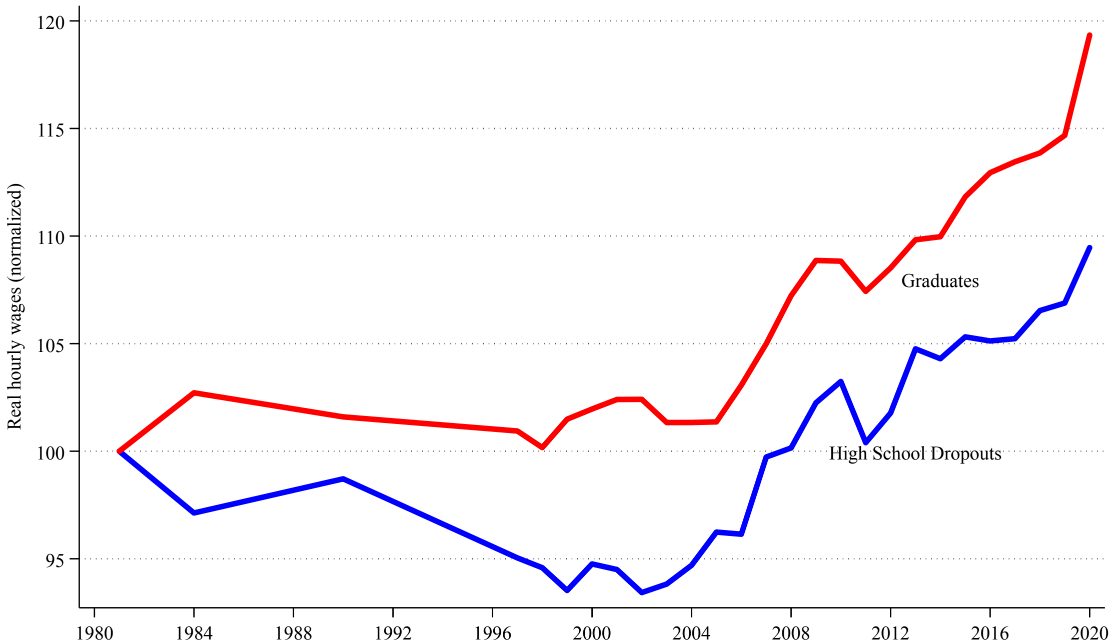

Figure A-3: Real Wages by Education, 1980-2020

Notes: This figure depicts the trends in real wages (base year 2015) normalized at 100 in 1980. The wages data is from Survey of Work History 1981, Survey of Union Membership, 1984, Labour Market Activity Survey, 1986-1990, Labour Force Survey, 1996-2020.

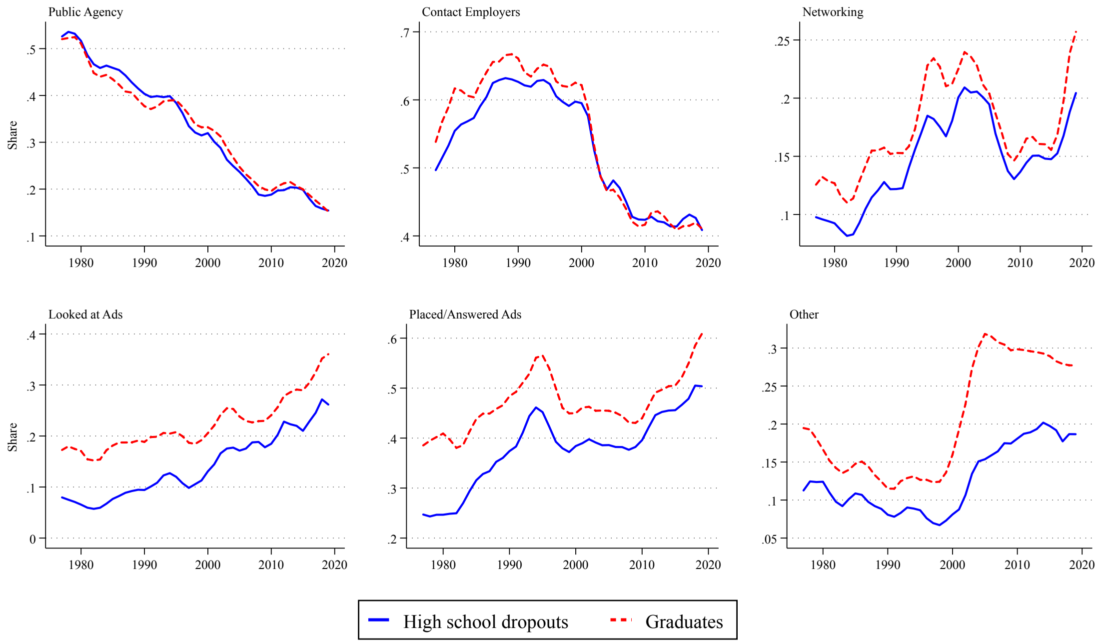

Figure A-4: Share of Prime-age Men Using Different Search Methods Over Time

Notes: The plots show shares of unemployment men using each search method to search for work. *Data source:* Labour Force Survey, 1976-2020.

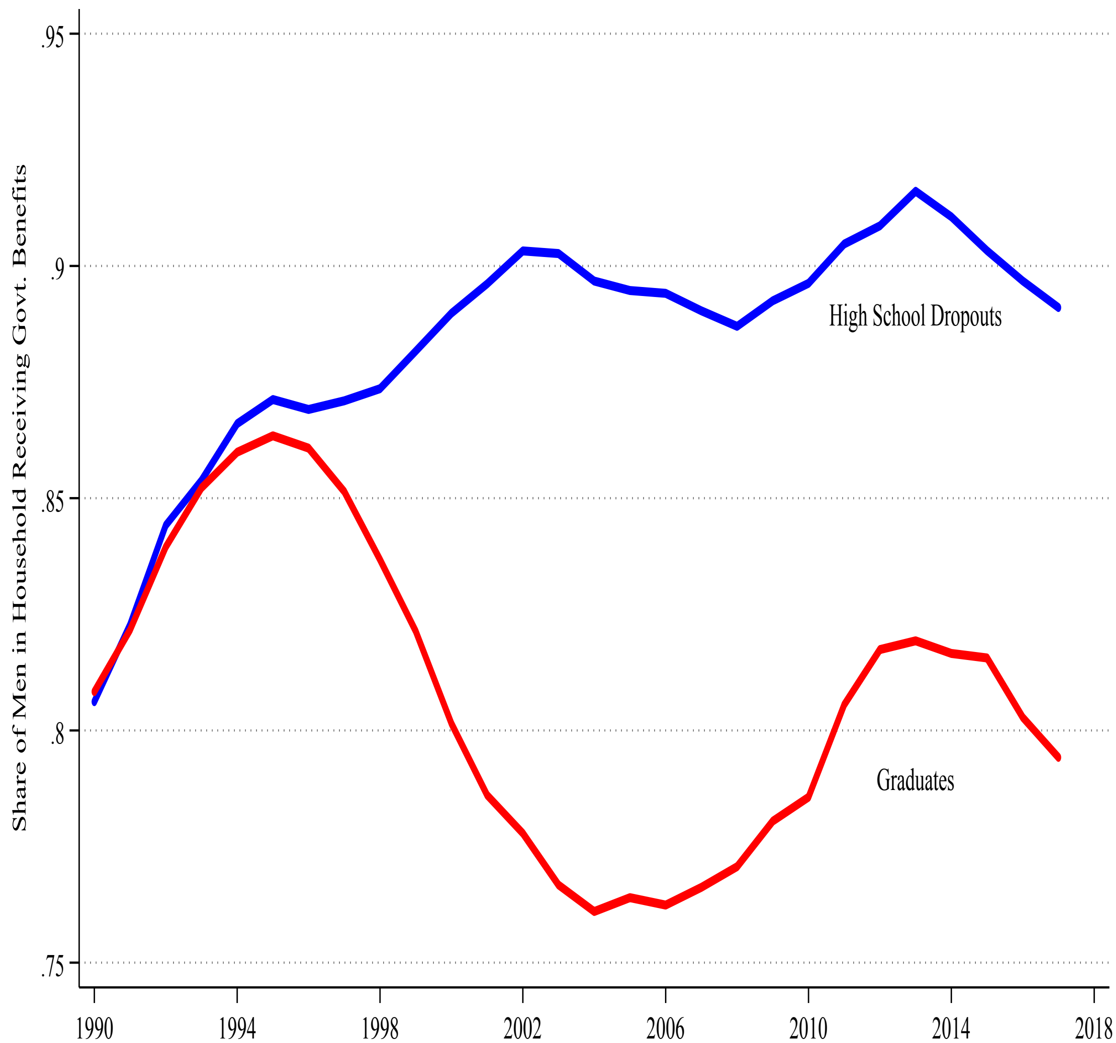

Figure A-5: Share of Men Living in Households Receiving Public Transfers

Notes: The figure shows the share of men living in households receiving some form of government transfers: social assistance, unemployment insurance, child tax credit, provincial transfers. *Data source:* Survey of Consumer Finances (1990-1998), Survey of Labour and Income Dynamics (1999-2011), and Consumer Income Surveys (2012-2017).

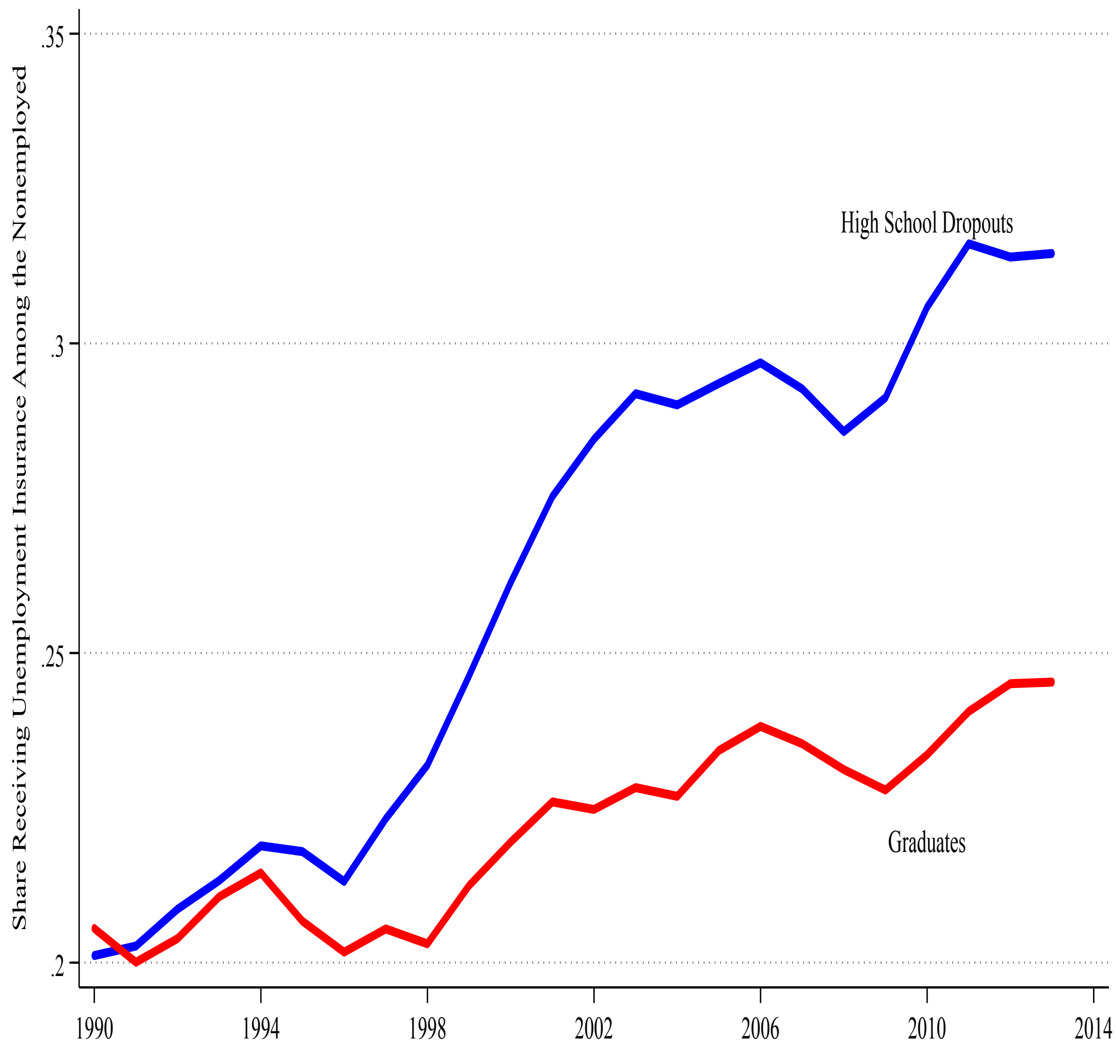

Figure A-6: Share of Nonemployed Men Receiving Employment Insurance

Notes: The figure shows the share of nonemployed men receiving unemployment insurance. *Data source:* Survey of Consumer Finances (1990-1998), Survey of Labour and Income Dynamics (1999-2011), and Consumer Income Surveys (2012-2017).

