Essays in Quantitative Macroeconomics

Essays in Quantitative Macroeconomics

By
THOMAS PALMER,
M.A. Economics (University of Western Ontario)
B.A. Economics (University of British Columbia)

A Thesis Submitted to the School of Graduate Studies
in the Partial Fulfillment of the Requirements for the Degree

Doctor of Philosophy
in
Economics

McMaster University
Hamilton, Ontario

McMaster University
Doctor of Philosophy (2024)
Economics
Hamilton, ON, Canada

TITLE: Essays in Quantitative Macroeconomics

AUTHOR:
Thomas Palmer,
M.A. (University of Western Ontario)
B.A. (University of British Columbia)

SUPERVISORS:
Dr. Pau S. Pujolas
Associate Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

Dr. Bettina Brüggemann
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

SUPERVISORY COMMITTEE MEMBERS:
Dr. Zachary L. Mahone
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

Dr. Gajendran Raveendranathan
Assistant Professor, Department of Economics
McMaster University, Hamilton, ON, Canada

NUMBER OF PAGES: xi, 107

Abstract

This thesis comprises three papers in quantitative macroeconomics that explore the following questions: (1) How does employer-provided training impact the college wage premium in the context of skill-biased technological change? (2) How does the option to sell a firm influence firm entry, exit, and growth dynamics? (3) How does college major selection impact occupational sorting and entrepreneurship? Chapter 1 combines matched employer-employee survey data from Canada with a quantitative model of the labour market featuring endogenous technology and training decisions to show that the rise in training, driven by technological advancements, attenuated the increase in the college wage premium by 63 percent between 1980 and the early 2000s. Chapter 2, co-authored with Bettina Brüggemann and Zachary Mahone, uses administrative matched employer-employee data from Canada and a quantitative model of firm dynamics to establish that transfers of business ownership significantly impact firm entry, exit, and growth dynamics, with 13 percent of new entrants surviving solely due to the option value of sale. Chapter 3 empirically establishes a negative relationship between STEM majors and entrepreneurship using micro-data from the 1997 National Longitudinal Survey of Youth. Through a quantitative model that links decisions regarding majors and entrepreneurship, I show that lowering STEM tuition increases STEM enrolment at the cost of reducing overall entrepreneurial activity.

Acknowledgements

The journey of pursuing my Ph.D. has been both challenging and rewarding, and I am deeply grateful to the many individuals who have supported and inspired me along the way.

First and foremost, I would like to express my sincerest gratitude to my co-supervisors, Pau Pujolas and Bettina Brüggemann, and my supervisory committee members, Zachary Mahone and Gajendran Raveendranthan, for their unwavering support and guidance throughout my time at McMaster. Their contributions extended far beyond academic mentorship, providing a solid foundation of personal support during this challenging pursuit. Pau went above and beyond to help me navigate the intricacies of academic research, generously welcoming me into his home and dedicating countless hours to discussions on various topics. This speaks volumes about his character and commitment to my success. I owe a tremendous debt of gratitude to Bettina and Zach for taking the risk and giving me my first opportunity to delve into quantitative macroeconomics research. Many of my fondest memories at McMaster were spent working alongside them as their research assistant and co-author. Gajen's extraordinary ability to articulate complex computational methods and his insightful questions constantly inspired me and pushed me to become a better economist. The depth of my gratitude to each of them far exceeds what these words can convey; their impact on my academic journey and personal growth is truly immeasurable.

I am also indebted to the many remarkable professors at McMaster who, while not part of my supervisory committee, were instrumental in my growth. Professors Alok Johri, Arthur Sweetman, Adam Lavecchia, Stephen Jones, Chris Muris, Angela Zheng, Katherine Cuff, Jeffrey Racine, Youngki Shin, Bradley Ruffle, Michael Veall, Marc-Andre Letendre, Jevan Cherniwchan, and many others always welcomed the opportunity to share their knowledge and advice. I am deeply grateful for their support. I must also extend my heartfelt thanks to Cynthia Zhao, Emma Beamson, Lihua Qian, Mena Petta-Lorenzini, Peter Kitchen, and Li Wang, whose assistance with numerous administrative tasks has significantly eased my journey. Their support has been invaluable, and I am truly grateful for their help.

One of the greatest rewards of my Ph.D. journey has been the friendships I forged with my peers. Andrew Leal, Oliver Loertscher, Daniel Tingskou, and Rabiul Islam each played a significant role in making this challenging experience more manageable and enjoyable. Their camaraderie, humour, and encouragement brought balance to my days, and I am grateful for the invaluable support and countless moments of laughter we shared. Their friendship has been a true highlight of this journey, and I feel fortunate to have had such exceptional people to share this process with.

Last but certainly not least, I want to express my heartfelt gratitude to my parents, Mark and Linda, and my brother, Patrick, for being a constant source of inspiration and unwavering support throughout this journey. Their encouragement and belief in me have been the driving force that propelled me forward, providing the strength I needed to navigate the many challenges of this experience. I am truly grateful for their love and support, and I dedicate this thesis to them.

Contents

List of Figures

List of Tables

Declaration of Academic Achievement

I, Thomas Palmer, declare that this thesis titled, **Essays in Quantitative Macroeconomics**, and works presented in it are my own. Chapters 1 and 3 are solo-authored papers. For these chapters, I independently collected the data, performed the analyses, and wrote the manuscripts. Chapter 2 is co-authored with Professor Bettina Brüggemann and Professor Zachary Mahone. We contributed equally to the data collection, analysis, and writing efforts for this chapter.

Introduction

Technological advancements, evolving labour market demands, and new entrepreneurial opportunities are transforming the choices individuals make in education, careers, and business. As these forces continue to reshape traditional career pathways and firm dynamics, understanding the mechanisms driving these changes is crucial for crafting policies that support inclusive growth and adaptability. This thesis contributes to this understanding through three independent yet interconnected papers that examine the effect of employer-provided training on wage inequality, the impact of ownership changes on firm dynamics, and the relationship between college majors and post-graduate career outcomes.

The papers in this thesis are grounded in a common methodological framework that integrates micro-level data on individuals and firms with heterogeneous-agent macroeconomic theory. This approach—combining empirical data with quantitative models—aims to provide a deeper understanding of economic phenomena and address key policy-relevant questions. Specifically, the thesis investigates: (1) the impact of employer-provided training on the college wage premium amidst skill-biased technological change; (2) the influence of firm sales and ownership transfers on entry, exit, and growth dynamics over a firm's life cycle; and (3) the effect of college major selection on occupational sorting and entrepreneurship following graduation. By examining these issues, this thesis provides novel insights into wage inequality, business dynamism, and entrepreneurial activity, offering guidance for future policies related to education financing, skill development, and business creation.

Chapter 1, titled "Skill-Biased Technological Change, Training, and the College Wage Premium," quantifies the impact of employer-provided training on the college wage premium in Canada between 1980 and the early 2000s. While a large body of research attributes the rise in the college wage premium—the wage gap between college and non-college educated workers—to the introduction of new information-and-communications technologies like computers and the Internet, I introduce novel evidence highlighting the significant role of employer-provided training in shaping this trend. In particular, I show that, in Canada, training participation among the working age population increased by more than 40 percent, with the largest increase in training experienced by workers without a college degree. In a quantitative model with endogenous technology adoption and training decisions, I show that the aggregate response of training to technology adoption attenuated the increase in the college wage premium by 63 percent over this period. This finding suggests that training can counteract the wage inequality generated by skill-biased technological change, offering valuable insights for education and workforce development policies.

Chapter 2, titled "Firm Sales and the Firm Life Cycle" and jointly written with Bettina Brüggemann and Zachary Mahone, explores the under-examined role of firm sales and ownership transfers in shaping entry, exit, and growth dynamics over the life cycle of a firm. A large quantitative literature has established that firm entry, exit, and growth are key drivers of aggregate productivity, output, and growth, and that there is substantial variation in exit and growth dynamics across individual firms. Yet, a key margin that has so far been overlooked in these analyses is the option to sell a firm; that is, the notion that businesses can—and often do—continue despite owners exiting the market, and vice versa. Using administrative matched employer-employee data on firms and business owners in Canada, we infer sales through changes in ownership and transfers of business equity. Empirically, we find that approximately 1.5 percent of firms are sold annually in Canada between 2000 and 2017, which is larger than annual exit rates for full-time employer businesses. Furthermore, we show that sales are risky: firms that are sold experience a persistently greater likelihood of exiting in the years following a sale; yet, conditional on survival, the average firm experiences a sustained increase in profits. Embedding firm sales into a quantitative model of firm dynamics, we find that 13 percent of new entrants survive exclusively due to the option value of sale and that realized ownership changes account for 18 percent of average log employment growth among small firms. These findings have important implications for policies supporting business growth and entrepreneurship.

Chapter 3, titled "College Majors, Occupations, and Entrepreneurship," quantitatively evaluates the relationship between college major selection and subsequent career outcomes, including occupational sorting and entrepreneurship. Using data from the 1997 National Longitudinal Survey of Youth (NLSY97), I show that college graduates from Science, Technology, Engineering, and Mathematics (STEM) fields are systematically more likely to pursue STEM-related careers and less likely to become entrepreneurs relative to their non-STEM counterparts. To understand the mechanisms behind this phenomenon and its implications for the broader economy, I develop a quantitative model that links college major selection to post-graduation occupational outcomes and entrepreneurship decisions. The model, calibrated to match key moments from the NLSY97 micro-data, successfully replicates the observed patterns in major selection, occupational sorting, and entrepreneurship. Through counterfactual experiments, I find that reducing STEM tuition by 50 percent would nearly double the share of STEM majors but decrease overall entrepreneurship rates. Conversely, reducing barriers to entrepreneurship increases entrepreneurial activity without any corresponding impact on STEM enrolment or employment. These findings highlight the complex relationships between education choices and career paths and, as such, offer novel insights for policies related to education financing and entrepreneurship.

Collectively, the three papers in this thesis advance our understanding of key economic phenomena shaping modern economies and labour markets. They emphasize the critical role of human capital development—through both formal education and job-related training—in shaping individual and aggregate economic outcomes. Additionally, the research underscores the complexity of entrepreneurship and firm dynamics, highlighting how educational background, policy environments, and the option value of firm sales can significantly impact these aggregate processes. By integrating microeconomic data with macroeconomic theory, this thesis offers a comprehensive examination of pressing issues such as wage inequality, skill formation, business growth, and entrepreneurship. The findings provide policymakers with evidence-based insights to design more effective strategies that address specific challenges, such as reducing wage disparities, enhancing workforce skills, and supporting sustainable business development in a rapidly-evolving economy.

Chapter 1

Skill-Biased Technological Change, Training, and the College Wage Premium: A Quantitative Evaluation

Skill-Biased Technological Change, Training, and the College Wage Premium: A Quantitative Evaluation

Thomas Palmer
McMaster University
palmet8@mcmaster.ca

Abstract

This paper establishes that the rise in employer-provided training due to technological change has dampened the college wage premium. Using unique survey micro-data, I show that high-technology firms provide more training overall, but the gap in training participation between high- and low-skill workers is smaller within these firms. To understand the aggregate implications of these patterns, I build a quantitative model of the labour market with endogenous technology and training investments. In a counterfactual exercise, I find that the increase in the college wage premium would be 63 percent greater if training costs remained constant between 1980 and the early 2000s.

Keywords: *Training, Technological Change, College Wage Premium, Education, Technology.*

1.1 Introduction

The proliferation of new information-and-communications technologies (ICT) beginning in the 1980s coincided with a rise in both the relative supply and price of post-secondary educated labour. The starkest changes occurred in the United States, where the college wage premium more-than-doubled and the share of post-secondary educated workers increased by 91 percent over this period. Other OECD countries experienced a similar fate, although typically to a lesser degree. For example, in Canada, the college wage premium increased by 22 percent between 1980 and 2000[1], while the share of post-secondary educated (henceforth, high-skill) workers increased by 67 percent. In this paper, I document a new channel connecting technological change to the college wage premium: employer-provided training. In particular, I empirically document that new technologies often require additional training, which itself generates an earnings premium independent of education. I then quantitatively show that increased training participation among low-skill workers has produced a dampening effect on the college wage premium.

Existing studies of the college wage premium, including Bound and Johnson (1992), Katz and Murphy (1992), and Krusell et al. (2000), often highlight the effect of new technologies on the productivity of high-skill workers and, hence, how skill-biased technological change (SBTC) raises the college wage premium through the response of labour demand. The main empirical contribution of this paper is to document a related link between technology and training provision at the workplace-level. Specifically, I use unique matched employer-employee survey data from Canada to establish three facts. First, I show that technology-intensive (henceforth, high-technology) firms are relatively more productive, employ more high-skill workers, and provide more training than low-technology firms. That is, there exists a positive relationship between work-related training and technological intensity in the workplace. Second, I demonstrate that training generates a significant earnings premium, conditional on various worker and firm characteristics. Third, I show that the difference in training participation rates between high- and low-skill workers is smaller within high-technology firms. As technological change intensifies, the relative training participation rate among low-skill workers accelerates which, in turn, raises the average wage among this group and reduces the college wage premium. I provide further evidence of this mechanism by estimating a college wage premium separately for trained and untrained workers within high-technology firms and show that the premium is indeed smaller for the trained group.

Together, Facts 1 and 2 suggest that high-skill workers disproportionately benefit from technological change: not only does the return to their formal credentials increase, but they are also more likely to work for high-technology employers and reap the benefits of training over their working career. When combined with Fact 3, however, the opposite

[1]See, for example, Krueger et al. (2010). In related work, Boudarbat, Lemieux, and Riddell (2010) document an increase of 25 percent for the university-to-high-school graduate premium among males between 1980 and 2005 using Canadian Census data. On the other hand, Kryvtsov and Ueberfeldt (2009) find no change in the premium between 1980 and 2000 when comparing males with at least a bachelor's degree to males without any post-secondary education using Canadian Survey of Consumer Finances data.

story begins to emerge. Therefore, to better understand the aggregate implications of these facts, I build a quantitative model of the labour market, which embeds endogenous investments in education (Flinn and Mullins, 2015; Shephard and Sidibe, 2019) for workers, and in training (Moen and Rosén, 2004; Flinn, Gemici, and Laufer, 2017) and technologies (Shi, 2002) for firms, into a general equilibrium directed search model. Assuming that the benefits of training are fully match-specific, search frictions allow workers to extract some of the rents from training.[2] To ensure interior solutions for the shares of high-skill, trained, and high-technology employees as in the data, I allow the cost functions for technology investment and training provision to depend on the education level of a worker. With directed search, solving the model remains tractable: given a set of taxes to finance the economy's unemployment insurance program, the worker and firm problems can be solved independently of the distribution of individuals across states.[3] The model generates an exact decomposition of the college wage premium into three channels: (1) the relative supply high-skill workers, (2) the relative complementarity between high-skill labor and technology, and (3) the training participation gap between high- and low-skill workers.

I calibrate the model starting with the final steady state and work backwards in time. I pool repeated cross-sections from the Canadian micro-data between 1999 and 2005 to obtain a set of 11 moments to exactly identify 11 parameters in the final steady state. To identify and estimate the parameters governing the initial steady state, I impose an additional restriction. Specifically, I sort the set of 11 parameters into a group of 7 parameters, which are held fixed across the steady states, and a group of 4 parameters, which vary across the steady states. The first group of parameters governs the productivity of education and training, the cost of posting vacancies, and the setup cost associated with the high technology. The second group of parameters governs the complementarity between technology and high-skill labour, the costs of training, and the cost of higher education; their values are chosen to match the increase in the college wage premium between 1980 and 2000, the share of training participants in 1980, and the share high-skill workers in 1980, respectively.

I use the calibrated model as a laboratory to decompose the college wage premium and, ultimately, measure the effect of training on the observed rise in the college wage premium for Canada between 1980 and the early 2000s. Consistent with existing quantitative work, I find that technological change is the primary driver of the college wage

[2]Put differently, in a frictionless environment with fully match-specific training, workers would not receive any premium from training, which is inconsistent with the data.

[3]See, also, Menzio and Shi (2010, 2011).

premium and explains 60 percent of its absolute variation over this period. Increased training participation explains a meaningful 28 percent of the absolute variation and, in fact, dampens the premium. In particular, I show that, if training costs were held fixed between the 1980 and the early 2000s, the increase in the college wage premium would have been 63 percent greater.[4]

The remainder of the paper is organized as follows. In Section 1.2, I discuss the related literature on training, technological change, and the college wage premium. In Section 1.3, I describe the data in detail and present the motivating facts. In Section 1.4, I formally describe the model and define the equilibrium concept. In Section 1.5, I discuss the calibration. In Section 1.6, I present the results of the decomposition exercise and discuss its implications. Section 1.7 concludes the paper.