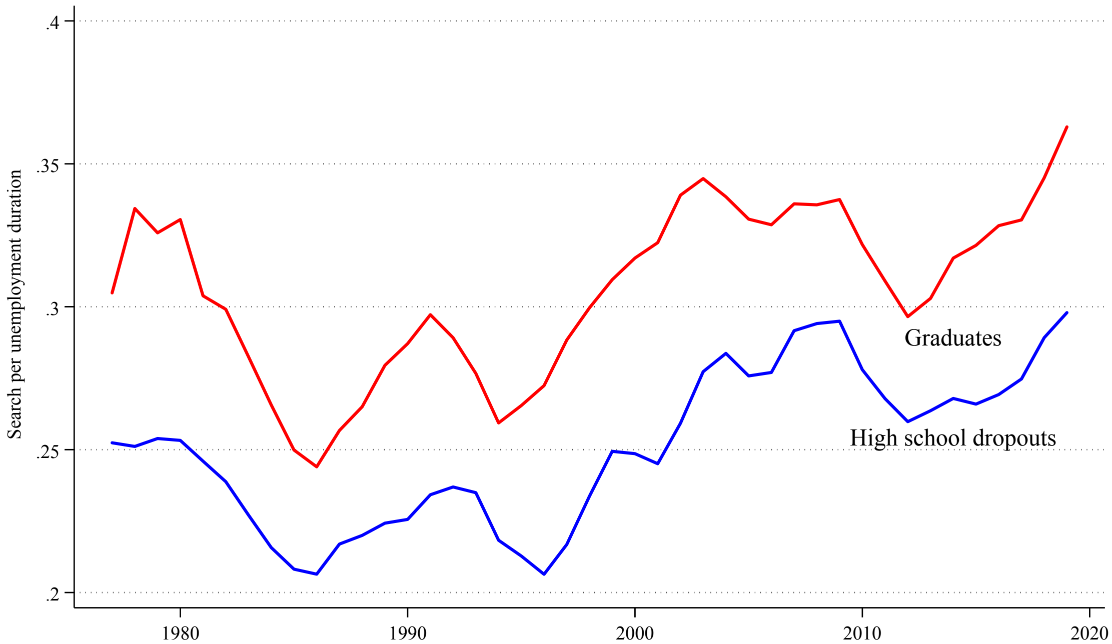

Figure A-7: Job Search Effort Normalized by Nonemployment Duration, 1976-2020

Notes: This figure presents annual averages of search effort computed as the number of search methods per nonemployment duration. *Data source:* Labour Force Survey.

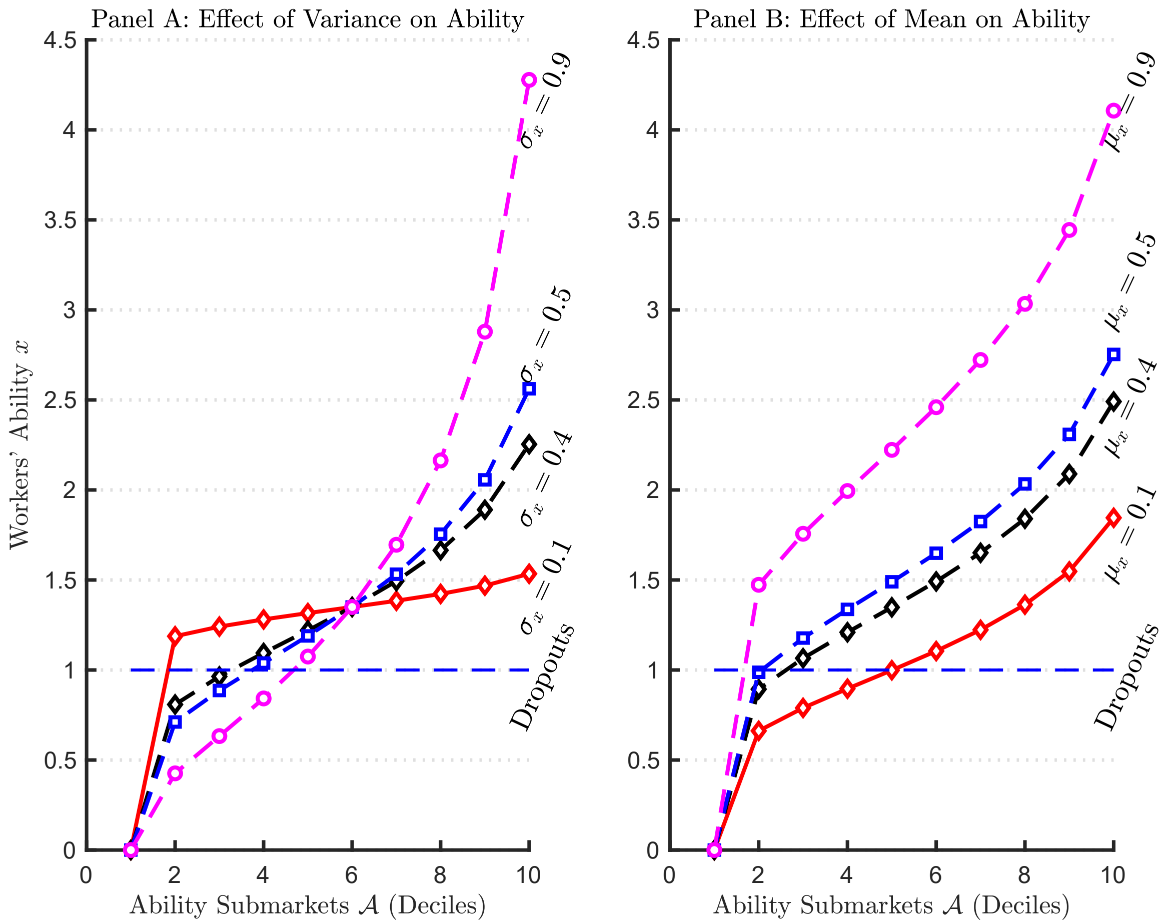

Figure A-8: Distribution of Ability Across Different Mean and Standard deviation by Deciles

Notes: The figure depicts the ability values in each submarket (deciles) for high school dropouts and graduates. Panel A plots the different ability values by deciles with μ_x fixed at 0.3 and varying σ_x from low to high. Panel B plots the different ability values by deciles, with σ_x fixed at 0.4 and varying μ_x from low to high.

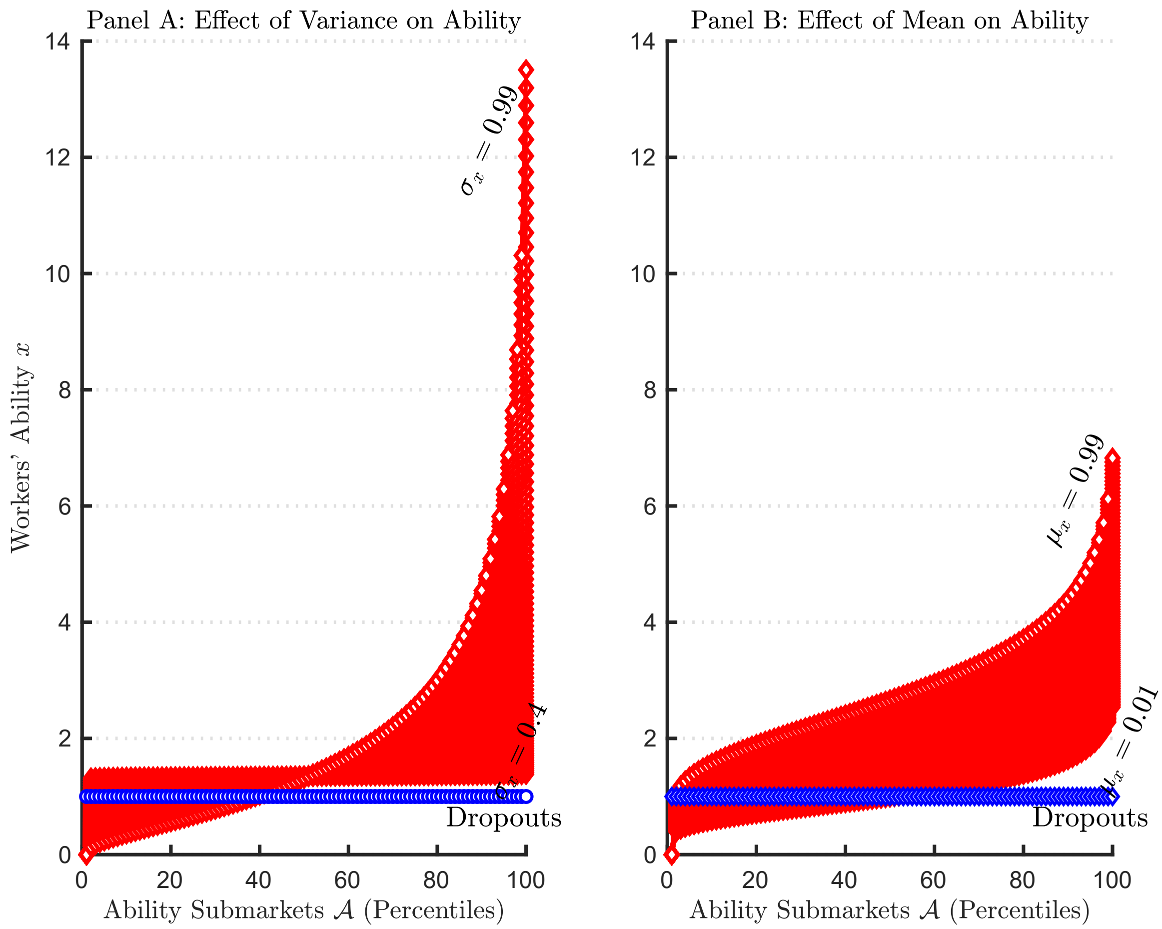

Figure A-9: Distribution of Ability Across Different Mean and Standard deviation by Percentiles

Notes: The figure depicts the ability values in each submarket (percentiles) for high school dropouts and graduates. Panel A plots the different ability values by percentiles with, μ_x fixed at 0.3 and varying σ_x from low to high. Panel B plots the different ability values by percentile with, σ_x fixed at 0.4 and varying μ_x from low to high.

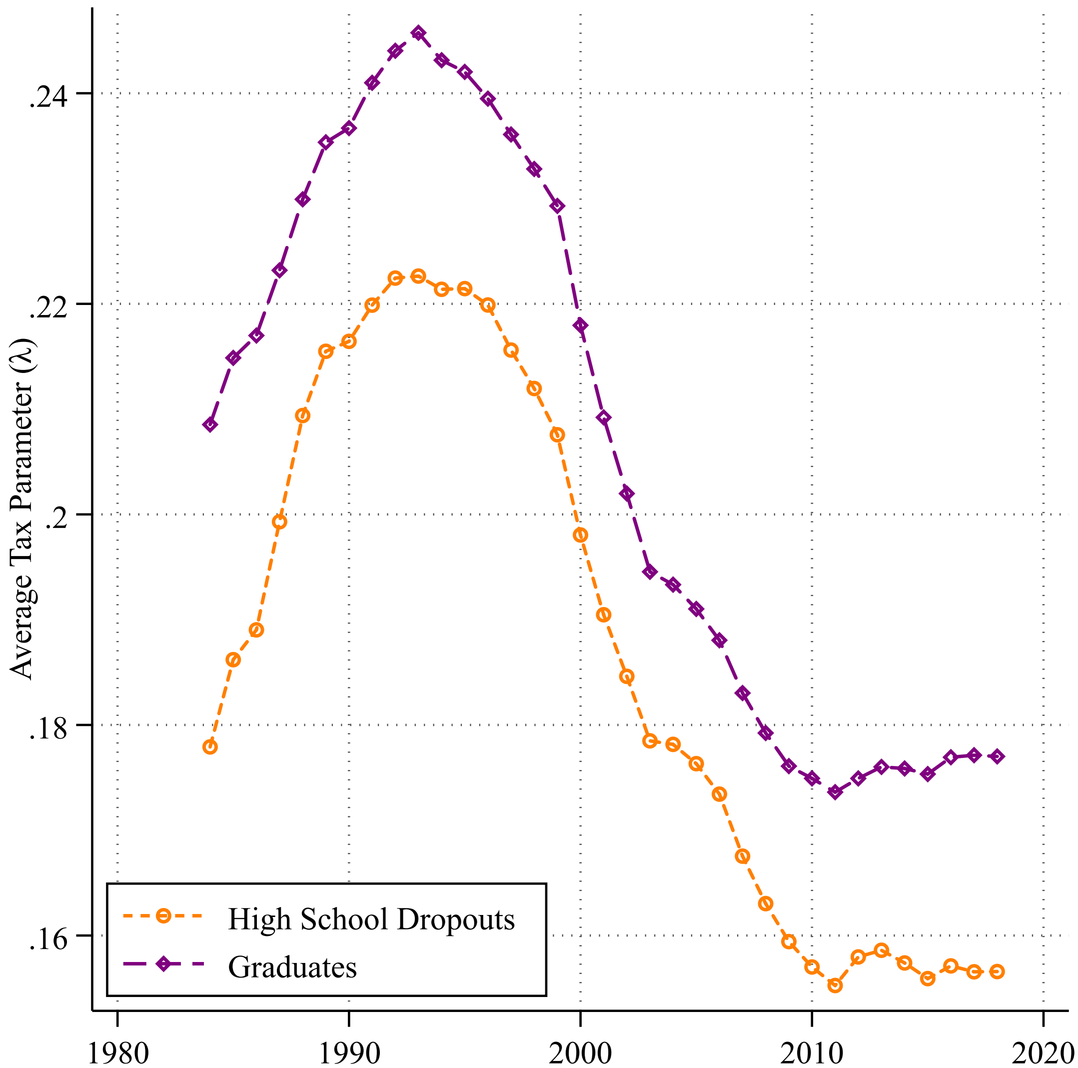

Figure A-10: Average Tax Level for Prime-Age Men Households in Canada, 1980-2020

Notes: The figure shows the changes in the average tax for prime-age men in Canada. The paramater λ is estimated using the Bénabou (2002) tax functions $T(\tilde{y}) = y - (1-\lambda)y^{(1-\bar{\tau})}$ where $T(\tilde{y})$ is tax and y is income before tax. The parameter λ indicates the average tax level in the income and $\bar{\tau}$ is a parameter that measures the progressivity of the tax system. The parameters are estimated using OLS and a three-year rolling data sample. Data source: Survey of Consumer Finances, 1980-1996, Survey of Labour and Income Dynamics, 1997-2011 and Consumer Income Surveys, 2012-2018.

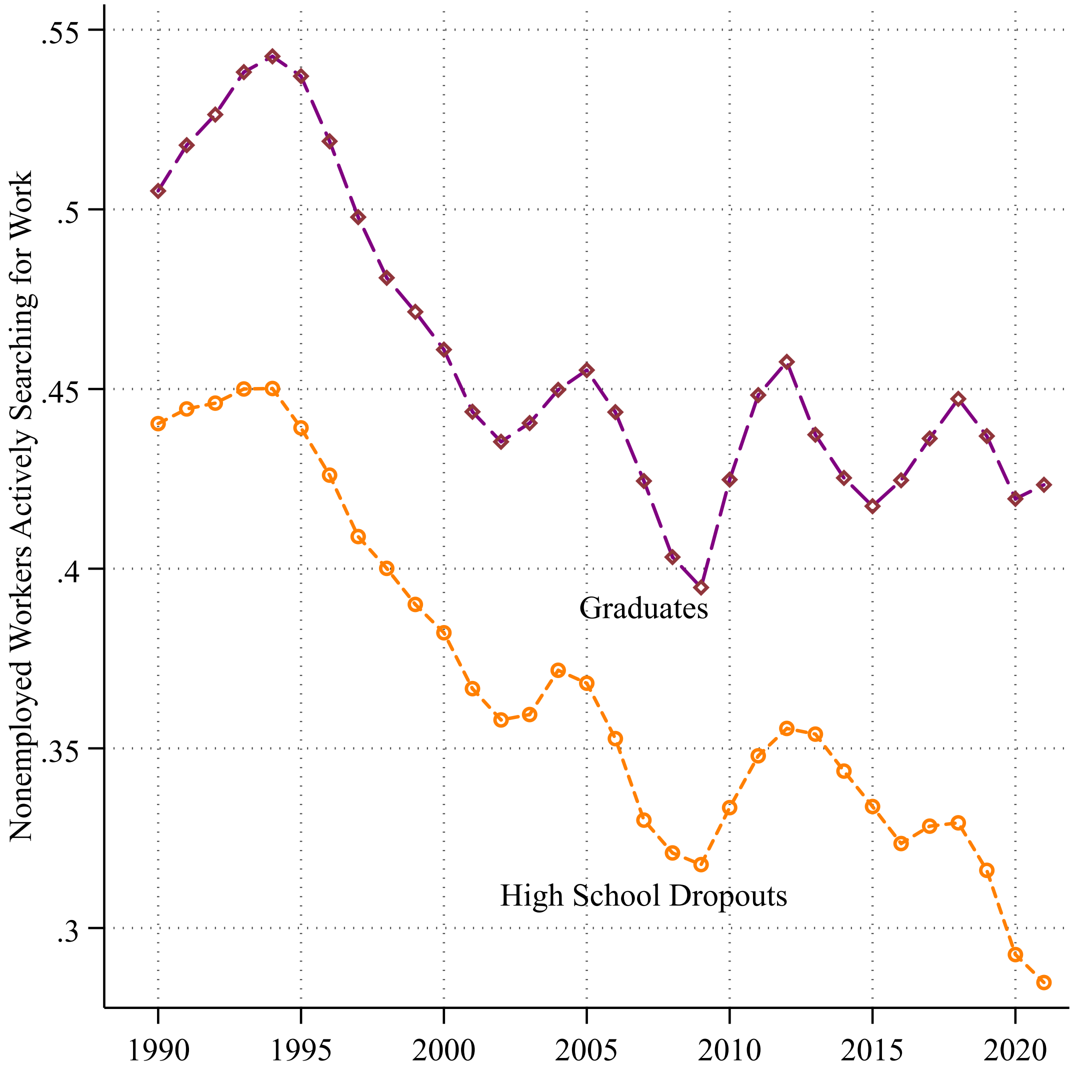

Figure A-11: Share of Nonemployed Workers Actively Searching for Work, 1990-2021

Notes: The figure depicts the trends in the share of nonemployed workers engaging in active job search. The shares are measured as the number of unemployed workers over the number of nonemployed workers. Data source: Labour Force Surveys, 1990-2021.

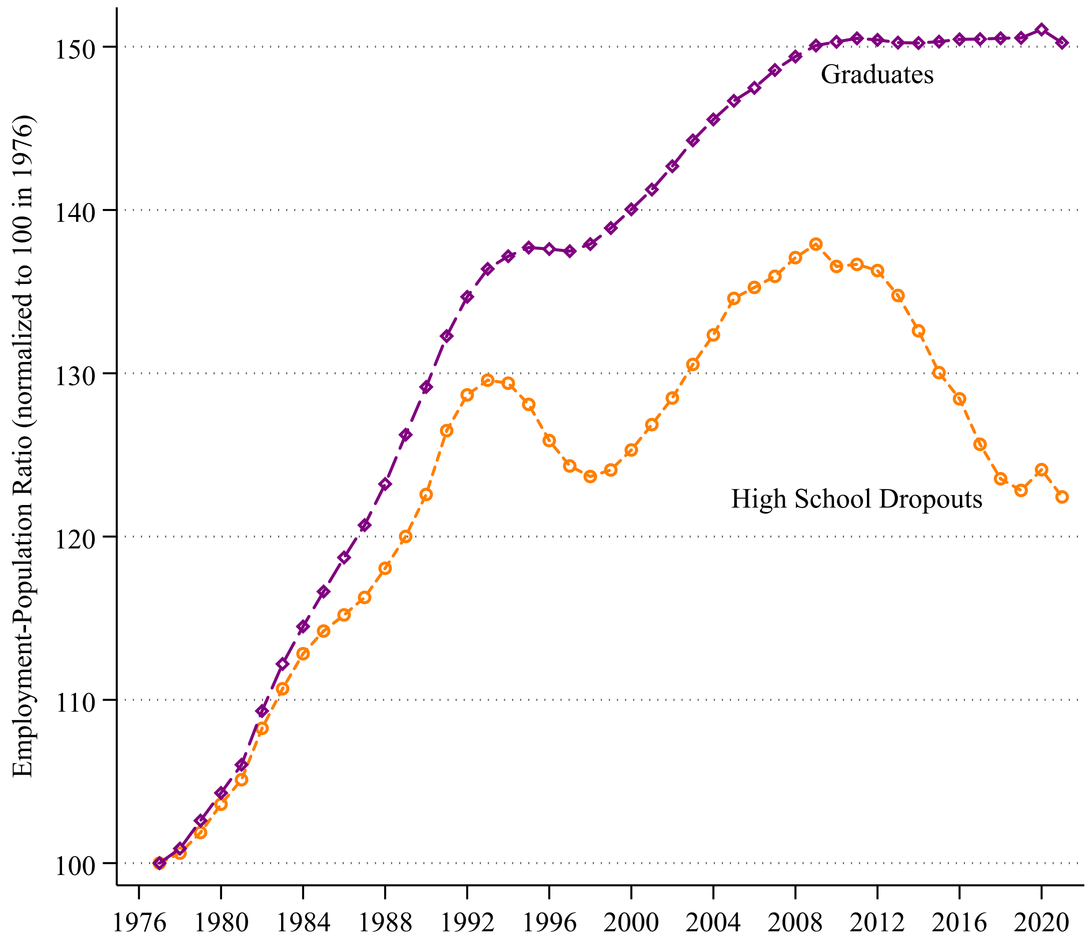

Figure A-12: Employment-Population Rate of Prime-Age Women in Canada, Age 25-54

Notes: The data is from annual averages of monthly Canadian Labour Force Surveys (LFS) labour market stocks. The trends are adjusted for age and province of residence. High school dropouts include women with no schooling, elementary education and those who dropped before completing high school. Graduates include women with at least a high school diploma. The data is weighted by survey weights. The base employment rates in 1976 are 35% and 52% for high school dropouts and graduates, respectively.