1.2 Related Literature

Starting with the seminal contributions of Bound and Johnson (1992), Katz and Murphy (1992), Autor, Katz, and Krueger (1998), and Krusell et al. (2000), an extensive literature has explained the rise in the college wage premium experienced by the U.S. between the 1960s and early 2000s as the result of skill-biased technological change (SBTC). The unifying theme of this literature is that the introduction of new production technologies—notably, computers—has disproportionately benefited high-skill workers because of the existence of capital-skill complementarities in production. As a result, a fall in the relative price of capital goods increases the relative demand for high-skill labour and exacerbates the college wage premium. Katz and Murphy (1992) formally demonstrate this mechanism by applying a competitive model of the labor market to data from the Current Population Survey for the period 1963-1987. Krusell et al. (2000) enrich the analysis by developing a model that links technological change to observables and use it to decompose the college wage premium over a longer time horizon. Ultimately, they find that the combination of cheaper capital goods and capital-skill complementarity accounts for approximately two-thirds of the growth in the college wage premium observed in the U.S. from 1963 to 1992.[5]

[4] In terms of its level, the college wage premium under fixed training costs would be 16 percent larger than its 1980 level or 6 percent larger than its 2000 level.

[5] Several papers have also applied similar frameworks to analyze trends in college wage premia in other countries, including Canada (Burbidge, Magee, and Robb, 2002; Boudarbat, Lemieux, and Riddell, 2010), the U.K. (Blundell, Green, and Jin, 2022), Indonesia (Amiti and Cameron, 2012), Japan (Lise et al., 2014; Takahashi and Yamada, 2022), Germany (Glitz and Wissmann, 2021), and many others.

Clearly, however, there are many factors beyond technological change that also affect the college wage premium.[6] For example, Walker and Zhu (2008), Velden and Bijlsma (2016), and Matsuda (2020) conduct a more thorough analysis of the college enrolment decision and document the importance of accounting for large changes in the relative supply of high-skill labor when analyzing changes in the college wage premium. Parro (2013), Dix-Carneiro and Kovak (2015), and Burstein and Vogel (2017) demonstrate that reductions in trade costs exacerbate the college wage premium at both an aggregate and a local level. Açıkgöz and Kaymak (2014) and Zentler-Munro (2021) study the role of bargaining power by drawing attention to the different rates of deunionization faced by high- and low-skill workers over time. Finally, He (2012) studies the impact of large-scale changes in the age structure of the economy—specifically, the baby boom and baby bust—on college enrollment and the college wage premium. While these alternative and complementary explanations certainly offer useful insights, the studies mentioned above generally still require a strong role for technology—and, in particular, capital-skill complementarities in production—to generate an empirically-consistent rise in the college wage premium. For this reason, I ultimately focus on technological change as the main driver of the college wage premium.

Within this literature, the most closely related papers are Lindner et al. (2022) and Doepke and Gaetani (2020). Lindner et al. (2022) study the impact of firm-level technological change on skill demand and aggregate inequality using a quantitative model of an imperfectly competitive labour market. In their framework, firms' wage policies internalize the fact that higher wages attract more workers. Therefore, in response to skill-biased technological change, firms increase the relative wage of high-skill workers, which generates a corresponding increase in both the firm's share of high-skill workers and the aggregate college wage premium. In my framework, skill-biased technological change similarly increases the share of high-skill workers and the college wage premium but for *two* reasons. First, SBTC directly increases the productivity of, and return to, being a high-skill worker. New firms respond to technological change by posting vacancies to attract high-skill workers, while more newborn individuals respond by enrolling in college at a greater rate. Second, the additional productivity gain from SBTC experienced by all firms encourages more of them to provide training. Workers in high-technology firms especially benefit, as they receive both a direct increase in productivity and a indirect increase in the probability of receiving training. Doepke and Gaetani (2020) study the differences in college wage premia between the U.S. and Germany through the lens of a

[6] See, also, Card and DiNardo (2002) for a more complete discussion of some prominent alternatives to the SBTC hypothesis.

competitive model of the labor market in which firms and workers make match-specific investments in skill accumulation. Because the incentive to invest in skills is increasing in the expected duration of the match, they argue that stricter employment protection laws in Germany have particularly benefited low-skill workers and have helped to moderate the German college wage premium over time. Contrary to Doepke and Gaetani (2020), I do not consider differences in employment protection. Instead, I show that technology itself generates a greater incentive to provide work-related training. The effect is most pronounced for low-skill workers, thus attenuating the college wage premium over time.

There is also a large empirical literature devoted to estimating the causal effect of work-related training on earnings. Studies in this literature generally find that training generates large and persistent returns for participants. For example, Blundell, Dearden, and Meghir (1996) apply a quasi-difference specification to a subset of individuals from the British National Child Development Survey and estimate a 3.6 percent wage increase to male participants of employer-provided on-the-job training courses and a 6.6 percent wage increase for employer-provided off-the-job training courses. Parent (1999) estimates a series of OLS and IV regressions using data from the National Longitudinal Survey of Youth (NLSY) and finds a wage effect ranging from 12 to 17 percent for on-the-job training and 7.5 to 14 percent for off-the-job training provided by an individual's current employer. Leuven and Oosterbeek (2008) consider a novel identification strategy by comparing training participants to non-participants who initially wanted to participate but were unable to for exogenous reasons. Under this environment, the authors estimate a near-zero return to training participation, suggesting that there may be some selection in terms of who does and does not receive training. In my empirical analysis of Canadian micro-data, I show that training participation is indeed skewed toward specific worker and firm types: those who are high-skill and employed at high-technology firms.

Finally, a few papers have attempted to link technology adoption, training, and labour market outcomes in their empirical work. Bartel and Sicherman (1998) evaluate the effects of technological change on individual training participation by combining the NLSY with cross-sectional industry-level measures of computer investment, total factor productivity, and R&D intensity. Ultimately, they find that workers are more likely to receive training as technological progress intensifies within the industry, and that training participation rates are increasing in education. Relative to Bartel and Sicherman (1998), the data I use allows me to identify technological change at a finer level, namely, the firm. As a result, I show that the positive association between technology and training carries over to the individual match-level and extends more generally to industries beyond manufacturing. In more recent work, Bresnahan, Brynjolfsson, and Hitt (2002), Black and Lynch (2004), and Boothby, Dufour, and Tang (2010) use firm-level survey data to study how firms' decisions to adopt new technologies affects productivity and innovation through its interaction with workplace organization and training provision. In contrast to this set of papers, I link information on technology use at the workplace level to individual employee characteristics. Accordingly, I not only provide further evidence that high-technology firms offer more training in absolute terms, but also show that the increased likelihood of receiving training in such firms is relatively greater for low-skill workers.

1.3 Data

In this section, I document new facts related to technology, training, and the college wage premium based on survey micro-data from Canada. Together, the facts demonstrate the key mechanism underlying the quantitative model of Section 1.4.

1.3.1 Workplace and Employee Survey (WES)

The empirical analysis is based on data from the Workplace and Employee Survey (WES). The WES is a matched employer-employee survey data set from Canada, which covers approximately 20,000 employees spread across 6,000 workplaces at an annual frequency from 1999 to 2006.[7] I focus, in particular, on the cross-sectional workplace and employee samples from 1999, 2001, 2003, and 2005.[8] I also restrict attention to workers aged 25 to 64 years old to limit variation in hours, employment status, and earnings arising from full- or part-time enrolment in education and retirement.

The WES data provides two main benefits for studying the link between technological change and the college wage premium. First, the data contain rich survey information on technological intensity and adoption at the workplace-level. Existing studies of technology adoption and wage inequality, such as Autor, Katz, and Krueger (1998) and Kristal (2020), are typically restricted to industry- or occupation-level analyses. Instead, I exploit the workplace-level variation in the WES to examine the effects of technology at a more granular level. Second, the data contain information on training participation and provision, thereby overcoming a major challenge in the existing empirical literature

[7] A "workplace" in this context means "establishment."

[8] Because sampled workers are only followed for two years at most, the survey-weighted statistics computed on the even-year samples of employees reflect the population statistics for the preceding (odd) survey year (minus attrition). Further details about the data and sampling design are contained in Appendix A1.

on human capital accumulation. Importantly, the ability to link employees to their employers is crucial to identify the responsiveness of training to technology, which is a key mechanism driving changes in the college wage premium in the model.

1.3.2 Definitions of Training, Education, and Technological Intensity

1.3.2.1 Training Participation and Provision

The WES collects information on two main types of employer-provided training: formal classroom and informal on-the-job training. For each mode of training, both employees and employers are asked about training duration, intensity[9], subject matter, and funding sources. Since both training types are identified as *employer-provided*, I define training participants as workers who have received *either* classroom *or* on-the-job training over the past survey year and training firms as workplaces that have provided *either* classroom *or* on-the-job training to at least one employee over the past survey year.

Under this classification, approximately 50 to 60 percent of workers are identified as training participants and 40 to 50 percent of workplaces are identified as training providers each year, on average.

1.3.2.2 High-Skill and Low-Skill Workers

In every odd survey year, employees are asked the following series of questions:

1. *Did you graduate high-school?*
2. *Have you received any other education?*
3. *What was that education?*

I define *high-skill* workers as employees that have graduated high school (answered "yes" to question 1) and have received additional education (answered "yes" to question 2). The set of high-skill employees thus includes college and university graduates, as well as employees with post-secondary diplomas; trade, vocational, or industry certificates; and, post-graduate or professional degrees. On the other hand, I define *low-skill* workers as employees that have either not graduated high-school (answered "no" to question 1) or have graduated high-school without any additional education (answered "yes" to question 1 and "no" to question 2).

[9] For example, a measure of how many courses were taken.

On average, high-skill workers account for approximately two-thirds of the labour force over the sample period. Roughly 25 to 30 percent of the overall sample are university-educated (Bachelor's, Master's, Ph.D., M.D., etc.); 30 to 35 percent are college-educated (college degrees, trade or vocational school, industry certified, etc.); and the remainder are high-school graduates and dropouts.

1.3.2.3 High-Tech and Low-Tech Firms

In each survey year, the WES questionnaire asks employers the following question:

> *"At this location, how many employees currently use computers as part of their normal working duties? By computers, we mean a microcomputer, personal computer, minicomputer, mainframe computer or laptop that can be programmed to perform a variety of operations."*

Using the reported answers to this question and the workplaces' total number of employees, I construct a variable that identifies the *share* of a workplace's employees who use computers regularly on the job. Formally, for each workplace j in year t, I compute:

$$\text{ShareCPU}_{j,t} = \frac{\text{Number of Computer Users}_{j,t}}{\text{Total Number of Employees}_{j,t}}. \tag{1.1}$$

I use the estimated shares to classify employers as high- or low-technology. I define *high-technology* firms as workplaces with at least 50 percent of employees using computers as part of their normal working duties: ShareCPU ≥ 0.50. Low-technology firms comprise the remaining firms, that is, workplaces with strictly less than 50 percent of employees using computers: ShareCPU < 0.50.[10]

1.3.3 Stylized Facts

Having defined the main variables of interest, I now establish empirically how technological change affects the college wage premium through the response of training. I report these findings as a series of three stylized facts.

Fact 1: High-technology firms are more productive, more likely to provide training, and more likely to employ high-skill workers than low-technology firms.

[10] Autor, Katz, and Krueger (1998) use a similar strategy to identify high-technology industries. However, rather than using a continuous measure of technological intensity, I adopt a binary classification for consistency with my quantitative model.

To understand the distinguishing characteristics of high- and low-technology firms, I evaluate their differences along three dimensions: (1) productivity, (2) training, and (3) employment structure.

Table 1.1 presents the results from estimating a workplace-level regression of (log) revenue productivity on a high-technology indicator; a vector of time-varying workplace-level control variables, including industry and firm size; and year fixed-effects.[11] High-technology firms are, on average, approximately 58 percent more productive than technology firms. This finding is consistent with existing models of technology investment and firm dynamics, such as Shi (2002), in which a firm invests in the high-technology whenever it meets or exceeds a threshold level of idiosyncratic productivity.

TABLE 1.1: Relative Productivity of High-Tech Firms

Dependent Variable: Log[Revenue Productivity]	
High-Tech	0.4538***
	(0.0306)
N (Unweighted)	22,392
N (Weighted)	2,627,197

Note: Standard errors in parentheses. Standard errors are bootstrapped using the workplace bootstrap weights provided by Statistics Canada and 100 replications. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$

To evaluate differences in training incidence between high- and low-technology firms, I estimate an employee-level discrete choice regression of training participation on a high-technology indicator. The top panel of Table 1.2 summarizes the estimated odds ratios, while the bottom panel reports the average marginal effects. Column (1) reports the results of a regression controlling only for the worker's skill level, while Column (2) includes a richer set of time-varying individual-level control variables (still including skill level), a set of time-varying workplace-level control variables, and year fixed-effects. The results here imply that high-technology employees are, on average, 10.4 percent more likely to receive training, conditional on education, and 3.4 percent more likely to participate in training, after controlling for a host of additional observables.

[11]The specifications for all regressions underlying Fact 1 are stated formally in Appendix A2.

Table 1.2: Probability of Training Participation

Dependent Variable: Training Indicator	**(1)**	**(2)**
Odds Ratios		
High-tech	0.4311***	0.1573**
	(0.0330)	(0.0617)
High-skill	0.6847***	0.2895***
	(0.0335)	(0.0385)
Average Marginal Effects		
High-tech	0.1036***	0.0337**
	(0.0184)	(0.0133)
High-skill	0.1673***	0.0626***
	(0.0184)	(0.0084)
Additional Controls?	–	✓
N (Unweighted)	81,622	65,814
N (Weighted)	41,405,006	31,047,424

Note: Standard errors in parentheses. Standard errors are bootstrapped using the workplace bootstrap weights provided by Statistics Canada and 100 replications. $^{***}p < 0.01, ^{**} p < 0.05, ^{*} p < 0.10$

Finally, I assess whether the definitions of high- and low-technology firms I use are consistent with the existing literature on technology-skill complementarities by estimating an employee-level logistic regression of the high-technology indicator on an indicator for the worker's skill level. As before, the top panel of Table **??** reports the estimated odds ratios and the bottom panel reports the average marginal effects. Column (1) controls only for training participation, while Column (2) controls for additional workplace- and worker-level variables including occupation, age, experience, gender, CBA coverage, immigration status, industry, firm size, and year. The estimated marginal effects imply a 19 percent greater likelihood for high-skill workers to be employed in a high-technology firm, conditional on receipt of training, and a 2.3 percent greater likelihood, conditional on a series of additional covariates.

Fact 2: Training participation generates a significant premium on hourly earnings.

While Fact 1 establishes a positive relationship between technology and training, it remains to show that training by itself has an effect on the college wage premium separate from technology. To address this issue, I perform a series of Mincer (1958) style regressions to quantify the impact of employer-provided training on hourly earnings at the worker-level, both conditionally and unconditionally. For a worker i employed by firm j in year t, the baseline specification is given by:

$$\begin{aligned} \ln(\text{Earnings}_{i,t}) = \beta_0 + \beta_1 \text{Train}_{i,t} + \beta_2 \text{HighSkill}_{i,t} + \beta_3 \text{HighTech}_{j,t} + \\ \beta_4(\text{HighSkill}_{i,t} \times \text{HighTech}_{j,t}) + \theta_t + \delta X_{i,t} + \xi Z_{j,t} + \varepsilon_{i,t} \end{aligned} \tag{1.2}$$

where hourly earnings are expressed in constant 1999 dollars; the vector $X_{i,t}$ of individual-level control variables includes occupation, usual weekly hours, tenure, experience, age, an indicator for non-employer-sponsored career-related training, CBA coverage, gender, immigrant status, and indicators for whether the worker uses computers or other types of information technologies in the workplace; and the vector $Z_{j,t}$ of time-varying firm-level controls includes firm size, productivity, and industry.

The coefficient of interest is β_1 and identifies the impact of training participation on (log) hourly earnings. Importantly, I also control for the effects of workers' skills (captured by β_2), the technology of worker i's employer (captured by β_3), and an interaction between worker skills and employer technology (captured by β_4) to control for potential earnings effects from technology-skill complementarities. Table 1.3 reports the results from performing this regression after accounting for the complex survey design of the WES.

Column (1) on the left reports the results from the regression omitting all control variables except for the ones specified—that is, without θ_t, $X_{i,t}$, and $Z_{j,t}$. Column (2) on the right reports the results from the regression with additional controls. In both cases, training participation is shown to have an economically and statistically significant impact on individual earnings. For the baseline case—conditioning only on education and employer type—training participation is associated with a roughly 15 percent increase in hourly earnings, while in the preferred specification with additional controls, training participation is associated with a 3 percent increase in hourly earnings. The positive wage effect of training is robust to various alternative sets of control variables and clustering standard errors by workplace.

Fact 3: Technology-induced training dampens the college wage premium.

TABLE 1.3: Impact of Training Participation on Earnings

Dependent Variable: Log[Hourly Earnings]	**(1)**	**(2)**
Trained	0.1407***	0.0292***
	(0.0119)	(0.0079)
High-skill	0.1804***	0.0641***
	(0.0110)	(0.0090)
High-tech	0.1310***	0.0120
	(0.0171)	(0.0125)
High-skill × High-tech	0.063***	0.0737***
	(0.0203)	(0.0166)
Additional Controls?	–	✓
N (Unweighted)	81,622	65,532
N (Weighted)	41,405,006	30,899,192

Note: Standard errors in parentheses. Standard errors are bootstrapped using the employee bootstrap weights provided by Statistics Canada and 100 replications. $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$

TABLE 1.4: Training Participation Rates by Education and Technology

	Training Participation Rates (%)		
Skill-Group	**Low-Tech**	**High-Tech**	**Change**
Low-Skill	37.48	50.62	35.04
High-Skill	55.95	65.07	16.30

Note: The column "Change" reports the difference in training participation rates (in percent) between low- and high-technology firms, conditional on skill group.

The main limitation of the WES micro-data is the short time horizon that it covers: it is simply not possible to obtain a direct measure of how training has evolved for high- and low-skill workers over time. However, it *is* possible to recover an indirect estimate of the evolution by exploiting the properties of technological change. In particular, if a relatively larger share of firms in 1980 were low-technology, then one way to indirectly assess the evolution of training participation by skill level is to measure how much the rate of training participation increases for each group when comparing low-technology to high-technology employees in the cross-section. Performing this exercise yields the results reported in Table 1.4.

Columns 2 and 3 list the share of training participants within low- and high-technology firms, respectively, conditional on a worker's skill level. Column 4 reports the relative

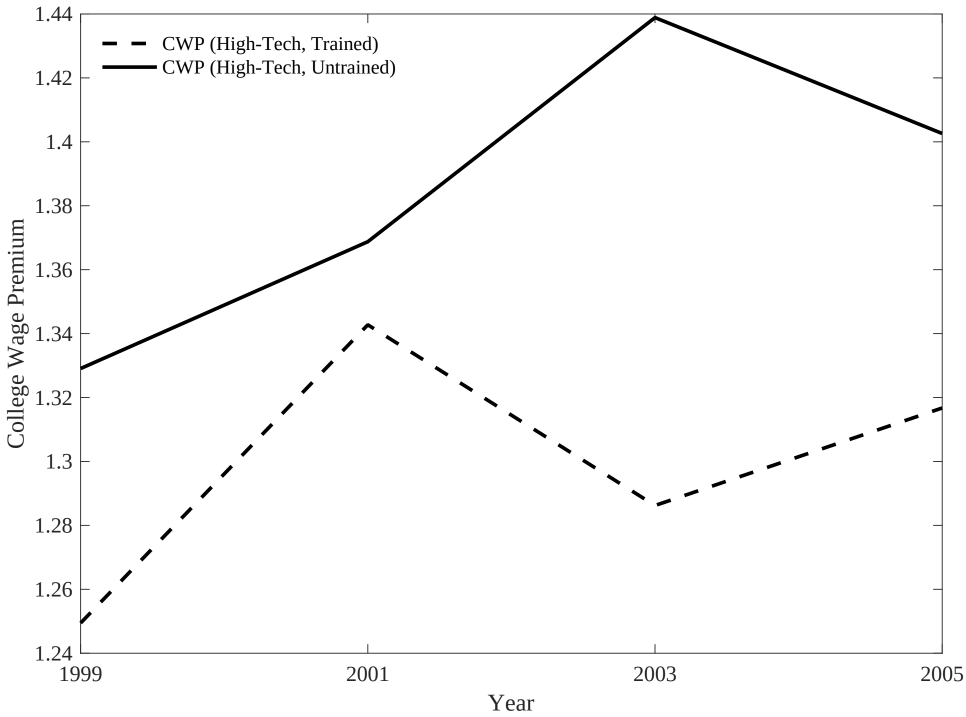

FIGURE 1.1: College Wage Premium by Technology and Training

increase in training participation for a worker of a given skill level when moving from a low- to a high-technology firm. As illustrated, the increased likelihood of receiving training in high-technology firms is much more pronounced for low-skill workers: low-skill workers experience a 35 percent increase in the probability of receiving training, while high-skill workers experience a 16 percent increase. Put differently, the relative training participation rate for low-skill workers is 67 percent in low-technology firms and 78 percent in high-technology firms. A broader implication of this result is that, conditional on being employed in a high-technology firm, the college wage premium is smaller among the subset of trained workers. This is precisely the result shown in Figure 1.1.

1.4 Model

To better understand the aggregate implications of the stylized facts presented in Section 1.3, I now develop a quantitative model of the labour market, which connects human capital investments in education and training to an aggregate process of skill-biased technological change in the spirit of Shi (2002) and Salgado (2020). The model features two levels of education and two levels of production technologies, as in the data. The shares of worker and firm types are determined endogenously in equilibrium. Workers decide their level of education prior to entering the labor market by solving a trade-off between the financial and non-pecuniary cost of enrollment and the expected return, which internalizes the probability of finding employment and receiving training. On the other side of the market, firms choose their production technologies at the time of posting vacancies by solving a trade-off between the higher setup cost associated with the high-technology and the expected return, which internalizes the firm's optimal training and wage policies conditional on matching.

1.4.1 Environment

Time is discrete and continues forever, $t = 0, 1, 2, \ldots$. Three types of agents populate the economy: workers, firms, and a government. The mass of workers is normalized to 1 and the mass of firms is endogenously determined through free entry into the labour market. Workers are ex-ante heterogeneous in innate ability a, risk neutral, and discount future income with factor $\beta \in (0, 1)$. I assume that innate ability a takes on one of m_a possible values, that is, $a \in \{a_1, a_2, \ldots, a_{m_a}\}$. Firms are ex-ante heterogeneous in productivity z, risk neutral, and share the same discount factor β. I assume that firm productivity z takes on one of m_z possible values, that is, $z \in \{z_1, \ldots, z_{m_z}\}$. The government sets lump-sum taxes τ on labor income to finance the economy's unemployment insurance program.

As workers choose their level of education $s \in \{0, 1\}$ prior to entering the labour market, it is convenient to envision each period unfolding in two stages: an education stage and a labor market stage. In the education stage, individuals decide—conditional on their innate ability a—whether to invest in post-secondary education by paying the cost of enrollment $c_s - \zeta_s > 0$ and enter the labor market as a type-$(a, 1)$ high-skill worker, or to forgo additional education and enter the labor market as a type-$(a, 0)$ low-skill worker. Once a worker enters the labor market, their education level s is fixed for the remainder of their life. In the labor market stage, workers simply shift between employment and unemployment depending on the shocks they experience over their working life.

The labour market is composed of a continuum of submarkets indexed by the set (a, s, z, k, x, ω), consisting of worker innate ability a, worker education $s \in \{0, 1\}$, firm productivity z, firm technology $k \in \{0,1\}$[12], firm training provision $x \in \{0,1\}$, and the piece-rate $\omega \in [0,1]$ posted by the firm. The piece-rate represents the fraction of output promised to the worker each period. For convenience, let $\mathbf{s}_w = (a,s)$ represent the vector of worker types and $\mathbf{s}_f = (z,k,x)$ represent the vector of firm types. There is no uncertainty. Search is directed in the sense that an unemployed worker of type $\mathbf{s}_w$ observes the firm types $\mathbf{s}_f$ and piece-rates ω posted within each submarket prior to making their application decision. The assumptions regarding the information environment imply that, when searching, an unemployed worker knows whether the firm has invested in the high-technology and whether they will receive training when computing their expected value from employment; and there is no incentive for either party to renegotiate contracts ex-post once a match is formed.

Matches within each sub-market are governed by a constant-returns-to-scale matching function $m(u,v)$, where u represents the mass of unemployed searchers and v represents the associated mass of vacancies in a given submarket. Let $\theta \equiv v/u$ denote the ratio of vacancies to unemployment, or tightness, of a submarket. Then, for a submarket with tightness $\theta(\mathbf{s}_w,\mathbf{s}_f,\omega)$, the probability that a worker successfully finds a job is given by:

$$p\left(\theta(\mathbf{s}_w,\mathbf{s}_f,\omega)\right) = \frac{m\left(u(\mathbf{s}_w,\mathbf{s}_f,\omega), v(\mathbf{s}_w,\mathbf{s}_f,\omega)\right)}{u(\mathbf{s}_w,\mathbf{s}_f,\omega)} \in (0,1). \tag{1.3}$$

On the other hand, the probability that a vacant job successfully finds a worker is given by:

$$q\left(\theta(\mathbf{s}_w,\mathbf{s}_f,\omega)\right) = \frac{m\left(u(\mathbf{s}_w,\mathbf{s}_f,\omega), v(\mathbf{s}_w,\mathbf{s}_f,\omega)\right)}{v(\mathbf{s}_w,\mathbf{s}_f,\omega)} \in (0,1). \tag{1.4}$$

Matches only end for exogenous reasons. With probability $\eta_s \in (0,1)$, an existing match is exogenously separated each period; and, with probability $\delta \in (0,1)$, an existing worker exogenously leaves the economy each period. When a match receives a separation shock η_s, the job is destroyed and the worker transitions to unemployment. When a worker receives an exit shock δ, the job is destroyed and the worker is replaced by a newborn worker of the same type, who enters the education stage next period.[13]

Timing within a period occurs as follows. First, a mass δ of newborn workers make their education decisions. In the following period, these workers will enter the labour market as a type-$\mathbf{s}_w$ unemployed worker. Second, idiosyncratic shocks are realized. Existing matches are exogenously separated with probability η_s and existing labor market

[12]Here, $k = 1$ represents the high-technology while $k = 0$ represents the low-technology.

[13]Hence, with a mass of workers normalized to 1, δ represents both the probability of exiting and the mass of entering workers each period.

participants exogenously exit the market with probability δ. Third, potential entrant firms pay cost κ to draw productivity z and, conditional on the realization of productivity, decide whether to invest in the high-technology. Given its productivity and technology, the firm chooses a contract consisting of a piece-rate ω and training policy x to post in a submarket to attract type-$\mathbf{s}_w$ unemployed workers. On the other side of the market, unemployed type-$\mathbf{s}_w$ workers direct their search to the submarket that maximizes their lifetime value. Fourth, matching occurs. Unfilled vacancies and unemployed job seekers are matched in each submarket according to the matching function $m(u, v)$ described above. Fifth, production occurs. Firms execute their contractual obligations by paying the training cost $c_x(s)$ and produce output according to the production function $y(\mathbf{s}_w, \mathbf{s}_f)$. Finally, payments are made. Firms earn operating profits $(1-\omega)y(\mathbf{s}_w, \mathbf{s}_f)$, employed workers earn net income $\omega y(\mathbf{s}_w, \mathbf{s}_f) - \tau$, and unemployed workers earn unemployment insurance b_s.

1.4.2 Firms

1.4.2.1 Production

There is a large mass λ of firms each period, which is determined by free entry. Each firm employs at most one worker. When matched with a worker of type-$\mathbf{s}_w$, a firm of type-$\mathbf{s}_f$ produces output $y(\mathbf{s}_w, \mathbf{s}_f)$ according to the following production technology:

$$y(\mathbf{s}_w, \mathbf{s}_f) = \varepsilon_x(s)^x \varepsilon_k(s)^k zah_s, \tag{1.5}$$

where $\varepsilon_x(s)$ and $\varepsilon_k(s)$ capture, in a reduced-form way, the productivity gains earned through training and technological investments by the firm; and, h_s represents the productivity gain earned through investments in education by the worker. I assume that $\varepsilon_x(s) > 1$ and $\varepsilon_k(s) > 1$ for all s, $h_0 = 1$, and $h_1 > 1$.

1.4.2.2 Firm Value Function

Consider a match of type $(\mathbf{s}_w, \mathbf{s}_f, \omega)$ and let $\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega)$ denote the continuation profit for the firm—that is, the profit accrued to the firm in every period after paying the training cost in period 1. Then, $\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega)$ solves the following recursive equation:

$$\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega) = (1-\omega)y(\mathbf{s}_w, \mathbf{s}_f) + \beta(1-\delta)(1-\eta_s)\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega). \tag{1.6}$$

Let $J_F(\mathbf{s}_w, \mathbf{s}_f, \omega)$ denote the lifetime value for the same firm. The only difference between $\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega)$ and $J_F(\mathbf{s}_w, \mathbf{s}_f, \omega)$ is that the latter includes a one-time training cost, conditional on providing training. That is, $J_F(\mathbf{s}_w, \mathbf{s}_f, \omega)$ solves:

$$J_F(\mathbf{s}_w, \mathbf{s}_f, \omega) = (1-\omega)y(\mathbf{s}_w, \mathbf{s}_f) - c_x(s)x + \beta(1-\delta)(1-\eta_s)\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega), \tag{1.7}$$

where $c_x(s) > 0$ is the lump-sum cost of training. In the current period, a firm earns revenue from output $y(\mathbf{s}_w, \mathbf{s}_f)$ net of the wage bill $\omega y(\mathbf{s}_w, \mathbf{s}_f)$ and training cost $c_x(s)$, conditional on training ($x = 1$). In the following period, the match survives with probability $(1-\delta)(1-\eta_s)$, in which case the firm earns continuation value $\Pi(\mathbf{s}_w, \mathbf{s}_f, \omega)$. Hence, when a firm commits to providing training, $x = 1$, it must pay for training *once* even though the return to training—captured in $\varepsilon_x(s)$—accrues to the match for as long as it survives.

1.4.2.3 Free Entry and the Zero Profit Condition

Each period, firms enter sub-markets until the value of a vacancy in each submarket is driven to zero. By paying cost κ, a potential entrant draws a level of productivity z and decides whether to invest in the high-technology, $k \in \{0, 1\}$. Conditional on the choice of technology k and its productivity z, the firm must also decide the terms of the contract (x, ω) to post along with its vacancy. Hence, the zero profit condition for submarkets is given by:

$$\kappa + c_k(s)k = q(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))J_F(\mathbf{s}_w, \mathbf{s}_f, \omega), \tag{1.8}$$

where $q(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))$ is the job-filling probability in sub-market $(\mathbf{s}_w, \mathbf{s}_f, \omega)$. Therefore, firms who invest in the high-technology pay cost $\kappa + c_k(s)$ to post a vacancy, while firms who maintain the low-technology only pay κ. The solution to this problem yields tightness $\theta(\mathbf{s}_w, \mathbf{s}_f, \omega)$ for each submarket and job finding probabilities $p(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))$, which are required to solve the worker's problem.

1.4.3 Workers

The mass of workers is normalized to 1. Over their lifetime, workers occupy three possible states: employed at a firm of type-$\mathbf{s}_f$ and receiving piece-rate ω; unemployed and earning unemployment insurance b_s; or out of the labour force. Note that workers who are out of the labor force only consist of newborn workers who have yet to make their education decision—there is no explicit labor force participation decision.

1.4.3.1 Employed Value Function

Consider first an employed worker of type-$\mathbf{s}_w$, who is working at a firm of type-$\mathbf{s}_f$ that provides a piece-rate of ω each period. Let $V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$ denote the lifetime value for this worker. Each period, the worker simply consumes their earnings, while facing exogenous probabilities $\eta_s \in (0, 1)$ and $\delta \in (0, 1)$ of getting hit by a job separation or exit shock, respectively. In the case of receiving a separation shock η_s, the worker transitions to unemployment and receives continuation value $V_U(\mathbf{s}_w)$; in the case of receiving an exit shock, the worker receives a value of 0 and is replaced by a newborn worker of the same ability a in the following period. Conditional on surviving both shocks, however, the worker retains the value $V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$ from employment into the next period. Accordingly, the value function $V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$ solves the following recursive equation:

$$V_E(\mathbf{s}_w, \mathbf{s}_f, \omega) = \omega y(\mathbf{s}_w, \mathbf{s}_f) - \tau + \beta(1 - \delta)\Big[\eta_s V_U(\mathbf{s}_w) + (1 - \eta_s)V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)\Big], \qquad (1.9)$$

where $\tau \geq 0$ is a lump-sum tax on labour earnings collected by the government.