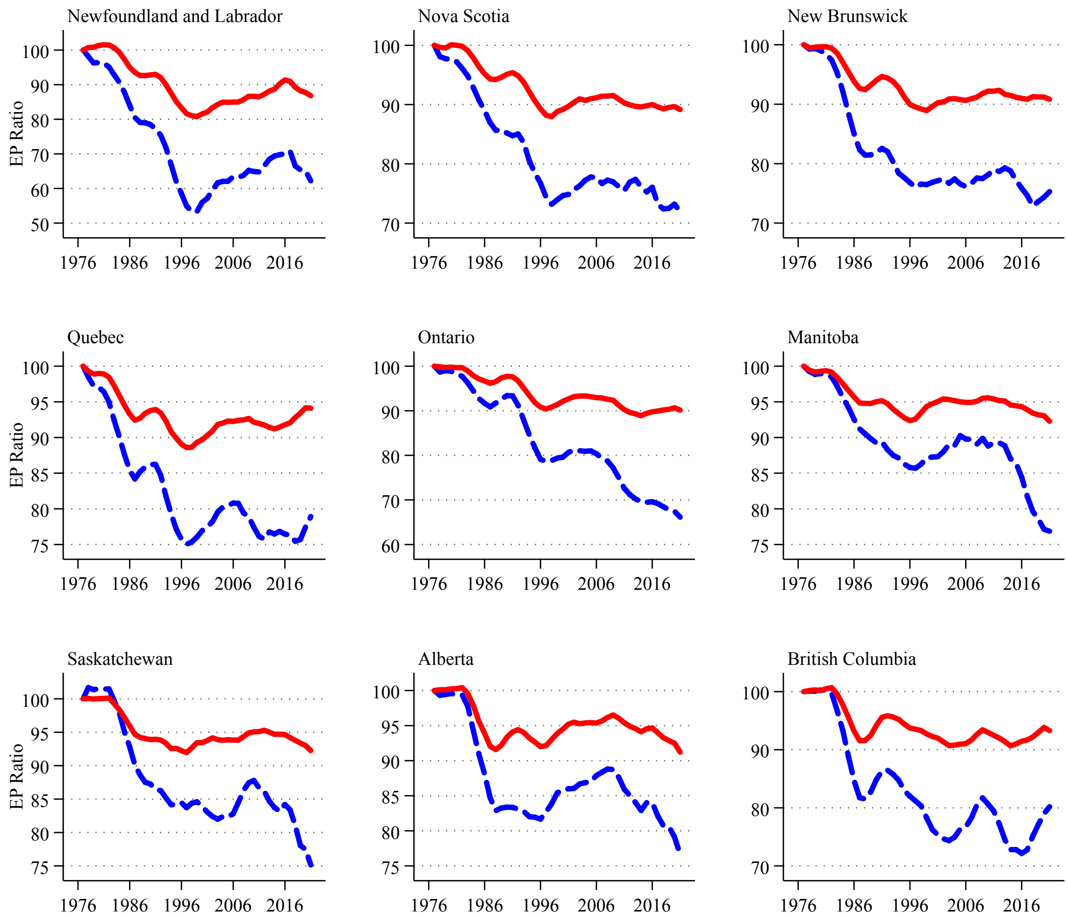

Figure A-13: Employment-Population Rate of Prime-Age Men Across Provinces, Age 25-54

Notes: The data is from annual averages of monthly Canadian Labour Force Surveys (LFS) labour market stocks. The trends are adjusted for age and province of residence. High school dropouts include men with no schooling, elementary education and those who dropped before completing high school. Graduates include men with at least a high school diploma. The data is weighted by survey weights.

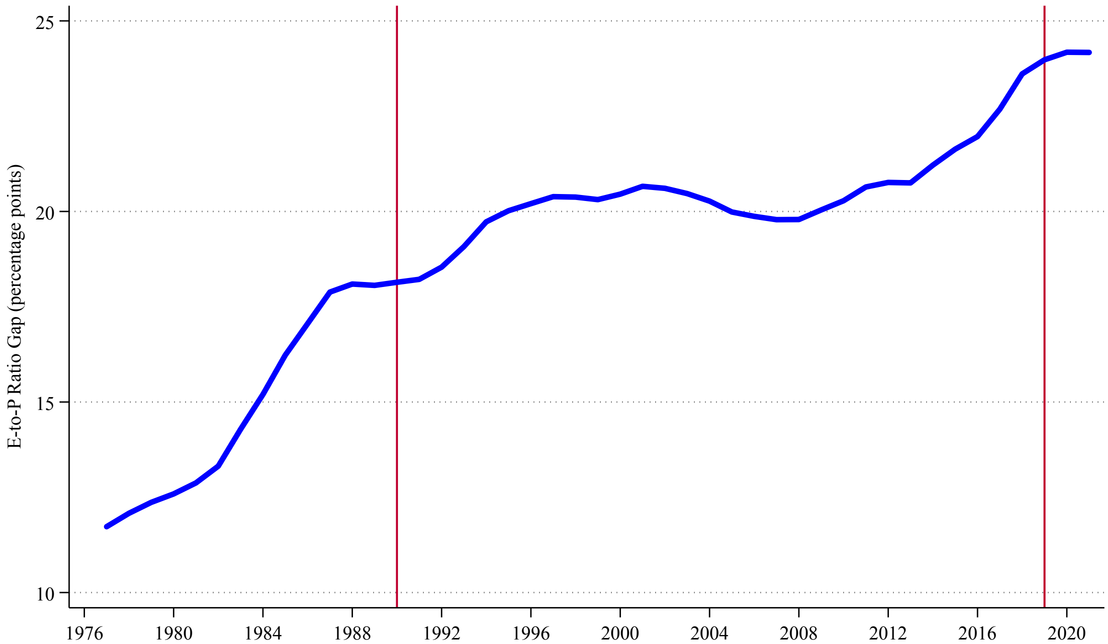

Figure A-14: Employment-Population Ratio Gap for Men, Ages 25-54

Notes: The plot shows the rise in the gap in employment rates between low-educated men and high-educated men (employment inequality). The data is annual averages of the monthly Canadian Labour Force Surveys (LFS). The plots are demographically adjusted for age.

A.2 Tables

Table A-1: Model Non-targeted Moments

Moment	Explanation	Period	Model	Data	Model Gap	Data Gap
			(1)	(2)	(3)	(4)
$e_{L,90}$	L employment rate	1990	71%	71%		
$e_{H,90}$	H employment rate	1990	88%	88%	17.1 pp	17.1 pp
$e_{L,19}$	L employment rate	2019	60%	60%		
$e_{H,19}$	H employment rate	2019	84%	84%	23.9 pp	24.0 pp
	Difference			6.8 pp	6.9 pp	

Notes: The table tabulates the performance of the model against moments from the data. The employment data is from the 1990 and 2019 LFS.

Appendix B: A Guide to Estimating the Canonical Income Process in Quasidifferences

with Dmytro Hryshko

B.1 Biases in the variances fixed effects and shocks

Tables B-1–B-3 document the results from regression analyses for the biases in the variance of fixed effects, permanent, and transitory shocks, respectively. Since we consider more than one true value for those variables, we express biases in each variable in percentages to their true values.

The variance of fixed effects is poorly identified when the true persistence is close to unity[1]—there is a large upward bias of at least 150 percent regardless of the weighting matrix used and sample size in terms of the number of individuals; see columns (7)–(12) of Table B-1. Biases are smaller when the true persistence is low; see columns (1)–(6). When the number of individuals is large and true persistence is low, optimal weighting results in a small downward bias that does not vary much with changes in the model variances. The biases are positive but still not large for equal and diagonal weighting and go down when the variance of transitory shocks is lower or true variance of fixed effects is higher—Table B-1.

The biases are typically small for the variances of permanent and transitory shocks, especially so when N is large—Tables B-2–B-3. For example, the biggest bias for the variance of permanent shocks is about 8 percent when the true persistence is low, N is small, and the last fifteen periods are used in estimation—column (6) of Table B-2. Thus for the variance of permanent shocks of 0.01, its biased estimate using quasidifferences is 0.0108. The biases vary with the size of the model parameters, but these effects are small. The variance of transitory shocks is biased downward for all the experiments we considered, although those biases are very small and become negligible when N is large—see the estimated constants in Ta-

[1]This is expected since the variance of fixed effects is not identified when the true persistence equals one.

ble B-3. Biases typically become smaller when the variance of fixed effects is higher and the true variance of transitory shocks is smaller.

B.2 Tables

Table B-1: Bias in the estimated variance of fixed effects. Regression analysis

	$\rho = 0.90$						$\rho = 0.995$					
Weighting	Equal		Optimal		Diagonal		Equal		Optimal		Diagonal	
$t =$	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: $N = 1000$												
σ^2_α	-5.130***	-6.300***	-0.458	-9.434***	-4.514***	-6.701***	-32.439***	-78.226***	-107.474***	-67.056***	-40.088***	-60.505***
	(0.500)	(0.872)	(1.561)	(2.773)	(0.427)	(0.685)	(3.262)	(4.280)	(11.234)	(16.530)	(3.123)	(4.987)
σ^2_η	9.473***	8.570***	4.074***	6.716**	11.039***	17.222***	59.315***	14.237**	-9.208	-7.331	102.283***	46.486***
	(0.914)	(1.722)	(1.379)	(2.605)	(0.647)	(1.094)	(4.855)	(6.629)	(7.587)	(10.352)	(3.321)	(5.691)
σ^2_ϵ	17.496***	26.395***	0.364	-0.972	14.330***	18.515***	12.107**	10.773	-12.222	-10.482	-14.513***	17.329***
	(0.914)	(1.722)	(1.379)	(2.605)	(0.647)	(1.094)	(4.855)	(6.629)	(7.587)	(10.352)	(3.321)	(5.691)
Const.	19.580***	24.553***	-14.733***	-23.638***	27.716***	40.514***	269.445***	344.895***	204.682***	162.624***	266.978***	264.997***
	(0.758)	(1.415)	(1.282)	(2.386)	(0.541)	(0.909)	(4.100)	(5.588)	(7.989)	(11.315)	(2.969)	(5.061)
Panel B: $N = 10000$												
σ^2_α	-5.956***	-7.818***	0.822	-0.292	-5.964***	-7.552***	-56.570***	-98.346***	-133.699***	-125.736***	-33.838***	-80.047***
	(0.307)	(0.415)	(0.730)	(1.388)	(0.265)	(0.359)	(4.024)	(4.187)	(10.625)	(17.472)	(3.801)	(4.331)
σ^2_η	0.419	0.624	0.330	-0.448	3.565***	4.757***	-42.974***	-56.388***	-95.197***	-50.806***	14.416**	-31.424***
	(0.448)	(0.781)	(0.544)	(1.198)	(0.374)	(0.691)	(6.371)	(8.976)	(7.392)	(11.464)	(5.629)	(8.090)
σ^2_ϵ	8.272***	11.542***	-1.176**	-3.183***	9.807***	14.355***	6.915	5.040	-2.413	-7.924	31.126***	13.442*
	(0.448)	(0.781)	(0.544)	(1.198)	(0.375)	(0.691)	(6.371)	(8.976)	(7.392)	(11.464)	(5.629)	(8.090)
Const.	10.882***	15.252***	-3.455***	-7.245***	17.818***	24.678***	269.569***	319.788***	279.107***	236.318***	272.102***	303.122***
	(0.383)	(0.648)	(0.543)	(1.111)	(0.321)	(0.567)	(5.355)	(7.368)	(7.673)	(12.206)	(4.737)	(6.688)
No. obs.	1000	500	1000	500	1000	500	1000	500	1000	500	1000	500

Notes: The table contains the results from a regression of the bias in the variance of fixed effects measured in percent, $100 \cdot (\hat{\sigma}^2_\alpha - \sigma^2_\alpha)/\sigma^2_\alpha$, on the standardized variances of fixed effects, σ^2_α, permanent shocks, σ^2_η, and transitory shocks, σ^2_ϵ. Columns t =[1–15], [1–30] (t =[16–30]) utilize estimation data based on the first fifteen or thirty (last fifteen) observations from Tables 3.1–3.2. Standard errors are in parentheses. *** (**) [*] significant at the 1% (5%) [10%] level.