1.4.3.2 Unemployed Value Function

Each period, an unemployed worker of type-$\mathbf{s}_w$ decides which sub-market $(\mathbf{s}_w, \mathbf{s}_f, \omega)$ to search in. With probability $p(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))$, the worker successfully finds a job offering piece-rate ω and receives value $V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$ in the next period conditional on survival; otherwise, the worker remains unemployed and receives value $V_U(\mathbf{s}_w)$. Hence, $V_U(\mathbf{s}_w)$ solves:

$$V_U(\mathbf{s}_w) = \max_{(\mathbf{s}_f, \omega)} \; b_s + \beta(1 - \delta)\Big[p(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))V_E(\mathbf{s}_w, \mathbf{s}_f, \omega) + (1 - p(\theta(\mathbf{s}_w, \mathbf{s}_f, \omega))V_U(\mathbf{s}_w)\Big]. \qquad (1.10)$$

1.4.3.3 Education Decision

Newborn workers with innate ability a make their education decisions at the start of each period. This decision involves paying a cost $c_s - \zeta_s > 0$ to enrol in education level $s = 1$ and earn human capital h_s, where I assume that $h_0 = 1$ and $h_1 > 1$. To understand this formulation of the cost function, it is convenient to think of c_s capturing the financial costs of enrollment and ζ_s representing unmodeled preferences for higher education, which may include various non-monetary costs associated with enrollment.

I assume that the stochastic component of the cost function, ζ_s, follows a Type-I Extreme Value distribution with zero location parameter and scale parameter $\chi_s > 0$. Following this decision, a worker enters the labour market as a type-$\mathbf{s}_w$ unemployed worker in the following period. Accordingly, the optimal education policy for a worker with innate ability a, before the realization of the shock ζ_s, solves:

$$s(a) = \underset{s \in \{0,1\}}{\arg\max} \ \mathbb{E}_{\zeta_s} \Big\{ s \left[\beta(1-\delta) V_U(a,1) - c_s + \zeta_s \right] + (1-s)\beta(1-\delta) V_U(a,0) \Big\} \tag{1.11}$$

The introduction of Type-I Extreme Value shocks into the education decision implies that a worker's education policy is a probability.

1.4.4 Government

The government levies lump-sum taxes on labour income τ to finance unemployment insurance b_s. I assume that the government balances its budget in every period. Each period, total unemployment insurance is equal to

$$\sum_{\mathbf{s}_w} b_s u(\mathbf{s}_w, \mathbf{s}_f, \omega), \tag{1.12}$$

where $u(\mathbf{s}_w, \mathbf{s}_f, \omega)$ is the mass of type-$\mathbf{s}_w$ unemployed job searchers in sub-market $(\mathbf{s}_w, \mathbf{s}_f, \omega)$. Hence, the government's budget constraint is given by:

$$\tau \sum_{\mathbf{s}_w, \mathbf{s}_f, \omega} [1 - u(\mathbf{s}_w, \mathbf{s}_f, \omega)] = \sum_{\mathbf{s}_w} b_s u(\mathbf{s}_w, \mathbf{s}_f, \omega). \tag{1.13}$$

1.4.5 Equilibrium

A stationary equilibrium for this economy consists of a set of value functions $\{V_E, V_U\}$ for workers; value functions $\{J_F, \Pi\}$ for firms; search and education policy functions $\{(\mathbf{s}_f, \omega), s\}$ for workers; a distribution of matches $\Omega_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$; a distribution of unemployed workers over ability and education $\Omega_U(\mathbf{s}_w)$; a distribution of non-labour force participants $\Omega_D(a)$; taxes $\{\tau\}$; and tightness $\theta(\mathbf{s}_w, \mathbf{s}_f, \omega)$ for each submarket such that:

1. *Worker Optimization:* The value functions $V_E(\mathbf{s}_w, \mathbf{s}_f, \omega)$ and $V_U(\mathbf{s}_w)$ solve (1.9), (1.10), and (1.11), with associated policy functions $\{s(a),\ \omega(a, s(a)),\ x(a, s(a)),\ z(a, s(a)),\ k(a, s(a))\}$.

2. *Firm Optimization:* The value function $J_F(\mathbf{s}_w, \mathbf{s}_f, w)$ solves (1.7).

3. *Free Entry:* Firms enter sub-markets until the zero-profit condition (1.8) holds.

4. *Government Budget Balance:* The government's budget constraint (1.13) holds with equality in every period.

5. The distributions of workers across employment states ($\Omega_E(\mathbf{s}_w,\ \mathbf{s}_f, \omega)$, $\Omega_U(\mathbf{s}_w)$, $\Omega_D(a)$) are stationary.

1.5 Calibration

The model solution consists of an initial and final steady state. To calibrate the steady states, I drawn on some values from the literature, calculate others directly from data, and choose the remainder to match salient features of the Canadian economy in 1980 and the early 2000s.

1.5.1 Functional Forms

The only functional form required to specify prior to calibration is the matching function. For this, I follow Den Haan, Ramey, and Watson (2000) and use a constant returns to scale matching function of the following form:

$$m(u, v) = \frac{u \cdot v}{(u^{\xi} + v^{\xi})^{1/\xi}}, \tag{1.14}$$

where ξ is the matching elasticity, u is the mass of unemployed workers, and v is the associated mass of vacancies in a given sub-market.

1.5.2 Externally-Calibrated Parameters

Table 1.5 below lists the parameters determined outside of the model. I assume that all externally-calibrated parameters remain constant between the two steady states.

A period in the model corresponds to one-quarter in the data. As a result, I set the discount factor β to 0.987, implying an aggregate interest rate of 5 percent annually in steady state. The exogenous rate of job separation, η_s, varies by skill group: for high-skill workers, I use a value of $\eta_H = 0.0238$ and, for low-skill workers, I use a value of $\eta_L = 0.0501$. The labour market exit rate, δ, is set to 0.0063, so that the average working life is 40 years. The matching elasticity is set to 1.60, as in Schaal (2017). The shape parameter for the Extreme Value Type-I distribution governing the taste shocks is set to 0.40 to ensure an interior solution. Finally, as a normalization, I set $h_0 = 1$.

TABLE 1.5: Parameters determined outside of the model

	Parameter	Value	Source
β	Discount factor	0.987	Risk-free rate of 5%
η_{s_0}	Separation rate (Low-skill)	0.050	Flinn and Mullins (2015)
η_{s_1}	Separation rate (High-skill)	0.024	Flinn and Mullins (2015)
ξ	Matching elasticity	1.600	Schaal (2017)
δ	Exit rate	$\frac{1}{160}$	Average working life of 40 years
h_{s_0}	Productivity from base education	1.000	Normalization

1.5.3 Internally-Calibrated Parameters

The remaining parameters of the model are calibrated in two stages. In the first stage, I calibrate the entire set of parameters to empirical averages from the WES micro-data. This exercise yields a solution to the final steady state of the model. In the second stage, I fix a set of 7 parameters, $\{\kappa, h_{s_1}, \varepsilon_x(s_1), \varepsilon_x(s_0), \varepsilon_k(s_0), c_k(s_0), c_k(s_1))\}$, to their final steady state values and vary the remaining 4 parameters, $\{c_s, c_x(s_0), c_x(s_1), \varepsilon_k(s_1)\}$, such that the share of high-skill workers and training participants by education in the initial steady state match their 1980 empirical averages, and the model-implied increase in the college wage premium is consistent with the actual increase between 1980 and 2000.[14] Table 1.6 reports the results of the calibration exercise.

As indicated in the third panel of Table 1.6, I set the vacancy posting cost κ to match an average unemployment rate of 7 percent, which was obtained using publicly-available aggregate information on the unemployment rate from the OECD. I allow unemployment insurance, b_s, to vary across education groups and set b_0 and b_1 such that unemployment insurance covers 40% of the average earnings for low- and high-skill workers in the final steady state, respectively. The productivity gain earned from education, captured by h_{s_1}, is set to match a college wage premium of 1.37. The productivity gain earned from training, captured by $\varepsilon_x(s_1)$ and $\varepsilon_x(s_0)$, are set to match the mean earnings premium accrued by low- and high-skill training participants, respectively, which in the data amount to 1.14 and 1.16. The productivity gain from high-tech employment for low-skill

[14]The share of high-skill workers in 1980 comes from publicly-accessible OECD data. The share of training participants by education in 1980 comes from Statistics Canada and Human Resources Development Canada (2001) using the 1983-1985 Adult Education and Training Survey (AETS). Finally, the average increase in the college wage premium is taken to be an intermediate value across existing estimates.

workers, captured by $\varepsilon_k(s_0)$, is set to match the average wage premium of 1.19 for low-skill workers employed by high-tech firms. Finally, the setup costs associated with the high-technology varies by education to ensure that both types of workers are hired by high-technology firms, as in the data. The parameters $c_k(s_0)$ and $c_k(s_1)$ are set to match the average share of low- and high-skill employees at high-technology firms, respectively.

The parameters reported in the top and middle panels of Table 1.6 vary across steady states. These parameters govern the cost of post-secondary education c_s, the cost of training provision by education $c_x(s)$, and the productivity gain from high-tech employment for high-skill workers—that is, the complementarity between high-skill labour and technology. In both cases, the cost parameters are set to match the shares of high-skill workers and training participants by education in the given years. However, since the WES only includes information on high-technology employment for the final steady state, I am required to follow a different procedure for calibrating the complementarity parameter $\varepsilon_k(s_1)$. For the final steady state, I compute an average wage premium of 1.28 for high-skill workers in high-technology firms directly from the WES micro-data. For the initial steady state, I set $\varepsilon_k(s_1)$ to match an average increase in the college wage premium of 12.5 percent, given a college premium in the final steady state of 1.37. The targeted value of 12.5 percent is chosen as an intermediate value between the estimate of 0 percent reported in Kryvtsov and Ueberfeldt (2009) and 25 percent reported in Boudarbat, Lemieux, and Riddell (2010).

1.6 Quantifying the Importance of Training

The calibrated model provides an exact decomposition of the increase in the college wage premium into three parts: (1) changes in the relatively complementarity between skills and technology, which is governed by the parameter $\varepsilon_k(s_1)$; (2) changes in the share of high-skill workers, which is governed by the parameter c_s; and (3) changes in the shares of training participants by education, which are governed by the parameters $c_x(s_0)$ and $c_x(s_1)$. Accordingly, in this section, I use the model to perform a decomposition analysis and, ultimately, quantify the contribution of training on the rise in the college wage premium.

To this end, I start in the initial steady state and sequentially allow each subset of parameters to adjust to their final steady state values. Following each simulation, I compute the model-implied level of, and change in, the college wage premium and measure their difference from the counterparts computed in the previous step. Finally, I

Table 1.6: Parameters determined jointly in equilibrium

Parameter		Value	Target	Data	Model
1980 Calibration:					
Education cost	c_s	90.00	Share of high-skill workers	0.23	0.23
Low-skill training cost	$c_x(s_0)$	10.91	Share of low-skill training participants	0.11	0.13
High-skill training cost	$c_x(s_1)$	3.04	Share of high-skill training participants	0.35	0.36
High-skill prod. from high-tech	$\varepsilon_k(s_1)$	1.06	Average increase in the college wage premium	0.13	0.10
2000 Calibration:					
Education cost	c_s	75.86	Share of high-skill workers	0.68	0.67
Low-skill training cost	$c_x(s_0)$	9.07	Share of low-skill training participants	0.26	0.27
High-skill training cost	$c_x(s_1)$	3.45	Share of high-skill training participants	0.42	0.45
High-skill prod. from high-tech	$\varepsilon_k(s_1)$	1.07	Average high-skill high-tech premium	1.28	1.20
Both Years:					
Vacancy posting cost	κ	3.48	Average unemployment rate	0.07	0.06
Productivity from education	h_{s_1}	1.09	Average college wage premium	1.37	1.36
Low-skill prod. from training	$\varepsilon_x(s_0)$	1.05	Average low-skill training premium	1.16	1.17
High-skill prod. from training	$\varepsilon_x(s_1)$	1.01	Average high-skill training premium	1.14	1.19
Low-skill prod. from high-tech	$\varepsilon_k(s_0)$	1.07	Average low-skill high-tech premium	1.19	1.16
Low-skill cost of high-tech	$c_k(s_0)$	2.05	Share of low-skill high-tech workers	0.36	0.36
High-skill cost of high-tech	$c_k(s_1)$	16.60	Share of high-skill high-tech workers	0.57	0.58

Notes: The targets for all parameters in the middle (2000 Calibration) and bottom (Both Years) panels are drawn from the Workplace and Employee Survey (WES). For the 1980 calibration, the share of high-skill workers is taken from publicly-available OECD data, training participation rates by education are taken from Statistics Canada and Human Resources Development Canada (2001), and the average increase in the college wage premium is an intermediate value of the estimates reported by Kryvtsov and Ueberfeldt (2009) and Krueger et al. (2010).

Table 1.7: Decomposition Results

	Setting			
	(1)	**(2)**	**(3)**	**(4)**
Technology-Skill Complementarity, ε_k	–	✓	✓	✓
Education Cost, c_s	–	–	✓	✓
Training Costs, c_x	–	–	–	✓
College wage premium (%)	27.78	31.51	32.22	30.50
College wage premium (Ratio)	1.32	1.37	1.38	1.36
Share of high-skill workers	0.23	0.32	0.71	0.67
Share of high-skill trained workers	0.36	0.95	0.53	0.45
Share of low-skill trained workers	0.13	0.04	0.00	0.27
Share of high-skill high-tech workers	0.58	0.97	0.56	0.58
Share of low-skill high-tech workers	0.36	0.65	0.28	0.36

– indicates that the parameter value is held fixed at initial (1980) value.
✓indicates that the parameter value varies across steady states.

compute the relative contribution of each channel by comparing the estimated difference at each step to the overall change in the premium.[15] The results from performing this exercise are summarized in Table 1.7.

Consider first columns (1) and (2) of Table 1.7. The difference in moments reported between the two columns solely reflects changes in the relative complementarity between high-skill labour and technology, that is, $\varepsilon_k(s_1)$. Relative to the initial steady state, the college wage premium in this environment increases by 3.73 percentage points, or 13.4 percent over the 1980 level. The *increase* in the college wage premium therefore overshoots the model-attained value of 9.79 percent by an additional 37 percent. This phenomenon occurs because the increase in technology-skill complementarity encourages firms to invest in the high-technology and raise their demand for high-skill workers. Combined with the positive association between technology and training, high-skill workers experience an extra boost in the probability of receiving training on the job. Therefore, increased labour demand and training participation for high-skill workers jointly act to raise the college wage premium.

The difference in values reported in columns (2) and (3) of Table 1.7 reflect the joint impact of changes in technology-skill complementarity and the costs of post-secondary education. In this environment, the college wage premium increases even further from

[15]Because training has a dampening effect on the premium, I take the sum of the absolute changes as the measure of the overall change.

31.51 percent to 32.22 percent, implying an increase of 16 percent over the 1980 level. Although the lower cost of education generates a much larger equilibrium share of high-skill workers, its effect on the college wage premium is completely undone by the response of training. In particular, the large increase in the supply of high-skill workers encourages new entrant firms to shift from posting vacancies with training to posting vacancies for high-skill workers. In the new equilibrium, the training gap widens as virtually no low-skill workers participate.

Finally, Column (4) of Table 1.7 reports the estimated moments from the final steady state—that is, after allowing all 4 parameters to vary across steady states. By comparing the difference in moments between columns (3) and (4), I obtain an measure of how much changes in training costs have dampened the increase in the college wage premium. In particular, with training costs held fixed (Column (3)), the college wage premium increases by 16 percent over its 1980 value. Once training costs are allowed to vary (Column (4)), the college wage premium increases by only 9.79 percent over its 1980 value. Therefore, absent changes in training costs, the increase in the college wage premium would have been 63 percent larger between 1980 and the early 2000s. On the other hand, the measured *level* of the college wage premium would have been 6 percent larger in the final steady state.

Across all experiments, the level of the college wage premium experiences an absolute change of 8.33 percentage points. Most of this change—approximately 60 percent—is driven by changes in technology-skill complementarity, indicating that technological change is indeed the driving force of the college wage premium. Changes in training account for the next largest share—28 percent—and, in fact, dampen the premium. This finding is consistent with recent work by Doepke and Gaetani (2020), who find increased on-the-job human capital accumulation among low-skill workers helps to explain the relatively subdued growth in the college wage premium experienced in Germany compared to the United States.

1.7 Conclusion

A large quantitative literature has established a tight link between the introduction of new technologies and rising college wage premia around the globe. In this paper, I use matched employer-employee survey data from Canada to highlight that the diffusion of technological change across firms has also generated a meaningful increase in training participation among workers—particularly, low-skill workers—which, because of the earnings premium associated with training, has dampened the college wage premium over time.

To reach this conclusion, I start by documenting a new set of facts using the Canadian micro-data. I show that (1) high-technology firms tend to be relatively more productive, provide more training, and hire more high-skill employees than low-technology firms; (2) training participants earn between 3 to 14 percent more on hourly earnings relative to non-participants; and (3) relative increases in training participation among low-skill workers reduce the college wage premium.

I use these features of the Canadian data to discipline a quantitative model of the labour market and generate the structure necessary to analyze the main drivers of the college wage premium. Consistent with existing evidence, I find that technology-skill complementarity accounts for the greatest share (60 percent) of absolute variation in the college premium over time. Importantly, however, I also find that a relative expansion in training participation among low-skill workers played a crucial role in dampening the college wage premium between 1980 and the early 2000s. In particular, absent changes in training costs, the rise in the college wage premium would have been 63 percent larger over this period. This latter finding suggests that potential gains in equity between high- and low-skill workers may be recovered if policymakers develop programs to encourage training participation among low-skill workers.

References

Açıkgöz, Ömer Tuğrul and Barış Kaymak (2014). "The Rising Skill Premium and Deunionization". *Journal of Monetary Economics*, 63, 37–50.

Amiti, Mary and Lisa Cameron (2012). "Trade Liberalization and the Wage Skill Premium: Evidence from Indonesia". *Journal of International Economics*, 87 (2), 277–287.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger (1998). "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics*, 113 (4), 1169–1213.

Bartel, Ann P. and Nachum Sicherman (1998). "Technological Change and the Skill Acquisition of Young Workers". *Journal of Labor Economics*, 16 (4), 718–755.

Black, Sandra E. and Lisa M. Lynch (2004). "What's Driving the New Economy?: The Benefits of Workplace Innovation". *The Economic Journal*, 114 (493), F97–F116.

Blundell, Richard, Lorraine Dearden, and Costas Meghir (1996). *The Determinants and Effects of Work-Related Training in Britain*. R50. IFS report.

Blundell, Richard, David A. Green, and Wenchao Jin (2022). "The U.K. as a Technological Follower: Higher Education Expansion and the College Wage Premium". *Review of Economic Studies*, 89 (1), 142–180.

Boothby, Daniel, Anik Dufour, and Jianmin Tang (2010). "Technology Adoption, Training and Productivity Performance". *Research Policy*, 39 (5), 650–661.

Boudarbat, Brahim, Thomas Lemieux, and W. Craig Riddell (2010). "The Evolution of the Returns to Human Capital in Canada, 1980–2005". *Canadian Public Policy*, 36 (1), 63–89.

Bound, John and George Johnson (1992). "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations". *American Economic Review*, 82 (3), 371–392.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt (2002). "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence". *The Quarterly Journal of Economics*, 117 (1), 339–376.

Burbidge, John B., Lonnie Magee, and A. Leslie Robb (2002). "The Education Premium in Canada and the United States". *Canadian Public Policy*, 203–217.

Burstein, Ariel and Jonathan Vogel (2017). "International Trade, Technology, and the Skill Premium". *Journal of Political Economy*, 125 (5), 1356–1412.

Card, David and John E. DiNardo (2002). "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles". *Journal of Labor Economics*, 20 (4), 733–783.

Den Haan, Wouter J., Garey Ramey, and Joel Watson (2000). "Job Destruction and Propagation of Shocks". *American Economic Review*, 90 (3), 482–498.

Dix-Carneiro, Rafael and Brian K. Kovak (2015). "Trade Liberalization and the Skill Premium: A Local Labor Markets Approach". *American Economic Review*, 105 (5), 551–557.

Doepke, Matthias and Ruben Gaetani (2020). "Why Didn't the College Premium Rise Everywhere? Employment Protection and On-the-Job Investment in Skills". *NBER Working Paper Series*, (27331).

Flinn, Christopher, Ahu Gemici, and Steven Laufer (2017). "Search, Matching, and Training". *Review of Economic Dynamics*, 25, 260–297.

Flinn, Christopher and Joseph Mullins (2015). "Labor Market Search and Schooling Investment". *International Economic Review*, 56 (2), 359–398.

Glitz, Albrecht and Daniel Wissmann (2021). "Skill Premiums and the Supply of Young Workers in Germany". *Labour Economics*, 72, 102034.

He, Hui (2012). "What Drives the Skill Premium: Technological Change or Demographic Variation?" *European Economic Review*, 56 (8), 1546–1572.

Katz, Lawrence F. and Kevin M. Murphy (1992). "Changes in Relative Wages, 1963–1987: Supply and Demand Factors". *The Quarterly Journal of Economics*, 107 (1), 35–78.

Kristal, Tali (2020). "Why Has Computerization Increased Wage Inequality? Information, Cccupational Structural Power, and Wage Inequality". *Work and Occupations*, 47 (4), 466–503.

Krueger, Dirk et al. (2010). "Cross-sectional facts for macroeconomists". *Review of Economic dynamics*, 13 (1), 1–14.

Krusell, Per et al. (2000). "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis". *Econometrica*, 68 (5), 1029–1053.

Kryvtsov, Oleksiy and Alexander Ueberfeldt (2009). *What Accounts for the US-Canada Education-Premium Difference?* Tech. rep. Bank of Canada.

Leuven, Edwin and Hessel Oosterbeek (2008). "An Alternative Approach to Estimate the Wage Returns to Private-Sector Training". *Journal of Applied Econometrics*, 23 (4), 423–434.

Lindner, Attila et al. (2022). "Firm-Level Technological Change and Skill Demand". *Centre for Economic Performance Working Paper*, (1857).

Lise, Jeremy et al. (2014). "Wage, Income and Consumption Inequality in Japan, 1981–2008: From Boom to Lost decades". *Review of Economic Dynamics*, 17 (4), 582–612.

Matsuda, Kazushige (2020). "Optimal Timing of College Subsidies: Enrollment, Graduation, and the Skill Premium". *European Economic Review*, 129, 103549.

Menzio, Guido and Shouyong Shi (2010). "Block Recursive Equilibria for Stochastic Models of Search On the Job". *Journal of Economic Theory*, 145 (4), 1453–1494.

Menzio, Guido and Shouyong Shi (2011). "Efficient Search On the Job and the Business Cycle". *Journal of Political Economy*, 119 (3), 468–510.

Mincer, Jacob (1958). "Investment in Human Capital and Personal Income Distribution". *Journal of Political Economy*, 66 (4), 281–302.

Moen, Espen R. and Åsa Rosén (2004). "Does Poaching Distort Training?" *The Review of Economic Studies*, 71 (4), 1143–1162.

Parent, Daniel (1999). "Wages and Mobility: The Impact of Employer-Provided Training". *Journal of Labor Economics*, 17 (2), 298–317.

Parro, Fernando (2013). "Capital-Skill Complementarity and the Skill Premium in a Quantitative Model of Trade". *American Economic Journal: Macroeconomics*, 5 (2), 72–117.

Salgado, Sergio (2020). "Technical Change and Entrepreneurship". *SSRN 3616568*.

Schaal, Edouard (2017). "Uncertainty and unemployment". *Econometrica*, 85 (6), 1675–1721.

Shephard, Andrew and Modibo Sidibe (2019). "Schooling Investment, Mismatch, and Wage Inequality". *PIER Working Paper*.

Shi, Shouyong (2002). "A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology". *Review of Economic Studies*, 69 (2), 467–491.

Statistics Canada and Human Resources Development Canada (2001). *A Report on Adult Education and Training in Canada: Learning a Living*. Government of Canada, Ottawa, Ontario.

Takahashi, Shuhei and Ken Yamada (2022). "Understanding International Differences in the Skill Premium: The Role of Capital Taxes and Transfers". *Journal of Economic Dynamics and Control*, 143, 104511.

Velden, Rolf Van Der and Ineke Bijlsma (2016). "College Wage Premiums and Skills: A Cross-Country Analysis". *Oxford Review of Economic Policy*, 32 (4), 497–513.

Walker, Ian and Yu Zhu (2008). "The College Wage Premium and the Expansion of Higher Education in the U.K." *The Scandinavian Journal of Economics*, 110 (4), 695–709.

Zentler-Munro, David (2021). "Rising Wage Inequality: Technological Change and Search Frictions". *Working Paper*, 1–60.

Chapter 2

Firm Sales and the Firm Life Cycle

Firm Sales and the Firm Life Cycle

Bettina Brüggemann
McMaster University
bruеggeb@mcmaster.ca

Zachary L. Mahone
McMaster University
mahonez@mcmaster.ca

Thomas Palmer
McMaster University
palmet8@mcmaster.ca

Abstract

Ownership changes are common across firms of all sizes, and they have meaningful impacts on subsequent firm performance. Using a panel of Canadian administrative data we document that sales are an important margin in the firm life cycle, larger than exit rates for employer firms. Applying an event-study framework, we find that (a) survival rates initially decline post sale, levelling off after three years and (b) conditional on survival, profits are permanently higher. Embedding firm sales in a model of firm dynamics we show that sales are quantitatively important to understanding entry, growth and exit dynamics. We estimate that 13% of entrants survive exclusively due to the option value of sale. Among small firms, 18% of average log employment growth is accounted for by realized ownership changes. In the stationary distribution, firm sales operate at the tails, allowing smaller firms to survive and magnifying the size of top firms. Finally, we quantify the importance of incorporating ownership changes for understanding the aggregate response to policy changes.

Keywords: *Firm Dynamics, Entrepreneurship, Business Ownership, Firm Sales.*

2.1 Introduction

Why do some firms succeed and grow while others exit? Among startups, an existing literature points to the importance of ownership changes in driving firm performance (e.g., Ewens and Marx, 2018; Becker and Hvide, 2022). Even for large firms, high profile examples of ownership changes, such as the acquisition and rebranding of X (previously

Twitter), reinforce this idea.[1]. Nonetheless, relatively little is known about the frequency and impact of owner changes across the entire distribution of firms, and models of firm dynamics do not typically allow for owner changes as a margin for adjustment.

In this paper we ask to what extent do large ownership changes, which we think of as firm sales, drive firm dynamics and what are their macroeconomic implications? We provide empirical evidence of the widespread nature of ownership changes across firms of all sizes and their significant impact on both profitability and survival. Guided by this evidence, we then build a structural model of firm dynamics with changes in ownership and show that sales are important to how we understand the entry, continuation and growth of firms, as well as aggregate TFP. To highlight the macroeconomic importance of the firms sales margin, we compare policy implications across calibrated models with and without firm sales. We find that incorporating ownership changes in a model of firm dynamics, especially the option value to sell, even if unrealized, has a significant impact on macroeconomic responses to policy. In this sense, our paper builds on a recent literature (Pugsley, Sedlacek, and Sterk, 2021; Jaimovich, Terry, and Vincent, 2023) that emphasizes the need for richer modelling of firm-level growth dynamics.

Our data come from the Canadian Employer Employee Dynamics Database (CEEDD) and cover the period 2001-2017. The CEEDD is a collection of linkable administrative databases that rely on tax and other filings, representing the universe of firms and employees operating in Canada. In our analysis, we restrict attention to incorporated employer firms (see Section 2.2.1 for a detailed discussion). Importantly, all owners (corporate and individual) with a greater than 10% stake in a company must file an annual declaration, which we can link to firm-level information, allowing us to track ownership for the vast majority of firms over time.

We define a firm sale as occurring when more than 50% of equity changes ownership between two consecutive years, and we focus on external sales (equity transfers to new owners), as opposed to internal sales (equity changes among existing owners). We find that firm sales are widespread. Among employer firms, the average sales rate (0.86%) exceeds the average exit rate (0.67%). Sales increase in firm employment, both in absolute terms and relative to exit: Among firms with at least 100 employees, the fraction of firms sold is almost 3 times larger than the fraction of firms exiting (1.72% vs. 0.64%), and more than twice the share of small firms sold (0.7%).

[1]Employment at Twitter has reportedly fallen by 80% (https://www.cnbc.com/2023/01/20/twitter-is-down-to-fewer-than-550-full-time-engineers.html) and advertiser traffic declined by 16.5% in the year after acquisition (https://apnews.com/article/twitter-x-elon-musk-takeover-anniversary-ac2cb6419d93d64086cc9ad980c5a57a).

We next examine whether ownership changes have a meaningful impact on firm performance. To do this, we employ an event-study framework to compare sold firms with their non-sold counterparts. We follow a coarsened-matching framework similar to Smith et al. (2019), matching sold firms with non-sold firms according a number of criteria reflecting business and owner characteristics. We then study the impact of sales on measures of profit per worker and survival. The results are striking. The survival rates of sold firms initially decline post-sale, levelling off after three years. Conditional on survival, however, profits are persistently higher.

Of course, due to the endogenous nature of a firm sale we cannot directly treat these findings as causal. We consider a number of explanations for the observed patterns: risky ownership changes, selection and anti-competitive behaviour. We argue that the evidence supports the first interpretation, namely that new owners are risky: On average, they improve profitability but some will fail. This interpretation is also in line with recent empirical studies that leverage the death of owners Smith et al. (2019) and team members Choi et al. (Forthcoming), finding a significant causal impact on firm performance.

To evaluate the quantitative importance of firm sales for entry, exit and growth dynamics, we incorporate the possibility of sales within a standard model of firm dynamics (e.g., Hopenhayn and Rogerson, 1993). To do this, we separately model a fixed firm productivity and a firm-owner match productivity (in addition to a time-varying, idiosyncratic component). Firm sales, then, represent a firm climbing up the match productivity ladder, which is the source of gains from trade. To capture the riskiness that our data suggest, we treat a new firm-owner match as noisily observed ex-ante. Sales will occur as long as the expected gains from trade are sufficiently large, but ex-post realizations of the match productivity will lead to higher initial exit rates of sold firms, conditional on age and size. It is worth emphasizing here that we do not model owners/entrepreneurs explicitly, as in other recent papers that incorporate firm sales (Mahone, 2023; Guntin and Kochen, 2023; Bhandari, Martellini, and McGrattan, 2022; Gaillard and Kankanamge, 2020). In the data, the mapping between owners and firms is noisy, with many corporate owners and serial entrepreneurs. Relative to other papers, we treat owners as in infinite supply and focus exclusively on the firm side and the risk implied from changes in ownership.

We calibrate the model to match salient features of our data with respect to firm dynamics and ownership changes. Two key parameters for the model are the relative importance of the firm and match productivities as well as their correlation in the entrant sampling distribution. The first finding is that the match component is on average much smaller than the firm component, which is informed largely by our event study results. If the match component is too small, ownership changes will have no impact on firm performance, which is counterfactual to the data. By contrast, if the match component is too large, then firms with low match productivities will exit, as sales happen only occasionally. However, these are precisely the firms that should be sold in equilibrium, and their exit makes it hard for the model to reproduce the sales moments. The second quantitative finding is that firm and match productivities are negatively correlated, which implies that larger firms on average have higher to climb up the ladder from birth than small firms. In the model, we see that more productive firms search more, while firms with better owners sell less often. To understand why, consider the case of a positive correlation between firm and match productivities. In this setting, larger firms will search more but, as they have better matches on average, will sell less. These competing forces make it difficult to achieve the increasing rate of firm sales by size in the data. Thus, the calibration delivers a negative correlation.[2].

We next turn to quantifying the role of ownership changes along three dimensions: (i) survival and growth dynamics (ii) stationary firm size distribution (iii) aggregate policy responses. In all cases, firm sales operate through two distinct mechanisms. The first is the impact of realized sales, which on average raise the owner-match quality and improve firm productivity. The second is the option value of sales. Even absent a sale occurring, the option to improve owner-match has a significant impact on firm values, particularly among high productivity firms with low quality owner matches. In all of our counterfactual analysis, we distinguish between the impact of realized sales and their option value.

We begin by studying the role of ownership changes for survival and growth dynamics. Sales impact survival primarily through the option value of future improvements in owner matches for young firms. In the calibrated model 13% of entrants survive exclusively due to the continuation value arising from a possible future sale. This impact is persistent as well, remaining at 11.5% by age ten. The impact of ownership changes on growth dynamics, however, are primarily driven by realized sales. We find that, among small firms, ownership changes account for 18% of average log employment growth. This impact falls in firm size and actually becomes negative among the largest firms.