Table B-2: Bias in the estimated variance of permanent shocks. Regression analysis

	$\rho = 0.90$						$\rho = 0.995$					
Weighting	Equal		Optimal		Diagonal		Equal		Optimal		Diagonal	
$t =$	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: $N = 1000$												
σ^2_α	-1.124***	-1.215***	0.133**	-0.877***	-1.532***	-3.010***	-0.048	0.250**	0.195***	0.068	-0.033	-0.505**
	(0.077)	(0.134)	(0.056)	(0.222)	(0.093)	(0.479)	(0.067)	(0.111)	(0.034)	(0.077)	(0.122)	(0.229)
σ^2_η	-0.061	-0.178**	-0.697***	1.628***	-0.405***	2.145***	0.202***	0.273***	-0.685***	-0.047	0.111**	1.195***
	(0.044)	(0.074)	(0.045)	(0.215)	(0.047)	(0.324)	(0.036)	(0.064)	(0.035)	(0.076)	(0.051)	(0.114)
σ^2_ϵ	0.624***	0.359***	0.116**	-0.446**	1.016***	2.719***	0.129***	0.237***	-0.009	0.204***	0.588***	1.248***
	(0.044)	(0.074)	(0.045)	(0.215)	(0.047)	(0.324)	(0.036)	(0.064)	(0.035)	(0.076)	(0.051)	(0.114)
Const.	1.365***	1.224***	-0.898***	2.636***	1.996***	8.226***	0.162***	0.060	-1.048***	-0.039	0.523***	2.374***
	(0.050)	(0.086)	(0.044)	(0.201)	(0.057)	(0.342)	(0.043)	(0.073)	(0.032)	(0.070)	(0.070)	(0.140)
Panel B: $N = 10000$												
σ^2_α	-0.195***	-0.265***	-0.001	-0.425***	-0.419***	-2.011***	0.115***	0.184***	0.028**	-0.018	0.066***	0.191***
	(0.025)	(0.043)	(0.014)	(0.125)	(0.028)	(0.282)	(0.022)	(0.035)	(0.011)	(0.029)	(0.023)	(0.057)
σ^2_η	-0.074***	-0.029	-0.064***	1.298***	-0.125***	1.265***	0.167***	0.194***	-0.001	0.029	0.191***	0.381***
	(0.013)	(0.021)	(0.011)	(0.166)	(0.014)	(0.210)	(0.012)	(0.022)	(0.010)	(0.021)	(0.012)	(0.044)
σ^2_ϵ	0.126***	0.054***	0.004	0.127	0.243***	1.096***	0.006	-0.011	0.013	0.036*	0.024*	0.226***
	(0.013)	(0.021)	(0.011)	(0.166)	(0.014)	(0.210)	(0.012)	(0.022)	(0.010)	(0.021)	(0.012)	(0.044)
Const.	0.259***	0.240***	-0.093***	2.289***	0.522***	4.664***	-0.077***	-0.143***	-0.099***	0.041*	-0.085***	0.080*
	(0.016)	(0.026)	(0.011)	(0.144)	(0.017)	(0.214)	(0.014)	(0.024)	(0.010)	(0.021)	(0.014)	(0.044)
No. obs.	1000	500	1000	500	1000	500	1000	500	1000	500	1000	500

Notes: The table contains the results from a regression of the bias in the variance of permanent shocks measured in percent, $100 \cdot (\hat{\sigma}^2_\eta - \sigma^2_\eta)/\sigma^2_\eta$, on the standardized variances of fixed effects, σ^2_α, permanent shocks, σ^2_η, and transitory shocks, σ^2_ϵ. Columns t =[1–15], [1–30] (t =[16–30]) utilize estimation data based on the first fifteen or thirty (last fifteen) observations from Tables 3.1–3.2. Standard errors are in parentheses. *** (**) [*] significant at the 1% (5%) [10%] level.

Table B-3: Bias in the estimated variance of transitory shocks. Regression analysis

	$\rho = 0.90$						$\rho = 0.995$					
Weighting	Equal		Optimal		Diagonal		Equal		Optimal		Diagonal	
$t =$	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]	[1–15], [1–30]	[16–30]
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A: $N = 1000$												
σ^2_α	0.810***	0.917***	0.771***	0.653***	0.806***	0.895***	0.109*	-0.096	0.710***	0.311***	0.125**	0.009
	(0.063)	(0.108)	(0.058)	(0.082)	(0.064)	(0.109)	(0.058)	(0.092)	(0.058)	(0.070)	(0.059)	(0.094)
σ^2_η	-0.016	0.024	-0.017	-0.057	-0.045	-0.057	0.032	-0.023	0.054*	-0.042	0.068**	-0.022
	(0.035)	(0.055)	(0.034)	(0.050)	(0.035)	(0.056)	(0.027)	(0.045)	(0.030)	(0.039)	(0.028)	(0.047)
σ^2_ϵ	-0.505***	-0.437***	-0.702***	-0.385***	-0.494***	-0.479***	-0.141***	-0.216***	-0.660***	-0.520***	-0.187***	-0.327***
	(0.035)	(0.055)	(0.034)	(0.050)	(0.035)	(0.056)	(0.027)	(0.045)	(0.030)	(0.039)	(0.028)	(0.047)
Const.	-1.074***	-1.122***	-1.689***	-1.202***	-1.151***	-1.348***	-0.324***	-0.274***	-1.661***	-1.137***	-0.427***	-0.546***
	(0.040)	(0.067)	(0.038)	(0.055)	(0.041)	(0.068)	(0.035)	(0.055)	(0.037)	(0.045)	(0.036)	(0.057)
Panel B: $N = 10000$												
σ^2_α	0.185***	0.257***	0.074***	0.083***	0.325***	0.315***	0.025	0.028	0.105***	0.085***	0.113***	0.108***
	(0.020)	(0.035)	(0.016)	(0.024)	(0.020)	(0.033)	(0.018)	(0.029)	(0.012)	(0.020)	(0.018)	(0.032)
σ^2_η	0.051***	0.007	0.002	0.001	0.065***	0.046**	-0.011	-0.014	-0.016**	-0.026**	-0.012	0.004
	(0.011)	(0.017)	(0.009)	(0.015)	(0.010)	(0.018)	(0.009)	(0.015)	(0.007)	(0.012)	(0.008)	(0.014)
σ^2_ϵ	-0.134***	-0.131***	-0.073***	-0.023	-0.243***	-0.231***	-0.033***	-0.033**	-0.077***	-0.061***	-0.054***	-0.067***
	(0.011)	(0.017)	(0.009)	(0.015)	(0.010)	(0.018)	(0.009)	(0.015)	(0.007)	(0.012)	(0.008)	(0.014)
Const.	-0.254***	-0.296***	-0.171***	-0.098***	-0.499***	-0.492***	-0.108***	-0.106***	-0.210***	-0.182***	-0.173***	-0.198***
	(0.013)	(0.021)	(0.010)	(0.016)	(0.013)	(0.021)	(0.011)	(0.018)	(0.008)	(0.013)	(0.011)	(0.019)
No. obs.	1000	500	1000	500	1000	500	1000	500	1000	500	1000	500

Notes: The table contains the results from a regression of the bias in the variance of transitory shocks measured in percent, $100 \cdot (\hat{\sigma}^2_\epsilon - \sigma^2_\epsilon)/\sigma^2_\epsilon$, on the standardized variances of fixed effects, σ^2_α, permanent shocks, σ^2_η, and transitory shocks, σ^2_ϵ. Columns t =[1–15], [1–30] (t =[16–30]) utilize estimation data based on the first fifteen or thirty (last fifteen) observations from Tables 3.1–3.2. Standard errors are in parentheses. *** (**) [*] significant at the 1% (5%) [10%] level.

B.3 Monte Carlo and Minimum Distance Estimation Procedure

The schematic diagram below implements the minimum distance procedure and the Monte Carlo simulations. These estimates can then be used to compute the bias given true initial parameters.

Algorithm 1 Monte Carlo Minimum Distance Procedure

i) Set rhogrid with K points where $\rho \in [0.5, 1.1]$, we set $\mathbf{K} = \mathbf{100}$

ii) Suppose you want to check bias for $\mathbf{N} = 1000$ and true ρ in data takes values $\rho = \{0.9, 0.95, 0.995\}$

iii) create vector rhotrue $= \{0.9, 0.95, 0.995\}$, which means $\mathbf{R} = \mathbf{3}$

for $r = 1, \cdots, \mathrm{R}$ **do**

 Set true rho value to use to simulate data $\rho =$ rhotrue$[r]$

 Set the other true parameters $\{\sigma_\eta^2, \sigma_\alpha^2, \sigma_\epsilon^2\}$ such that $\Phi = \{\rho, \sigma_\eta^2, \sigma_\alpha^2, \sigma_\epsilon^2\}$

 Fix beginning seed number

 for $s = 1, \cdots, \mathrm{S}$ **do**

  a) Set seed at each iteration: seed $=$ seed $+$ (s $\times$ **100**)

  b) Simulate lifecycle panel data using Φ, with $z_{i,o} = 0$ or $z_{i,o} = \frac{\sigma_\eta^2}{1-\rho^2}$

  for $k = 1, \cdots, \mathrm{K}$ **do**

   1) Set $\tilde{\rho} =$ rhogrid$[k]$

   2) Calculate quasidifferences using the $\hat{\rho}$ from the grid

   3) Calculate data moments $\mathbf{m^d}$

   4) Derive variance-covariance matrix $\mathbf{V}$

   5) Calculate weighting matrix $\mathbf{W}$ using $\mathbf{V}$ (OMD and DWMD) or identity matrix with large value on the diagonal (EWMD)

   6) Create objective function to minimize the distance between weighted data and theoretical moments $m(\mathbf{\Phi}) - \mathbf{m^d}$ such that:

   $$\arg\min_{\tilde{\Phi}} \left[m(\mathbf{\Phi}) - \mathbf{m^d}\right]' \mathbf{W} \left[m(\mathbf{\Phi}) - \mathbf{m^d}\right]$$

   7) Optimize using global search algorithm

   8) Store the optimal estimates of $\tilde{\mathbf{\Phi}} = \{\tilde{\rho}, \tilde{\sigma}_\eta^2, \tilde{\sigma}_\alpha^2, \tilde{\sigma}_\epsilon^2\}$, and the value of the objective function in a matrix **GK**

  c) Reorder the matrix **GK** in ascending order using the values of the objective function and pick the first row.

  d) Store results from step c) in matrix **GS**

 e) Use matrix **GS** to derive $\hat{\Phi} = \frac{1}{S}\sum_{s=1}^{S} \tilde{\Phi}$ & standard errors for each ρ

 f) Store results from step e) in matrix **RS**

iii) Calculate bias in parameters as $\hat{\mathbf{\Phi}} - \mathbf{\Phi}$ using **RS**

Appendix C: The Role Private and Public Transfers for Consumption Insurance in South Africa