[2]The negative correlation is consistent with preliminary empirical evidence reported in Brüggemann, Mahone, and Muris (2024)

We next aggregate up from these firm-level dynamics to consider the impact of ownership changes for the stationary firm size distribution. To do so, we compute a counterfactual economy in which wages are held fixed while the option to sell is removed. We find that the entire impact of sales is concentrated in the upper and lower quartiles of the firm size distribution, both of which are substantially compressed in the absence of sales. However, each quartile responds to a different sales mechanism. The bottom quartile of the calibrated model is populated by firms that enter due to the option value of a future improvement in the owner match. Removing sales causes these firms to exit rapidly, compressing the lower end of the distribution by nearly two-thirds. Meanwhile, the upper quartile in the benchmark economy is driven by realized sales. While most of these firms have fairly good owner-matches, the scale of these large firms implies that even small changes in ownership can have a significant impact on their size. Absent sales, the upper tail compresses by 85%. Thus, we find that ownership changes contribute substantially to firm size dispersion in the aggregate data, operating primarily at the tails.

Finally, we turn to analyzing the macroeconomic implications of incorporating ownership changes. To do this, we compare the aggregate economy's response to firm-level policies between our benchmark model and a model without firm sales, calibrated to match the same moments on firm size and growth dynamics. We consider a subsidy to operating firms, equivalent to 25% of the average fixed cost, paid for by a lump sum tax on households. We find that the model with no ownership changes is substantially more sensitive to the subsidy policy. Absent firm sales, output falls by 10.03%, relative to 6.54% in the benchmark. Average TFP falls by 5.15% relative to 2.69% in the benchmark. Similarly,exit rates and average firm size decline more strongly in the model without firm sales. The key operating force is the option value of sales. In the benchmark model, the option to improvement owner matches contributes substantially to the continuation value of small firms. As a result, the impact of the subsidy on continuation decisions is more limited, as compared to the model without owner changes, where firm's have less scope for significant growth.

In summary, this paper makes two contributions to the literature. First, we document that ownership changes are widespread across the firm size distribution and that they meaningfully impact firm performance. Second, we use structural model of firm dynamics, augmented to allow for firm sales, to quantify the impact of ownership changes on entry, exit, growth and the aggregate distribution. We find that both realized sales as well as their option value have quantitatively significant implications for the nature of firm growth and aggregate policy.

The remainder of the paper is organized as follows. Section 2.2 introduces the CEEDD data and presents our main empirical results. Section 2.3 develops a model of firm dynamics with risky ownership changes. Section 2.4 discusses the model calibration, intuition for how the model and data map together, and computes our primary counterfactuals for the impacts of realized sales and their option value. Section 2.5 compares the policy implications of a model with and without ownership changes. We then conclude.

2.2 Empirical Evidence on Firm Sales

We start this section by describing the data sets we are using for our analysis of firm sales, before presenting key statistics on firm sales and statistically analyzing both the drivers and the effects of firm sales in the firm life cycle.

2.2.1 Canadian Employer Employee Dynamics Database

We use the Canadian Employer Employee Dynamics Database (CEEDD) for the years 2001-2017. The CEEDD is a collection of administrative datasets maintained by Statistics Canada that allows for linking employers and employees. Our firm level data comes from the National Accounts Longitudinal Data File (NALMF) which includes annual data on all incorporated or (unincorporated) employer firms operating in Canada. In the NALMF, we observe more than one million firms each year, with information on revenue, expenses, employment and many other firm characteristics such as industry or firm age. Individual data on workers and owners comes from T4 employment files, which provides data on labour income, as well as the T1 Family Files which includes information on age and family relationships. Finally, ownership information is obtained from T2 Schedule 50 (T2S50) filings, which are required of all equity holders (individual or corporate) with a stake that is 10% or larger.

The basis for firm files in NALMF is Statistics Canada's Business Register. When a new business registers in Canada it receives a Business Number (BN) from the Canada Revenue Agency (CRA). An *enterprise*, which contains one or more BNs, refers to the highest level of the Business Register statistical hierarchy and is associated with a complete set of financial statements.[3] The enterprise is the unit of analysis in the NALMF, and we will use the terms *firm* and *enterprise* interchangeably.

[3] See `https://www23.statcan.gc.ca/imdb/p3Var.pl?Function=UnitI&Id=140364`

FIGURE 2.1: Enterprise Ownership Structure in the NALMF

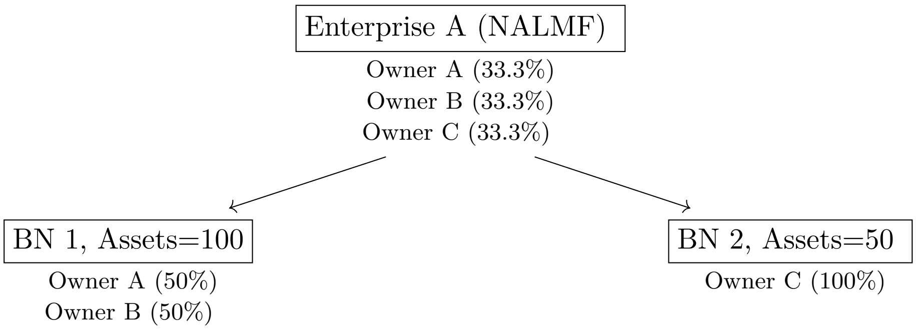

Ownership data from T2S50 filings is at the BN level and so must be aggregated to the enterprise level in order to be matched to the NALMF. We do this using BN-level asset information as ownership weights in the aggregated enterprise. Figure 2.1 presents a diagram example of an enterprise containing two business numbers. Ownership weights of Enterprise A are computed by multiplying the equity share of an individual by the share of assets their business represents for the enterprise. For example, Owner A's share of Enterprise A is computed as $0.5 \times \frac{100}{150} = 0.333$.

Our baseline sample selection requires a firm to have no missing years of observation or missing owner information. Because we are interested largely in employer firms we require a firm to have non-zero employment for at least one year in the sample (by which we mean a tax schedule T-4 is filed). Finally, as is common in the literature (e.g. Smith et al., 2019), we remove all government and FIRE NAICS categories. This avoids certain industries where new business creation may not be closely linked to new economic activity.

2.2.2 An Overview of Firm Sales

We define a firm as exiting in year t if they are not observed in $t+1$ and are not identified as a predecessor to a firm in year $t+1$.[4] Changes in equity ownership are inferred through T2S50 filings. We define sales as changes in the controlling share of a company (at least 50%) between t and $t+1$. We separately identify sales as *external* (sales to new owners) and *internal* (equity trades among existing owners).

[4]Through mergers and break ups of companies, new firms may be created (or disappear) that are in fact partial continuations of existing firms. The linked employee-employer nature of the CEEDD data allows us to identify these firms through shared employees, addresses, names and physical location. Statistics Canada provides a basic set of files identifying predecessors and successors.

FIGURE 2.2: Sales Rates Over Time (2001-2017)

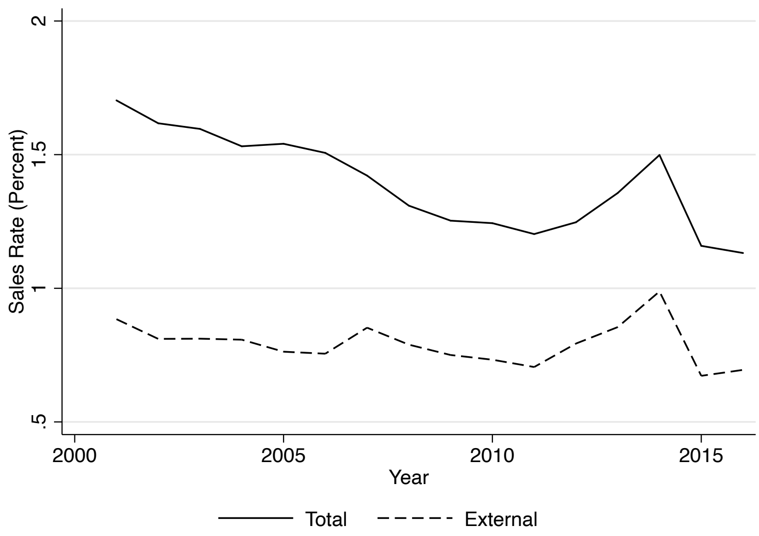

We begin by characterizing some basic features in the data. Figure 2.2 plots the pattern of sales rates over time. The average sales rate in our sample is 1.4 percent, while for external sales only it is 0.8 percent. Total sales rates have been declining over the sample, but this is largely driven by internal sales. The external sales pattern (dashed line) is roughly flat.

Figure 2.3 puts sales rates in the context of firm exit rates, a margin that has received much attention in the firm dynamics literature.[5] We see that for all but part-time employer firms, sales rates are at least as large as exit, and grow in firm size. We view this as evidence that the sales is a quantitatively important margin in the firm life cycle.

We also look at who buyers and sellers are in this market. We distinguish between corporate owners and individual owners, where we separately consider serial and non-serial owners. We follow Brandt et al. (2022) in defining a *serial* owner as an individual observed owning two or more firms in the sample. We define a firm as sold (bought) by a corporate owner if the majority of equity pre-sale (post-sale) is owned by a corporation. Looking at Table 2.1, we see that corporations are heavily over-represented in the market, as both buyers and sellers, while individuals are under-represented.

[5]We do not use age here because firms enter the NALMF upon incorporating or hiring. In the data, age information based on various founding year definitions are very noisy.

Figure 2.3: Sales and Exit Rates by Size

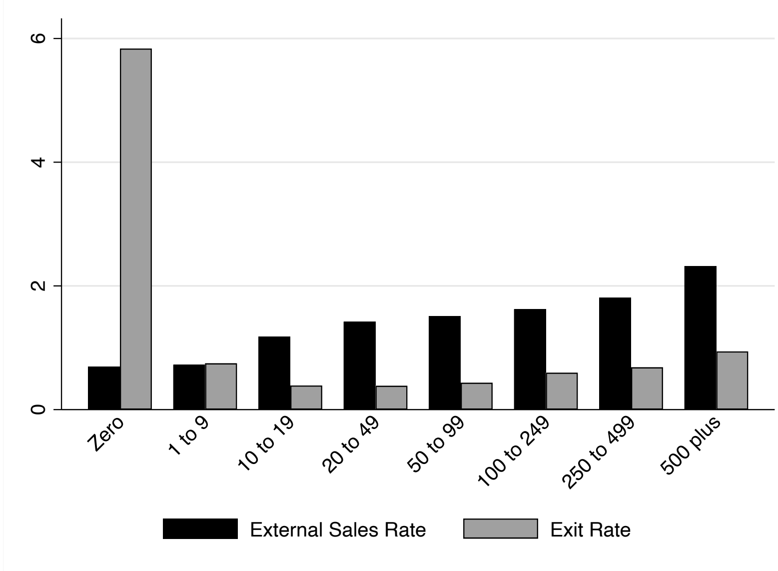

Table 2.1: Owner Types by Buyers and Sellers

	Firms		
Majority Type	All	Sold	Bought
Corporate	14.01	25.25	48.57
Individual, Non-Serial	50.60	29.82	26.90
Individual, Serial	35.40	44.92	24.53

Finally, to get an overview of what drives firm sales, we estimate a logit model on firm and owner characteristics. Firm variables include profit-per-employee, size, lagged growth, age, and (2 digit) NAICS. For our measure of profits, we use the difference of reported revenues and expenses. Owner variables include majority type (individual, individual serial, corporate) and, for individuals, owner age. The results in Table 2.2 report the results for key explanatory variables of interest.

On the firm side, profits appear only weakly positively associated with sales likelihood. Notably, while sales rates appear increasing in firm size (as suggested by Figure 2.3), they are decreasing in firm age. This suggests that age and size, while correlated, contain

TABLE 2.2: Predictors of Firm Sales

Dependent Variable: Pr[External Sale = 1]	
Profit Per Employee	$4E-9^{*}$
Firm Size	
1-9 employees	0.090^{***}
10-19 employees	0.315^{***}
20-49 employees	0.359^{***}
50-99 employees	0.325^{***}
100-249 employees	0.329^{***}
250-499 employees	0.361^{***}
>500 employees	0.527^{***}
Average Pre-Sale Employment Growth	
$> 2\%$	0.014
$< -2\%$	0.067^{***}
Firm Age	
2-4 years	-0.243^{***}
4-6 years	-0.341^{***}
6-8 years	-0.434^{***}
8-10 years	-0.471^{***}
10-15 years	-0.511^{***}
15-20 years	-0.520^{***}
> 20 years	-0.502^{***}
Owner Type	
Non-serial, Individual	-1.008^{***}
Serial, Individual	-0.245^{***}
Average Owner Age	
25-35 years	-0.296^{***}
35-45 years	-0.399^{***}
45-55 years	-0.382^{***}
55-65 years	-0.116^{***}
65-75 years	0.191^{***}
> 75 years	0.100^{***}
Year dummies	✓
NAICS dummies	✓
N	10,719,140

$^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.10$. Time fixed effects included.

different information, a point emphasized in the literature on firm learning.[6]

For owners, where the base category is corporate, we see that firms which are majority corporate-owned are the most likely to be sold, followed by those owned by serial entrepreneurs and then individuals. This is consistent with the over-representation of corporations in the market for firms, observed in Table 2.1. Lastly, the point estimates

[6] See Jovanovic (1982).

for owner age suggest a U-shape: young owners are initially more likely to sell than middle aged, after which the sales likelihood increases again as owners age into retirement.

2.2.3 The Dynamic Impact of Sales

We next turn to analyzing the impact of sales. To do this, we use an event study framework to compare sold firms with similar non-sold firms in the sample. To create a sample of similar, non-sold, firms, we follow a coarsened-matching approach, as in Smith et al. (2019). We then focus on two outcomes: profit per average worker and survival. For inclusion in our sample, we require sold firms to be sold only once and focus here on external sales. A non-sold firm must have never been sold in our sample and cannot exit before the time of the match. For inclusion in the sample, firms must revenue greater than \$10,000 the year prior to sale, have positive revenue in the four years prior to sale and have positive employment in at least one of the four years prior to sale. With the resulting set of firms, we then construct a coarsened matching sample using similar criteria as in Smith et al. (2019). Firms in the sold group are matched with the non-sold group in the year before sale according to five criteria (i) Revenue quintile, (ii) Number of owners, (iii) Average owner age, (iv) Firm age, and (v) 2-digit NAICS.

For each sold firm j and match j' in panel year t, define the difference:

$$Y_{j,j',t} = y_{j,t} - y_{j',t}, \tag{2.1}$$

where the outcome variable y represents profit per average pre-sale employees or survival (1=survive). We then estimate the following regression equation:

$$Y_{j,j',t} = \sum_{k \in \{-4,-3,-2,0,1,2,3,4\}} \beta_k I^k_{j,t} + \epsilon_{j,t}, \tag{2.2}$$

where $I^k_{j,t}$ is an indicator variable for firm j's event year k in panel year t. Year dummies are also included and standard errors are clustered on the match. It is worth emphasizing here that we are not interpreting the point estimates of Equation 2.2 as causal. The treatment in such an analysis is commonly argued to be exogenous, as in the death of owners studied in Smith et al. (2019). In our case, firm sales are clearly endogenous, non-random outcomes. We view our results as summarizing the impact of sales for firms after controlling for characteristics commonly thought to account for firm dynamics (our matching criteria).