C.1 Figures

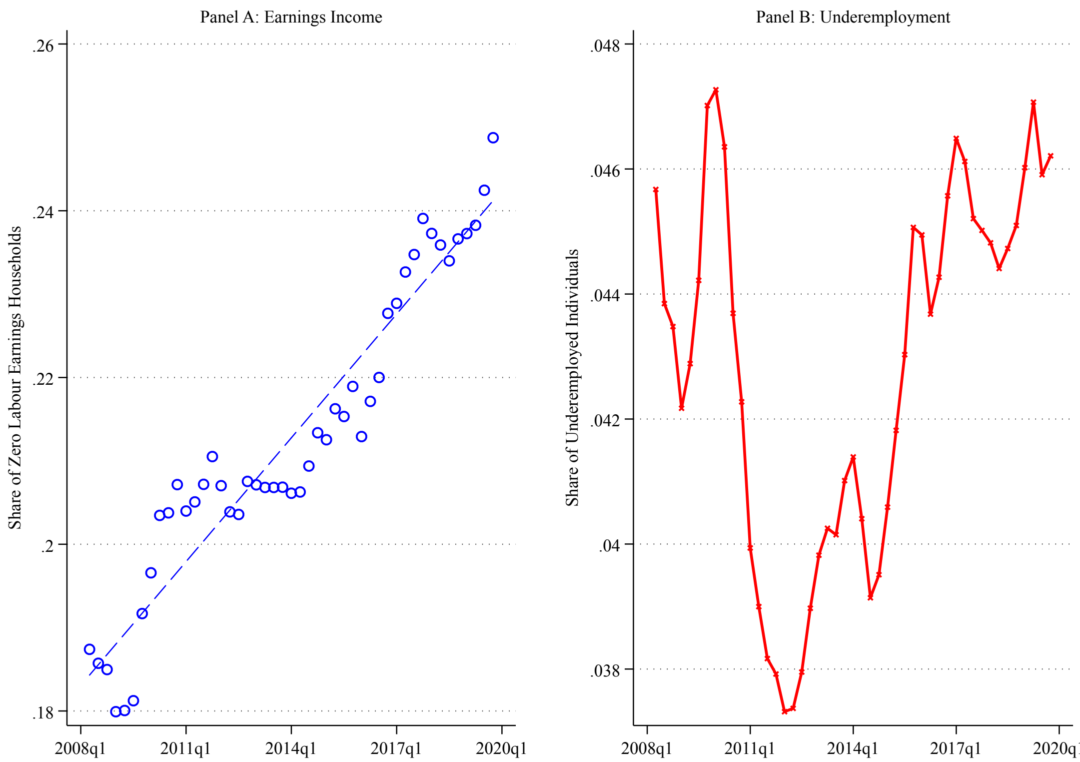

Figure C-1: Zero Labour Earnings Households & Underemployed Individuals

Notes: Panel A depicts the share of households with zero income from the labour market (no one formally or informally employed). Panel B plots the share of employed individuals who feel that their current employment does not match their qualification. *Data source:* Quarterly Labour Force Survey, 2008-2020.

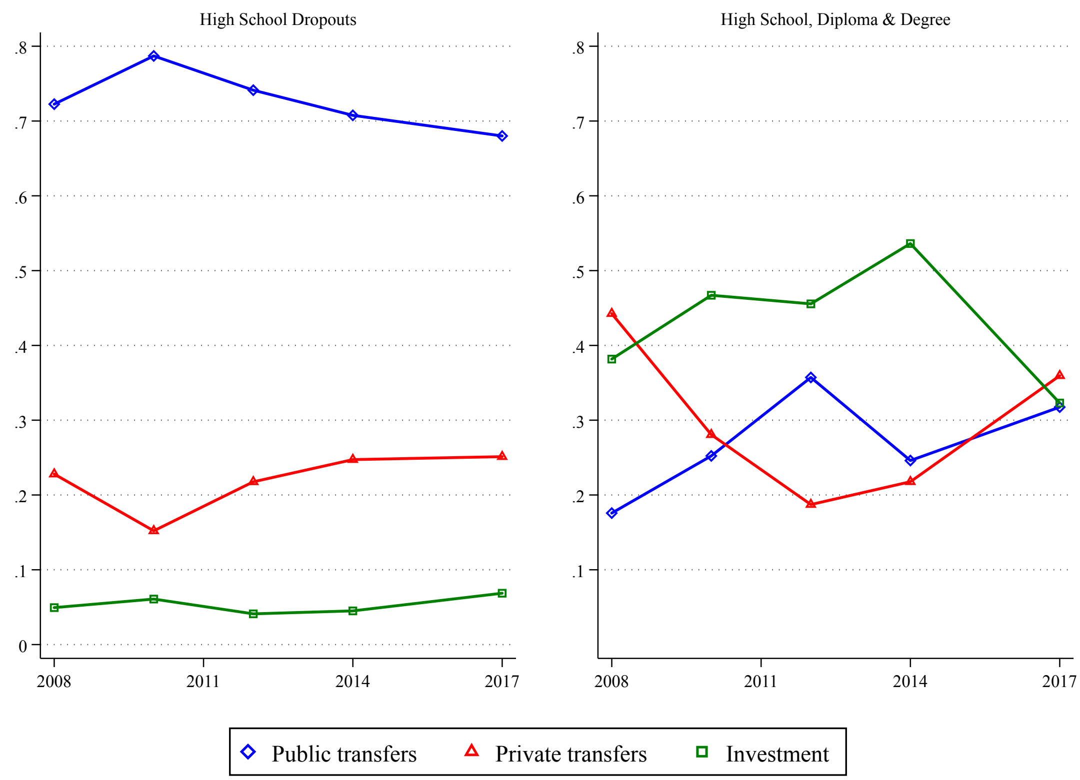

Figure C-2: Household Income Shares for Households with Zero Labour Income by Education

Notes: The figure depicts the proportions of public transfers (social grants), private transfers from family and friends, and investment in household income for households without labour income (agriculture farming income excluded). The shares are computed from amounts in real Rand terms, deflated using a monthly CPI corresponding to the month and year of the interview.

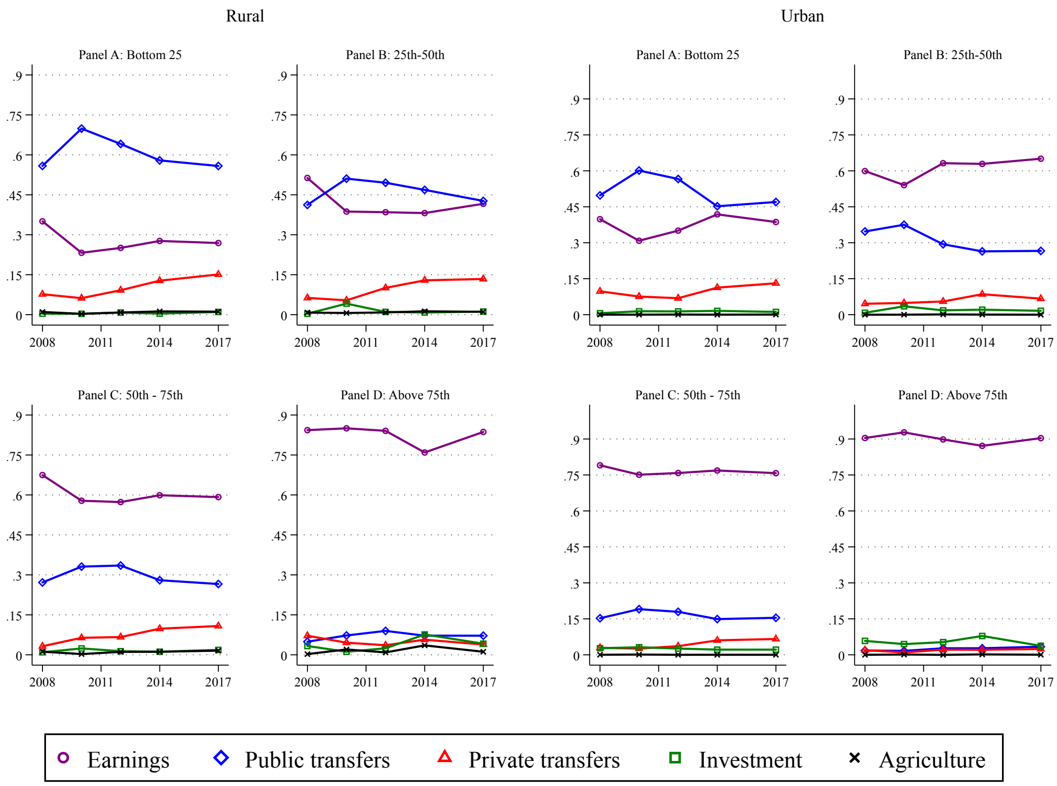

Figure C-3: Source of Household Income by Income Percentile and Region of Residence

Notes: The panels show trends in the share of earnings, investment, receipts from social assistance, agriculture, and private transfers from family and friends in household income by region of residence and income percentile. The household income aggregates resources from the five income sources. All amounts are in real Rand terms, deflated using a monthly CPI corresponding to the month and year of the interview. *Data source:* National Income Dynamics Survey (NIDS), 2008-2017.

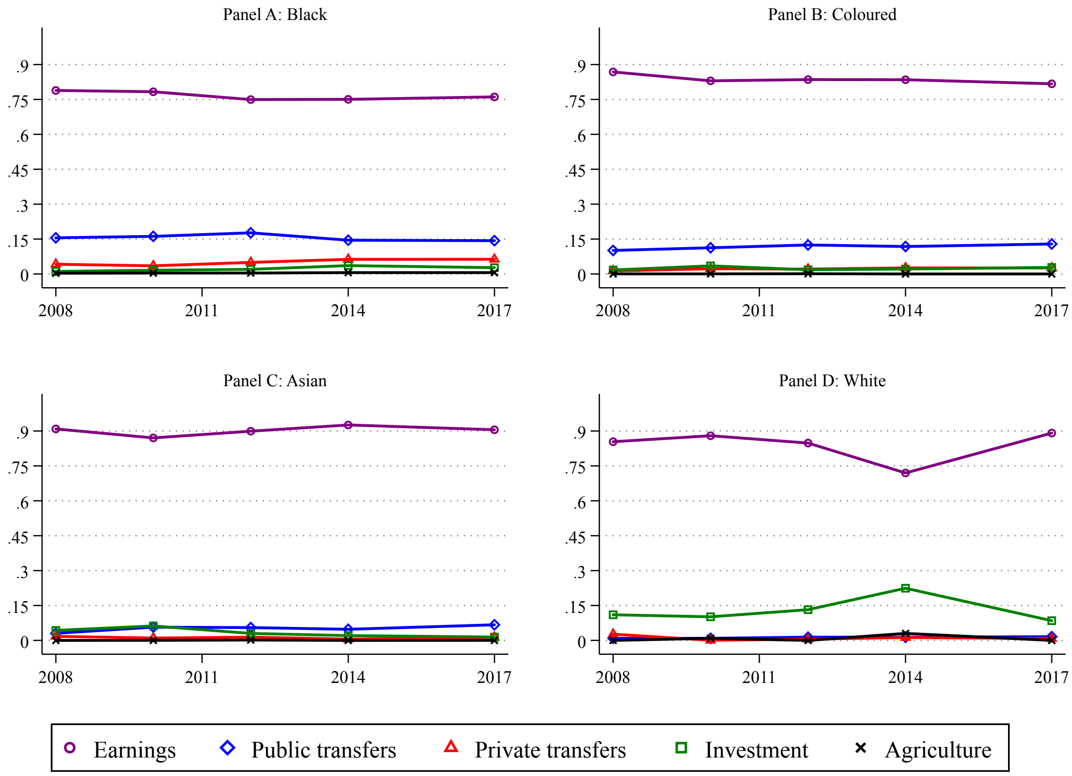

Figure C-4: Source of Household Income by Race

Notes: The panels show trends in the share of earnings, investment, receipts from social assistance, agriculture, and private transfers from family and friends in household income by race. The household income aggregates resources from the five income sources. All amounts are in real Rand terms, deflated using a monthly CPI corresponding to the month and year of the interview. *Data source:* National Income Dynamics Survey (NIDS), 2008-2017.

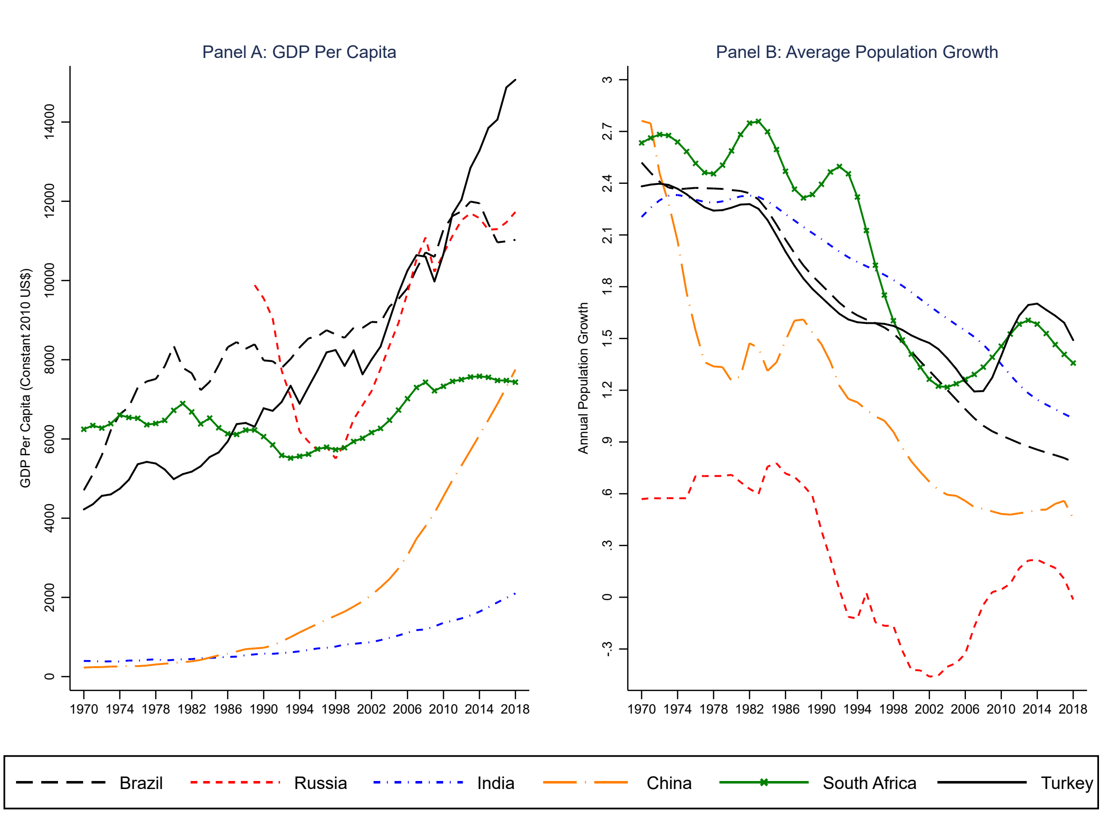

Figure C-5: Per Capita Income and Population Growth in Emerging Economies

Notes: The figure depicts the average GDP per capita and population growth among emerging economies popularly known as BRICS (Brazil, Russia, India, China, and South Africa) and Turkey. Panel A profiles GDP per capita computed as GDP divided by midyear country population size where the GDP is in constant 2010 U.S.$, and is the sum of domestic gross value added (plus product taxes less subsidies) before adjustments for depreciation. Similarly, annual population growth rates in Panel B are the exponential rate of growth of midyear population (all residents) between two consecutive years. The data source is the World Bank's World Development Indicators.

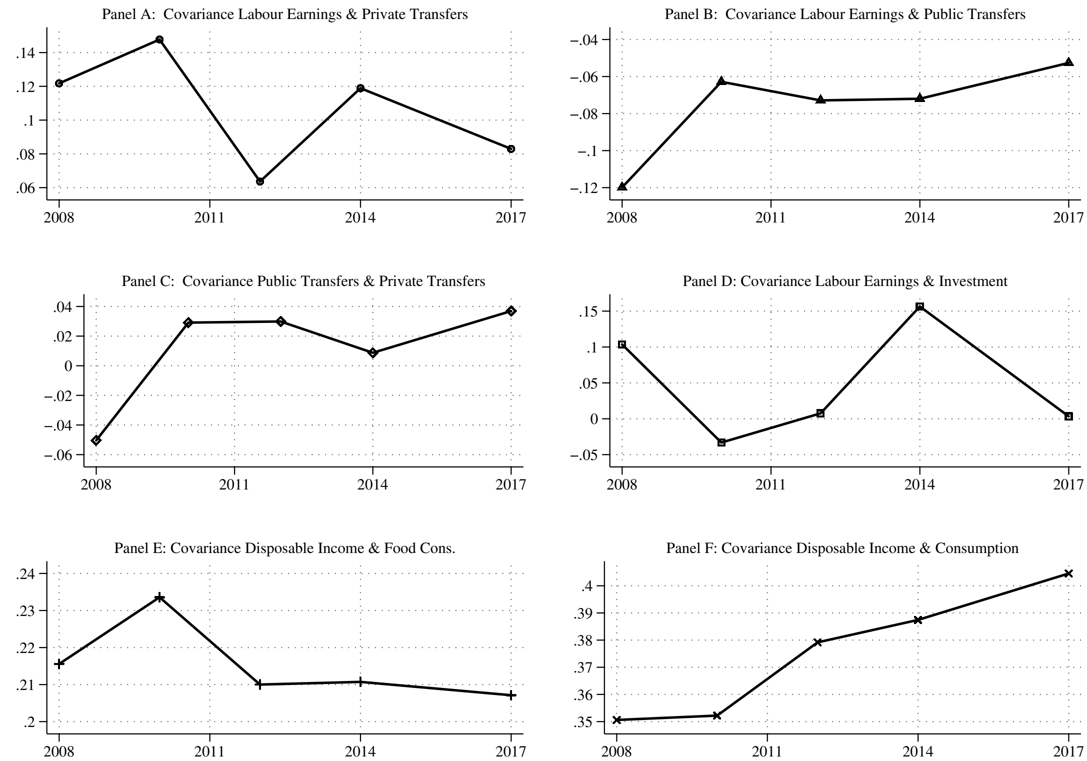

Figure C-6: Covariance Between Income Components and Consumption

Notes: The figure depicts the covariance between labour earnings, private transfers, public transfers and investment income. The bottom panels plot the covariance of food and nondurable consumption with disposable household income.

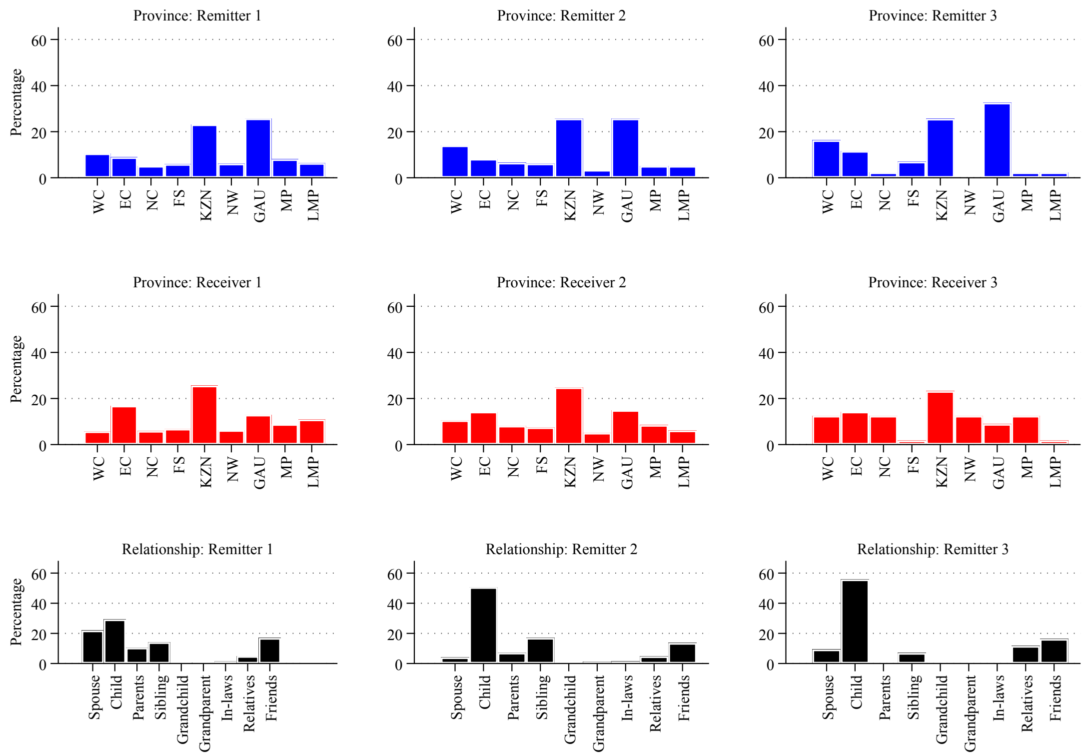

Figure C-7: Relationship of Remitter to the Receiving Household

Notes: The figure presents the relationship between remitters (senders of private transfers) and the households that receive the private transfers as well as source of remittances by province. Private transfers from spouse is money sent by husband or wife who live outside the household; child includes transfers sent by biological, adopted, or foster children; relatives represents income from any other relatives, which includes uncles, aunts, cousins, e.t.c.

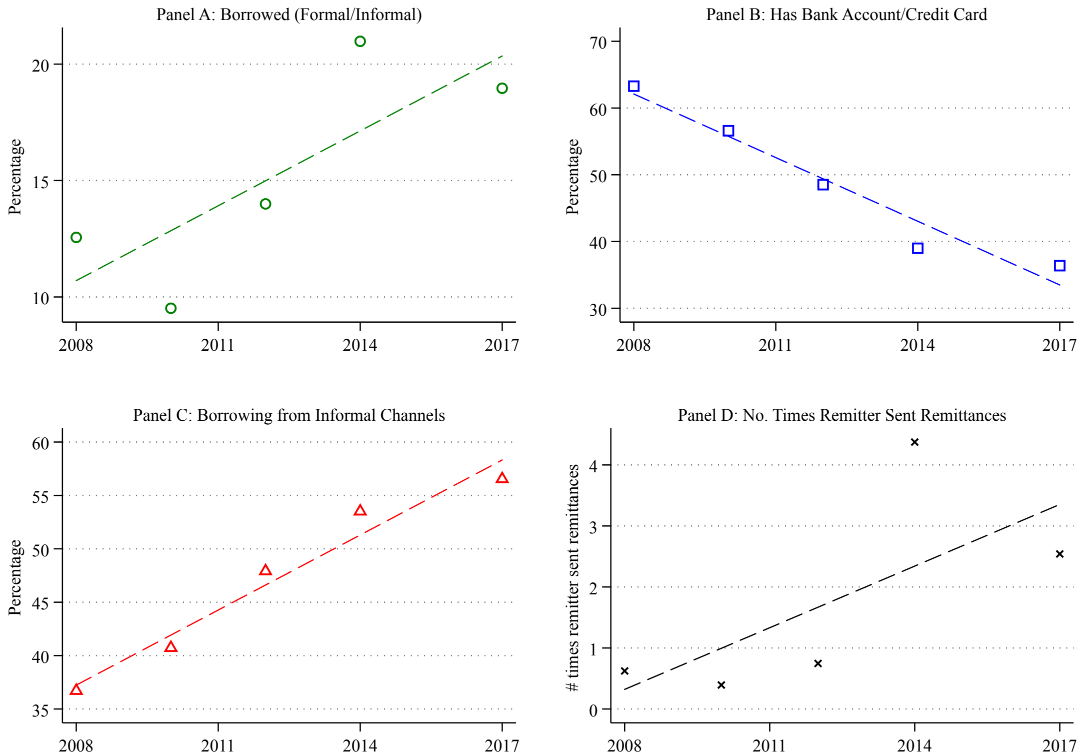

Figure C-8: Access to Formal and Informal Borrowing Channels

Notes: Panel A shows the share of households who borrowed using formal institutions or informal channels in the previous year. Panel B plots the share of individuals with a bank account or credit card. Among those who borrowed, Panel C depicts the share of individuals who borrowed using informal channels (from friends, loan sharks, etc.)