FIGURE 2.4: Firm Survival Upon Sale

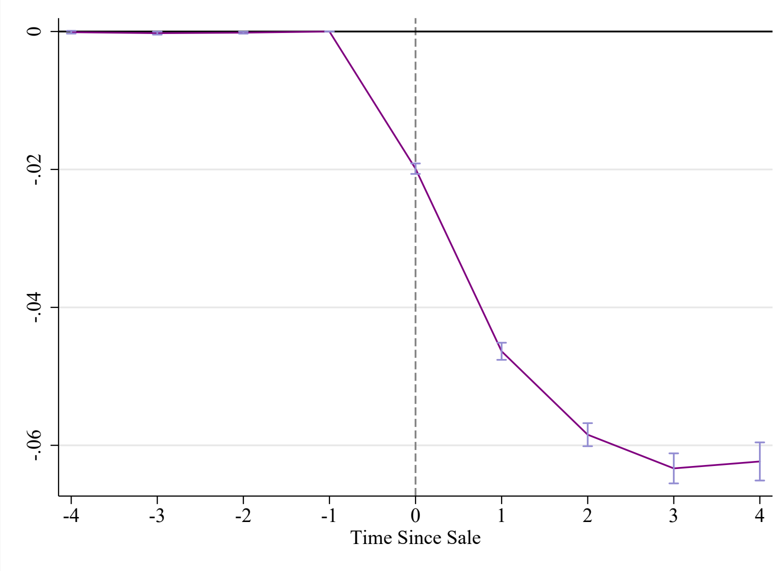

FIGURE 2.5: Firm Profits Upon Sale

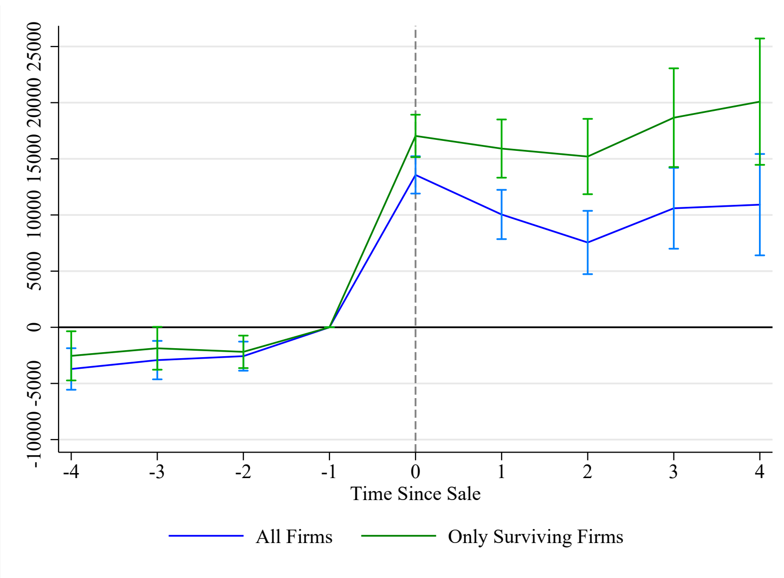

Figure 2.4 displays survival of sold firms relative to the non-sold, matched sample over four years pre- and post-sale. Point estimates, which are the differences in conditional means, show that exit rates are higher for sold firms in the first 3 years post sale. The flattening of the plot in years three to four indicates that eventually sold firms exit at the rate of their counterparts (so that the differential exit rate is not changing).

Our second main results concerns profit per average pre-sale worker (i.e. the denominator is held fixed). Figure 2.5 shows the impact of sales on profits. The blue line denotes the impact of sales on all firms, while the green line conditions only firms surviving in the four years post-sale. There is a clear jump in firm profits upon sale, which is persistent and sizeable, with point estimates around $15,000 per worker among surviving firms.

2.2.3.1 Interpreting the Findings

The shape of survival in the initial years after sale recalls the steep decline and then flattening of new-firm exit rates that is well known. In the same way that models studying entry and growth dynamics have emphasized learning, a natural way to interpret the results above is a story of ex-ante uncertainty about an owner-firm match (or, alternatively, owner ability). Through this lens, the shape of the survival function reflects that some new owners perform poorly, due to i.e. a failed growth strategy (see X for an example). For sales to take place, the expected match quality of a new owner must be higher than the existing one, explaining the persistent increase in profits conditional on survival observed in Figure 2.5.

It is of course also possible that ex-ante selection drives our results. For example, it may be that firms with a higher risk of exit tend to be sold. One way to test for this is to re-run the coarsened matching algorithm including an additional matching criteria based on pre-sale trends. We separately consider *turnaround* firms, firms that are on average shrinking in the four years prior to sale, and *growth* firms, firms that are instead growing during that same pre-sale period. We then run the same event study, comparing sold turnaround firms with non-sold, turnaround matches and similarly for growth firms. The results as shown in Appendix Figure A2.1 exhibit the same patterns as Figure 2.4, which suggests that selection (as captured by pre-sale growth rates) does not appear to explain the results. As an additional way to handle this, the model developed in Section 1.4 will also feature selection, allowing us to separate its role in the results.

Another concern that arises with our findings is a story of competition. There is ample anecdotal evidence of larger firms buying up and “killing" competitors in order to maintain market power. To examine this, we run the event study on separate samples conditional on buyer type. We divide the sample into firms sold to individuals or to corporate buyers. If the sales patterns found above were explained by anti-competitive behaviour, we would expect to see corporations driving the decline in survival. Appendix Figure A2.2 shows that both subgroups feature the same general patterns of survival and profits. To be sure, the negative impact on survival lags the sales date and is more attenuated for individual owners, but both types experience a decline in survival. We interpret this as evidence that, while such killer acquisitions exist (and may be important in certain industries), they cannot explain the broad pattern in sales that we observe.

2.3 Model

In this section, we embed firm sales into a model of firm dynamics in the spirit of Hopenhayn and Rogerson (1993). To do so, we model firm-level productivity as a combination of owner and firm fixed-effects, as well as a time-varying idiosyncratic component. This approach allows us to distinguish owners from firms, while avoiding separately modelling owner entry and exit.[7]

2.3.1 Environment

The economy is populated by an endogenous mass M of heterogeneous incumbent and potential entrant firms.

Incumbent firms are heterogeneous in productivity a_{ijt}, which consists of a fixed owner component ε_i, a fixed firm component γ_j, and a time-varying idiosyncratic component z_{ijt}. We assume that the owner and firm fixed effects are drawn from a known joint distribution $G(\varepsilon, \gamma)$ with continuous density $g(\varepsilon, \gamma)$.[8] In each period, incumbents make decisions over whether to operate or exit (subject to a fixed operating cost $f_c \geq 0$), how much labour to hire (subject to the market wage $w \geq 0$), whether to search for a new owner (subject to a search cost $c_s \geq 0$), and, conditional on searching, whether to sell in order to maximize the present value of profits. In doing so, incumbents face a series of shocks that affects their ability to operate, search, and sell. Specifically, with probability δ, an incumbent exogenously exits the economy and, with probability $(1-\eta)$, an incumbent is exogenously precluded from searching—and, therefore, selling—to a new owner.

In addition, there is a mass of potential entrants μ_e that decides whether to enter to maximize the expected value of entry each period subject to an entry cost κ. To do so, a potential entrant draws a signal s_i of owner productivity and a realization γ_j of

[7]Empirically, the average firm has more than one owner and many firms have corporate owners. While these features are interesting, we abstract from modelling owner dynamics to focus exclusively on the firm.

[8]The specification of $G(\varepsilon, \gamma)$ is general enough to allow for non-zero correlation between owner and firm effects.

firm productivity from a joint sampling distribution $G(s, \gamma)$, which has an equivalent support to the joint distribution $G(\varepsilon, \gamma)$. In making their entry decisions, potential entrants face two sources of idiosyncratic uncertainty. First, with probability ψ, the signal s_i of owner ability reflects true owner ability ($s_i = \varepsilon_i$), while with probability $(1 - \psi)$, the signal s_i is random, forcing the firm to draw a new value of owner ability ε_i from $G(\varepsilon, \gamma)$ prior to production. Second, potential entrants do not know the value of idiosyncratic productivity z_{ijt} prior to making their entry decision; instead, they must take expectations over the sampling distribution of idiosyncratic productivities $\overline{\Gamma}(z)$.

The timing within a period occurs as follows. First, production takes place: incumbents combine labour ℓ and productivity a_{ijt} to produce output $y_{ijt}(a_{ijt}, \ell)$ and pay workers the market wage w. Second, a fraction δ of incumbent firms are forced to exit the economy. Third, potential entrants pay make their entry decisions. Fourth, exogenous search shocks arise: a fraction $(1 - \eta)$ of incumbents are precluded from search activity. Fifth, conditional on surviving the exit and search shocks, incumbents make their endogenous exit and search decisions. Finally, idiosyncratic productivities z_{ijt} are updated according to the transition matrix $\Gamma(z, z')$.

2.3.2 Incumbents

In each period, incumbents make up to four decisions: whether to exit or operate, how much labour to hire, whether to search, and, conditional on searching, whether to sell. To produce output y_{ijt}, we assume that incumbents hire labor ℓ and combine it with firm productivity a_{ijt} as follows:

$$y_{ijt} = a_{ijt}\ell^{\alpha}, \tag{2.3}$$

where α represents the scale of production. In addition, we assume that firm productivity a_{ijt} aggregates in the following way:

$$a_{ijt} = \varepsilon_i \gamma_j z_{ijt}, \tag{2.4}$$

where ε_i represents the fixed owner component of firm productivity, γ_j the fixed firm component, and z_{ijt} the time-varying idiosyncratic component. Labour is paid a constant market wage rate w, which is determined by labor market clearing.[9]

Let $J(\varepsilon, \gamma, z)$ denote the value of operation for a firm with owner ability ε, firm productivity γ, and idiosyncratic productivity z. Then, $J(\varepsilon, \gamma, z)$ solves the following

[9]Absent the fixed cost, profit per employee in the model would be equal to the wage times a constant.

recursive equation:

$$J(\varepsilon,\gamma,z) = \max_{\ell}\ pa_{ijt}\ell^{\alpha} - w\ell + \int \max\left\{0, -f_c + \frac{1}{R}(1-\delta)\widetilde{J}(\varepsilon,\gamma,z)\right\} g(f_c)df_c, \quad (2.5)$$

where $f_c \geq 0$ is the per-period operating cost, $R > 0$ is the gross rate of return, and $\delta \geq 0$ is the probability of exogenous exit. We assume that each firm draws an i.i.d. operating cost f_c each period—denominated in units of the final good—from the exogenous distribution $G(f_c)$. The above problem implicitly defines a cut-off, f_c^*, such that an incumbent will exit if and only if its fixed cost exceeds the cutoff:

$$f_c^*(\varepsilon,\gamma,z) = \frac{1}{R}(1-\delta)\widetilde{J}(\varepsilon,\gamma,z). \quad (2.6)$$

The function $\widetilde{J}(\varepsilon,\gamma,z)$ represents the expected continuation value for the firm, which internalizes the transition probabilities over idiosyncratic productivity states and the exogenous and endogenous choices over search and sales. Specifically,

$$\widetilde{J}(\varepsilon,\gamma,z) = \sum_{z'}\Gamma(z,z')J(\varepsilon,\gamma,z') + \eta\max\left\{S(\varepsilon,\gamma,z) - \sum_{z'}\Gamma(z,z')J(\varepsilon,\gamma,z'), 0\right\}, \quad (2.7)$$

where $\eta \geq 0$ is the exogenous probability of being allowed to search for a new owner and $S(\varepsilon,\gamma,z)$ captures the value of search. In other words, if an incumbent is unallowed to search for exogenous reasons—which occurs with probability $(1-\eta)$—then its continuation value is simply $\sum_{z'}\Gamma(z,z')J(\varepsilon,\gamma,z')$.

Conditional on being allowed to search—which occurs with probability η—firms may still elect not to search for a new owner. In the model, this occurs because search is both costly and uncertain. Specifically, an incumbent that wishes to search must pay a cost worth c_s in units of the final good to draw a *signal* s of owner ability from the (marginal) sampling distribution $G_s(s)$. This latter point is crucial because it means that firms must make their sales decisions prior to observing the realization of the new owner's ability. In other words, a sale occurs only if the *expected* value of moving to a new owner, given the signal s, is higher than continuing under the current owner with (true) ability ε. Formally, the value of search $S(\varepsilon,\gamma,z)$ solves the following optimization problem:

$$S(\varepsilon,\gamma,z) = \int_s \max\left\{\sum_{z'}\Gamma(z,z')\left(\widehat{J}(s,\gamma,z') - J(\varepsilon,\gamma,z')\right), 0\right\} g_s(s)ds - c_s, \quad (2.8)$$

where $\widehat{J}(s,\gamma,z')$ represents the expected value of a the firm who draws signal s from the marginal distribution $G_s(s)$, that is,

$$\widehat{J}(s,\gamma,z') = \psi J(\varepsilon = s,\gamma,z') + (1-\psi)\int_x J(\varepsilon = x,\gamma,z')g_s(x)dx \tag{2.9}$$

where $\psi \in (0,1)$ represents the precision of the signal s. In words, Equation (2.9) states that, with probability ψ, the draw of signal s from the marginal distribution $G_s(s)$ reveals the true owner ability and, with probability $(1-\psi)$, the draw reveals nothing and the firm must draw a new owner ability ε prior to production.

2.3.3 Free Entry

In each period, there is a mass μ_e of potential entrants. Prior to making their entry decisions, potential entrants draw a signal s and firm fixed effect γ from the joint sampling distribution $G(s,\gamma)$. As noted in the previous subsection, the draw s will reveal the true owner ability with probability $\psi \in (0,1)$. In equilibrium, potential entrants will enter the economy until the expected value of entry equals the entry cost $\kappa \geq 0$:

$$\int\int\Big(\sum_z \overline{\Gamma}(z)\widehat{J}(s,\gamma,z)\Big)g(s,\gamma)dsd\gamma = \kappa. \tag{2.10}$$

2.3.4 Households

The economy is populated by an infinitely-lived representative household. In each period, the household supplies labour inelastically at the exogenous amount $\overline{N} > 0$, consumes the final good, and receives dividends from operating firms. The household discounts the future with factor $\beta \in (0,1)$ and chooses consumption C to maximize lifetime utility:

$$V = \max_C \{\ln C + \beta V'\} \tag{2.11}$$

subject to

$$pC = w\overline{N} + \Pi \tag{2.12}$$

where Π is firm profits.

2.3.5 Equilibrium

A stationary equilibrium for this economy consists of a wage w, mass of entrants μ_e, an endogenous distribution of firms $\Omega(\varepsilon, \gamma, z)$, value functions, and policy functions over exit, search, sales, and hiring such that, given the wage w and mass of entrants μ_e:

1. *Firm optimization:* The value functions $\left\{J(\varepsilon,\gamma,z), \widehat{J}(s,\gamma,z), \widetilde{J}(\varepsilon,\gamma,z), S(\varepsilon,\gamma,z)\right\}$ solve (2.5)–(2.9) and $\{x(\varepsilon,\gamma,z), s(\varepsilon,\gamma,z), t(s,\varepsilon,\gamma,z), \ell(\varepsilon,\gamma,z)\}$ are the associated policy functions.

2. *Free entry:* The free entry condition (2.10) holds.

3. *Stationarity:* The endogenous distribution $\Omega(\varepsilon,\gamma,z)$ of firms is stationary.

4. *Market clearing:* The goods and labour markets clear:

$$\begin{aligned}
\overline{N} &= \mu_e \frac{\kappa}{w} + \int \left\{ \ell(\epsilon,\gamma,z) + (1-\delta)(1-x^*(\varepsilon,\gamma,z))s^*(\epsilon,\gamma,z)\frac{c_s}{w} \right\} d\Omega(\varepsilon,\gamma,z) \\
&\quad + \int \int_{f_c \le f_c^*} (1-\delta)(1-x^*(\varepsilon,\gamma,z))\frac{f_c}{w}\, dG(f_c) d\Omega(\varepsilon,\gamma,z) \\
Y &= C.
\end{aligned}$$

2.3.6 Profit Per Worker (PPW)

Before calibrating the model, it is instructive to compute PPW, as it is a key variable measured in the event study. In the model, PPW takes the following form, letting $a = \epsilon\gamma z$ be the productivity term:

$$PPW = w\frac{1-\alpha}{\alpha} - f_c\left(\frac{w}{\alpha a}\right)^{\frac{1}{1-\alpha}} = w\frac{1-\alpha}{\alpha} - f_c\frac{w}{\alpha a \ell^\alpha}$$

In our event study, we measure PPW holding labour constant at the pre-sale level. In the model, holding labor constant, PPW only changes because of fixed costs and productivity terms (a). What we evaluate in the event study is the difference between matched sold and non-sold firms. This can be expressed as:

$$\begin{aligned}
PPW_s - PPW_{ns} &= (f_{c,s} - f_{c,ns})\frac{w}{\alpha a_s \ell_s^\alpha} + f_{c,ns}\left(\frac{w}{\alpha a_s \ell_s^\alpha} - \frac{w}{\alpha a_{ns}\ell_{ns}^\alpha}\right) \\
E[PPW_s - PPW_{ns}] &= \bar{f}_c\left(\frac{w}{\alpha a_s \ell_s^\alpha} - \frac{w}{\alpha a_{ns}\ell_{ns}^\alpha}\right)
\end{aligned} \tag{2.13}$$

Equation 2.13 relies on the fact that fixed costs are i.i.d. and uncorrelated with all components of the productivity term, a. Thus, when we measure changes in the average difference of profit per worker (holding labour fixed), we are capturing exclusively changes in productivity, which can result both from changes in owner abilities as well as the idiosyncratic terms. Thus the model will inherently capture the potential for selection on idiosyncratic terms.

2.4 Calibration and Results

In this section we discuss the model calibration and investigate the role of firm sales in the benchmark model.