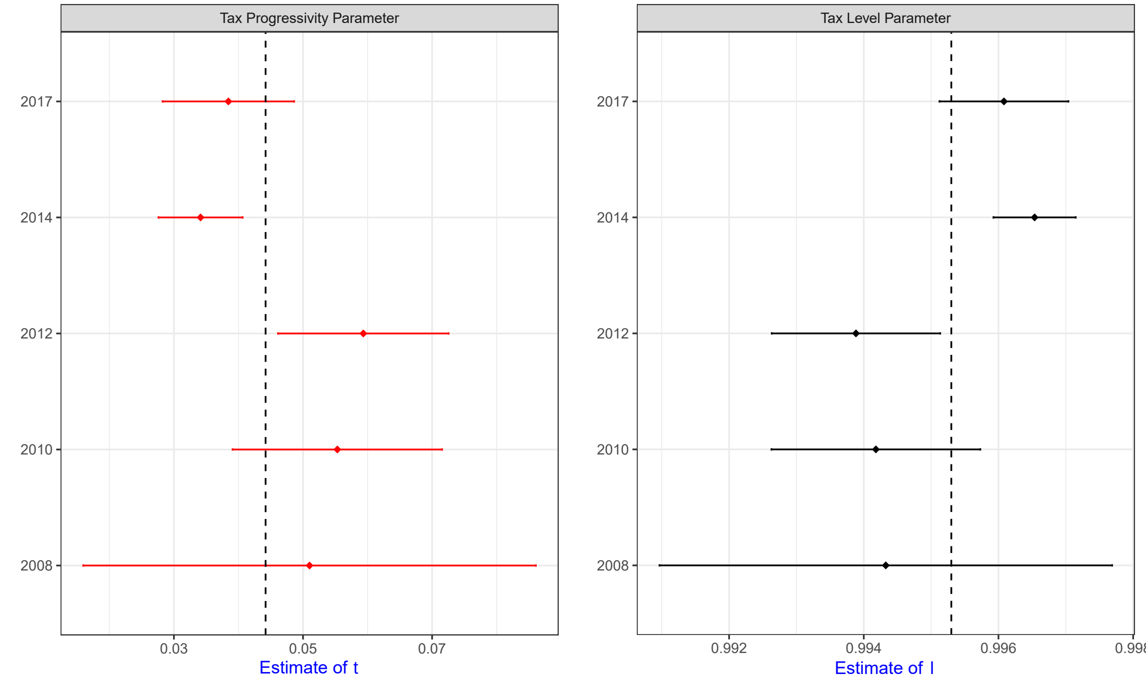

Figure C-9: Tax Function Parameters, 2008-2017

Notes: The panels report point estimates and 95% confidence intervals of τ and λ that measure the progressivity and average level of the tax system, respectively. The estimates are from ordinary least squares. Dotted line in each panel (dashed) corresponds to estimates from the pooled data (2008-2017).

C.2 Tables

Table C-1: The Role of Different Income Sources: Minimum-Distance Partial Insurance Estimates

Consumption:		Nondurable	Nondurable	Nondurable	Nondurable	Nondurable
Income:		Net income (Y)	Before tax	Y + fin. income	Excluding private transfers (R)	Excluding transfers (T)
Sample:		Baseline	Baseline	Baseline	Baseline	Baseline
		(1)	(2)	(3)	(4)	(5)
σ^2_ζ	2010-2012	0.0976	0.1144	0.1091	0.1199	0.0823
(Variance perm. shock)		(0.0060)	(0.0068)	(0.0066)	(0.0112	(0.0084)
	2014-2017	0.1219	0.1462	0.1392	0.1275	0.0980
		(0.0095)	(0.0101)	(0.0099)	(0.0119)	(0.0149)
σ^2_ϵ	2010	0.3900	0.3873	0.4005	0.3803	0.5489
(Variance trans. shock)		(0.0119)	(0.0124)	(0.0123)	(0.0163)	(0.0159)
	2012	0.3021	0.3009	0.2949	0.2952	0.4732
		(0.0142)	(0.0149)	(0.0155)	(0.0203)	(0.0208)
	2014-2017	0.1969	0.1692	0.1791	0.2536	0.3069
		(0.0105)	(0.0105)	(0.0114)	(0.0142)	(0.0159)
ϕ		0.6892	0.5822	0.6180	0.5388	0.7694
(Transmission perm. shock)		(0.0507)	(0.0401)	(0.0402)	(0.0527)	(0.1147)
ψ		0.0215	0.0297	0.0329	0.0715	0.0478
(Transmission trans. shock)		(0.0132)	(0.0137)	(0.0140)	(0.0130)	(0.0092)
Households		2680	2685	2679	2687	2680

Notes: This table reports the permanent and transitory parameters. Y represents disposable income; Fin., financial income; R, private transfers; T, transfer payments. The estimates are computed using diagonally weighted minimum distance (DWMD) estimation. Standard errors are in parentheses, and they were derived from via bootstrapping with 500 replications.

Table C-2: Minimum-Distance Partial Insurance Estimates for Different Consumption Aggregates

Consumption:		Nondurable	Food	Excl. hel + educ.	Necessities	Total Cons.
Income:		Net income	Net income	Net Income	Net income	Net income
Sample:		Baseline	Baseline	Baseline	Baseline	Baseline
		(1)	(2)	(3)	(4)	(5)
σ^2_η	2010-2012	0.0976	0.1033	0.0923	0.0955	0.0979
(Variance perm. shock)		(0.0060)	(0.0079)	(0.0063)	(0.0132)	(0.0060)
	2014-2017	0.1219	0.1079	0.1219	0.0876	0.1211
		(0.0095)	(0.0109)	(0.0097)	(0.0143)	(0.0095)
σ^2_ϵ	2010	0.3900	0.3796	0.3918	0.3933	0.3907
(Variance trans. shock)		(0.0119)	(0.0118)	(0.0116)	(0.0155)	(0.0120)
	2012	0.3021	0.3032	0.3070	0.3121	0.3085
		(0.0142)	(0.0149)	(0.0145)	(0.0164)	(0.0153)
	2014-2017	0.1969	0.2080	0.1966	0.2241	0.1976
		(0.0105)	(0.0122)	(0.0110)	(0.0145)	(0.0112)
ϕ		0.6892	0.4144	0.6293	0.5143	0.7252
(Transmission perm. shock)		(0.0507)	(0.0468)	(0.0480)	(0.0882)	(0.0559)
ψ		0.0215	0.0357	0.0768	0.0250	0.0598
(Transmission trans. shock)		(0.0132)	(0.0140)	(0.0121)	(0.0099)	(0.0145)
Households		2680	2675	2677	2597	2645

Notes: This table reports the permanent and transitory parameters. The estimates are computed using diagonally weighted minimum distance (DWMD) estimation. Standard errors are in parentheses, and they were derived via bootstrapping with 500 replications.

Table C-3: Minimum-Distance Partial Insurance Estimates by Education, Area of Residence & Race

	Cohort	Benchmark	NHS	HC	Rural	Urban	Black	Coloured	White	Black & NHS	Coloured & NHS	Young	Old
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
σ^2_η	2010-2012	0.0976	0.0936	0.0957	0.1187	0.0645	0.1160	0.0000	0.0792	0.1045	0.0093	0.1200	0.0856
(Variance perm. shock)		(0.0060)	(0.0074)	(0.0087)	(0.0100)	(0.0072)	(0.0097)	(0.0004)	(0.0215)	(0.0096)	(0.0125)	(0.0117)	(0.0092)
	2014-2017	0.1219	0.1204	0.1207	0.1286	0.1192	0.1192	0.1659	0.0868	0.1160	0.1367	0.1227	0.1642
		(0.0095)	(0.0115)	(0.0163)	(0.0106)	(0.0115)	(0.0101)	(0.0154)	(0.0248)	(0.0118)	(0.0209)	(0.0149)	(0.0107)
σ^2_ϵ	2010	0.3900	0.3919	0.3965	0.5459	0.3028	0.4298	0.4124	0.1770	0.4033	0.3880	0.4201	0.2106
(Variance trans. shock)		(0.0119)	(0.0123)	(0.0249)	(0.0301)	(0.0108)	(0.0181)	(0.0186)	(0.0214)	(0.0148)	(0.0356)	(0.0170)	(0.0161)
	2012	0.3021	0.3656	0.1900	0.3201	0.3213	0.3638	0.1830	0.2263	0.4750	0.1282	0.3002	0.1821
		(0.0142)	(0.0230)	(0.0158)	(0.0212)	(0.0199)	(0.0206)	(0.0154)	(0.0243)	(0.0240)	(0.0156)	(0.0167)	(0.0172)
	2014-2017	0.1969	0.2265	0.1582	0.1503	0.2269	0.2096	0.1005	0.2764	0.2641	0.1356	0.2848	0.0254
		(0.0105)	(0.0120)	(0.0191)	(0.0124)	(0.0135)	(0.0113)	(0.0113)	(0.0438)	(0.0146)	(0.0190)	(0.0182)	(0.0095)
ϕ		0.6892	0.7863	0.5793	0.8791	0.4820	0.8060	0.4423	0.3754	0.9039	0.5004	0.6873	0.4636
(Transmission perm. shock)		(0.0507)	(0.0734)	(0.0823)	(0.0702)	(0.0378)	(0.0708)	(0.0403)	(0.1966)	(0.0806)	(0.0803)	(0.0871)	(0.0309)
ψ		0.0215	0.0412	0.0085	0.0006	0.1017	0.0646	0.0000	0.0013	0.0492	0.0000	0.0615	0.0075
(Transmission trans. shock)		(0.0132)	(0.0156)	(0.0149)	(0.0034)	(0.0147)	(0.0157)	(0.0006)	(0.0114)	(0.0187)	(0.0000)	(0.0155)	(0.0323)
Households		2680	1694	986	1061	1652	1913	545	221	1298	347	1534	1391

Notes: This table reports the variance of permanent and transitory shocks, and the transmission magnitudes for the two shocks. The estimates are computed using diagonally weighted minimum distance (DWMD) estimation. Standard errors are in parentheses, and they were derived via bootstrapping with 500 replications.

Table C-4: Minimum-Distance Partial Insurance Estimates by Marital Status and Household Head Age

Cohort		Benchmark	Married	Cont. Married	Balanced	Male Heads	Age 30-65
		(1)	(2)	(3)	(4)	(5)	(6)
σ_η^2	2010-2012	0.0976	0.0906	0.0818	0.0658	0.0960	0.0952
(Variance perm. shock)		0.0060	0.0058	0.0057	0.0111	0.0063	0.0061
	2016-2016	0.1219	0.1176	0.1113	0.0697	0.1197	0.1221
		0.0095	0.0095	0.0096	0.0121	0.0095	0.0095
σ_ϵ^2	2010	0.3900	0.3583	0.3540	0.4465	0.3950	0.3995
(Variance trans. shock)		0.0119	0.0125	0.0139	0.0212	0.0134	0.0129
	2012	0.3021	0.2923	0.3202	0.2178	0.3055	0.3060
		0.0142	0.0136	0.0165	0.0234	0.0139	0.0149
	2014-2017	0.1969	0.1973	0.2115	0.1323	0.1990	0.2020
		0.0105	0.0117	0.0110	0.0135	0.0113	0.0106
ϕ		0.6892	0.7349	0.7847	0.5723	0.6973	0.6750
(Transmission perm. shock)		0.0507	0.0554	0.0638	0.1233	0.0526	0.0511
ψ		0.0215	0.0442	0.0761	0.0007	0.0372	0.0451
(Transmission trans. shock)		0.0132	0.0130	0.0140	0.0049	0.0135	0.0139

Notes: This table reports the permanent and transitory parameters. The estimates are computed using diagonally weighted minimum distance (DWMD) estimation. Standard errors are in parentheses, and they were derived from via bootstrapping with 500 replications.