2.4.1 Calibration

Our goal in bringing the model to the data is to replicate the patterns on firm sales, exit, and growth as closely as possible. Before detailing our calibration procedure, we first need to specify the distributional assumptions we make. Following convention in the firm dynamics literature, we assume that idiosyncratic productivity, z, follows an AR(1) process in logs:

$$\ln z_t = \mu_z + \rho_z \ln z_{t-1} + \epsilon_{z,t},\ \ \epsilon_z \sim N(0, \sigma^2_{\epsilon_z}), \tag{2.14}$$

where $\rho_z \in (0,1)$ represents the persistence, $\epsilon_{z,t}$ the innovation, and μ_z the mean.

In addition, we assume that the firm owner-match fixed effects are jointly log-normally distributed. Specifically,

$$(\varepsilon, \gamma) \sim \mathcal{BVLN}(\mu_\varepsilon, \mu_\gamma, \sigma_\varepsilon, \sigma_\gamma, \rho_{\varepsilon\gamma}), \tag{2.15}$$

where μ denotes the mean of the logarithms of the two fixed effects, σ their standard deviation and $\rho_{\varepsilon\gamma}$ the correlation between the two.

In total, the model contains a set of 15 parameters. We set the interest rate to 4% annually and the labour share in production to 0.67. The remaining thirteen parameters are calibrated to match a set of moments from the CEEDD micro-data. The values and targets for these parameters are listed in Table 2.3 below. Due to vetting requirements, several empirical moments cannot yet be released publicly in these tables, but the model does quite well in reproducing the remaining targeted features of the data.

TABLE 2.3: Parameters determined jointly in equilibrium

Parameter		**Value**	**Target**	**Data**	**Model**
Search option	η	1.3	Sales rate	0.70	0.71
Search cost	c_s	7,500	Mean emp sold/non-sold	2.0^a	2.5
Owner/firm FE correlation	$\rho_{\varepsilon\gamma}$	-0.75	Sales by size	1.90	1.89
Persistence (z)	ρ_z	0.5	Autocorr. emp	0.95^b	0.96
SD (z)	σ_ϵ	0.01	Variance emp growth	0.3^c	0.36
Entry cost	κ	47mm	Avg PPW	45800.82	46208.38
Exit shock	δ	2.2	Avg exit	2.50	2.50
Operating cost	f_c	300000	Share neg. profit	*vet*	17.6
Signal reliability	ψ	0.7	2 yr surv	*vet*	-3.6%
Mean firm FE	μ_γ	1.55	Median emp	*vet*	3.90
SD firm FE	σ_γ	0.6	25th/75th emp	*vet*	4.40
Mean owner FE	μ_ε	0.55	Med PPW change at sale	*vet*	16600
SD owner FE	σ_ε	0.5	Med neg PP change at sale	*vet*	-63268

Note: "vet" identifies empirical moments in the table that require vetting from Statistics Canada before release. a - computed in SBO, b - Jaimovich, Terry, and Vincent, 2023, c - Statistics Canada.

Figure 2.6 presents the two event study figures and sales by firm size computed on simulated data. Note that in the calibration we target raw changes in profit at sale, the difference in survival between sold and non-sold firms, and the relative sales rates of the top- and bottom-quartiles in firm size. These are clearly informative for the three graphs below but do not directly map to them, as the event study matches similar firms by size. Qualitatively, the model is able to capture the basic dynamics of the model, undershooting the change in PPW and slightly overshooting the decline in survival. Firm sales increase in age.

2.4.2 Inspecting the Sales Mechanism

In the model, there are four key determinants of the search and sales decisions of a given firm. Owner ability (ϵ) and firm productivity (γ) determine the size of the expected gains to searching for a new owner. Search costs c_s, paid up front, imply that expected gains must be sufficiently large before search takes place. Finally, conditional on searching and receiving a signal s, uncertainty about the signal (ψ) determines whether or not such a proposed sale is accepted. Figure 2.7 presents two graphs of policy functions to gain intuition for how these four elements interact.

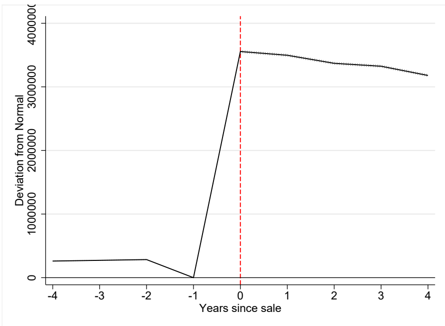

(A) Event Study: Model Profits

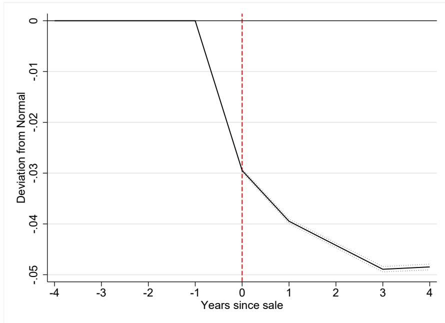

(B) Event Study: Model Survival

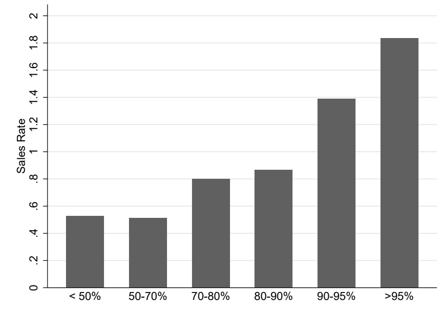

(C) Model Sales by Firm Size

FIGURE 2.6: Benchmark Model: Event Studies and Sales

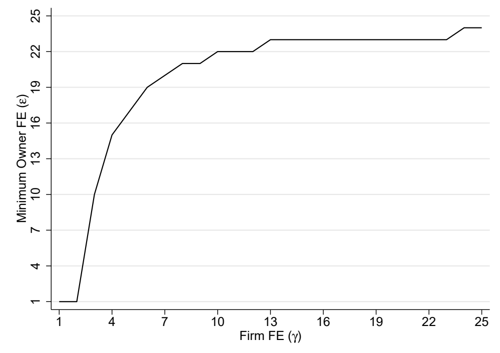

(A) Firm Productivity and Search

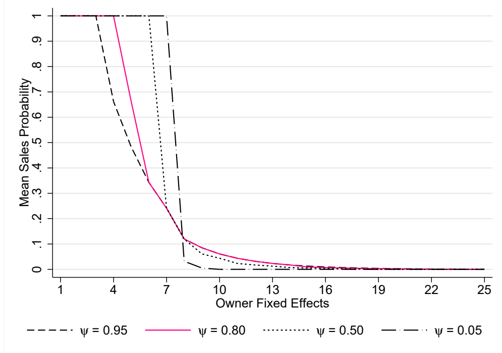

(B) Sales Probability and Uncertainty

FIGURE 2.7: Sales Mechanism: Firms, Owners and Uncertainty

Panel 2.7a of Figure 2.7 illustrates how owner ability, firm productivity and search costs interact. For each value of firm productivity, the y-axis graphs the level of owner ability above which no search occurs. Thus, below the line is the search region. For the lowest productivity firms, search never occurs. This is because, even with the worst owner, the expected value of search is simply not high enough to justify paying the search cost c_s. As firm productivity rises, so too does the expected value of searching for any level of owner ability. More productive firms search for better owners even if they are fairly high in the distribution. Of course, while the value of search rises in firm productivity, it falls in owner ability (the expected gain to finding a new owner is smaller). No firm searches when they are at the top of the owner ladder.

Panel 2.7b of Figure 2.7 illustrates the impact of uncertainty on sales probabilities (conditional on being able to search). For a fixed firm productivity (the midpoint of the distribution) we vary owner ability and compute the implied sales probability as the likelihood of accepting a sale conditional on the signal distribution. An uninformative signal ($\psi = 0.05$) means that sales are effectively a random draw from the sampling distribution. In this case, firms with owners below the average ability in the sampling distribution sell, while those above average do not, leading sales probabilities that look like a step-function (dotted line). As uncertainty is removed and the signal quality improves, we see that around the average ability level, sales become less likely for worse owners and more likely for better owners. Firms with better owners are willing to sell if they receive a high signal because the expected owner quality is in fact high and vice versa for firms with lower quality owners when they receive a low signal. Because new owners are drawn from a distribution, it will always be the case that sales probabilities are one for the lowest owner and zero for the highest, but signal precision gives the model scope to induce higher quality owners to sell in some cases.

Understanding these mechanics help provide intuition for how the model maps to the data. The panels in Figure 2.7 show that search increases in firm productivity while sales (conditional on search) decrease in owner ability. The model needs to square these forces with the fact that firm sales are increasing in firm size in the data (Figure 2.3). If the underlying equilibrium distribution implied a positive correlation between firm productivity and owner ability, these forces would be competing, as larger firms would search more often but sell less often, conditional on search. As a result, our calibration finds that initial draws must be negatively correlated, so that larger firms (with high business productivity) have greater scope to grow their owner talent through equity sales.[10]. It is

[10]One interpretation for this negative correlation is to think of the joint distribution as containing conditional ladders. From this perspective, the founders of small firms tend to have most of the skills

Table 2.4: Sold v. Non-Sold Characteristics

Moment	**Sold**	**Non-Sold**
Average TFP	536,178.90	779,678.90
Average ϵ	1.49	2.95
Average γ	6.42	4.55
Average z	60,615.19	60,504.55

worth noting that, in equilibrium, a negative correlation between owner ability and firm productivity will arise through selection, even in the absence of a negative correlation in the sampling distribution. The option value of changing owners allows good firms with bad owners to survive, pushing the equilibrium correlation negative.

Selection on Sales

With the calibrated model in hand we can return to the question of selection, raised in the empirical section. Table 2.4 reports average TFP, owner match quality (ϵ), business productivity (γ) and the idiosyncratic shock (z) for the set of sold and non-sold firms (computed annually and then averaged).

In the simulated data, sold firms have better business productivities and worse owner matches, as might be expected. On average they are less productive than non-sold firms (and therefore somewhat smaller, pre-sale). Importantly, there is no difference in the average idiosyncratic component, consistent with our empirical finding that profit and survival patterns are similar among firms that were growing or shrinking pre-sale.

2.4.3 Firm Sales in the Life Cycle

We begin our analysis by considering the impacts of sales over the firm life cycle. As noted previously, incorporating firm sales into the model affects firm dynamics in two ways. First, the option value to sell, even if unrealized, increases the continuation value of firms with low owner match quality, inducing them to survive. Second, the realized change in owner match quality for firms that are able to sell leads to improved profitability and size on average. Figure 2.8 traces out these impacts by firm age.

Figure 2.8a reports, for each age, the share of surviving firms that would not choose to continue if the option to sell were removed. To compute this, we solve a counterfactual

required to run them, while the founders of big firms tend not to (are lower down their ladder), however these owners are not directly comparable. Misallocation then would refer only to owners within a certain firm productivity rather than between productivity levels.

Figure 2.8: Survival and Dispersion Dynamics

(a) Survival from Sales Option

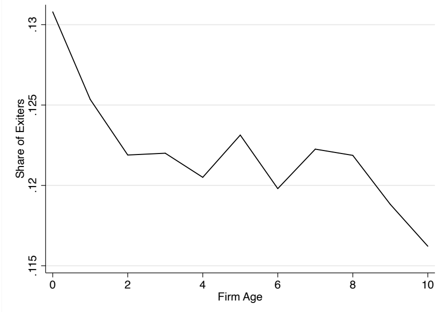

(b) Share of Dispersion from Realized Sales

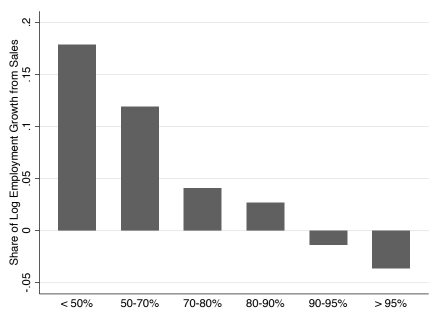

economy in which the wage is held at the benchmark value but we remove the option to sell by setting $\eta = 0$. We then extract the continuation policy function from this counterfactual economy and evaluate it among the firms in our benchmark simulation. Among entering firms, approximately 13% survive due to the option value of sale alone, with this share declining to 10% by age 30. These estimates are particularly striking given the very low likelihood of a sale occurring, something internalized by firms themselves.

Turning to the firm size distribution, Figure 2.8b reports, by firm age, the share of the standard deviation in employment size accounted for by realized firm sales. This is obtained by computing counterfactual employment at all firms based on the initial owner match quality they entered with. We see that sales rapidly gain in importance, accounting for 80% of the standard deviation in firm size after just five years.

2.4.4 Stationary Distribution

We next turn to the implications of the lifecycle dynamics analyzed above for the stationary distribution. To highlight the role of sales, we first report in Table 2.5 the aggregate impacts of sales on the economy using the counterfactual economy where we do not allow for sales (setting $\eta = 0$) while holding wages at the benchmark level. Absent the option to sell, less productive firms are more likely to exit, raising the exit rate as well as the average productivity (owner plus firm) and PPW in the economy. Average owner productivity declines, as lower ability owners are not being replaced. Conversely, average firm productivity rises – greater firm productivity has to compensate for the decline in owner productivity now that owners cannot be replaced. Perhaps most striking is the response of the firm size distribution to eliminating sales, the impact of which is entirely concentrated in the tails. Removing sales compresses the bottom of the distribution by

	Benchmark	Counterfactual
Exit	1.99	2.04
Owner Productivity	1.48	1.37
Firm Productivity	6.31	6.54
Mean PPW	44,115	53,834
Labour p25/p1	3.51	1.28
Labour p75/p25	4.67	4.79
Labour p99/p75	47.5	6.81

TABLE 2.5: Impacts from removing sales

eliminating highly productive firms with poor owner matches. In the benchmark model, these firms enter because the option value of climbing up the match ladder is very high. Absent sales, these firms tend to exit quickly. The top of the distribution is also compressed in the counterfactual, but this effect operates through realized sales. Ownership changes have significant level impacts in revenue and employment among the largest firms. Removing firm sales collapses the upper quartile of the distribution as these large firms are no longer able to improve owner matches.

We illustrate the impact of these two dimensions on the firm size distribution in Figure 2.9. We begin by dividing the benchmark firm distribution into rough quintiles. We then keep the underlying distribution but replace realized owner matches with the initial match a firm is born with - thus we undo sales while keeping intact the selection implied by the option value. This distribution (“No Realized Sales") simply shifts all mass to the left. Firms become uniformly smaller, by construction, and now more than a quarter of firms are in the smallest quintile (as defined by the benchmark distribution). In the final step, we remove both realized sales and the option value of sales by running the counterfactual reported in Table 2.5. At the top, this distribution (“No Option to Sell") is almost identical to the one obtained by removing realized sales. At the bottom however, we see that removing the option value dramatically reduces the mass of firms in the smallest quintile, shifting the distribution to the middle.

2.5 The Macroeconomic Implications of Firm Sales

In the final part of our analysis, we ask turn to the macroeconomic implications of firm sales. How does incorporating ownership changes into a model of firm dynamics change the aggregate response of the economy to policy? We answer this question in two steps. First, we remove firm sales ($\eta = 0$) and recalibrate the model to match firm size and

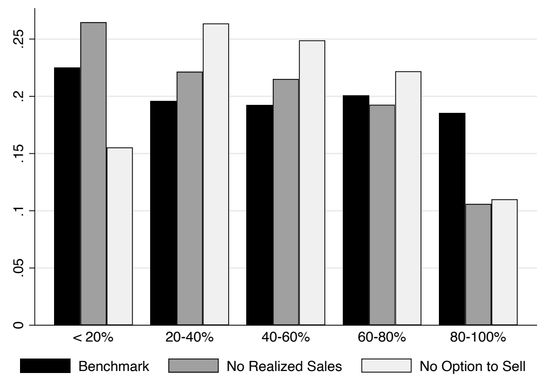

FIGURE 2.9: Counterfactual firm size distributions

growth moments. Then we impose the same policy in the calibrated models with and without firm sales and study the aggregate responses.

Moment	Sales	No Sales
Median Emp.	3.9	4.02
Emp p75/p25	4.40	4.51
Autocorr Emp	0.96	0.96
SD Log emp growth	0.36	0.36

TABLE 2.6: Calibration Moments: Sales (Benchmark), No Sales ($\eta = 0$)

The calibrated moments for the model with and without sales are reported for comparison in Table 2.6. Calibrating a structural model of firm dynamics, absent sales, absorbs an important, endogenous dimension of growth into the exogenous productivity process. To understand the impact of this for how the model responds to policy, we consider an operating subsidy to firms, paid for by taxing the household – and examine the responsiveness of the aggregate economy to such a policy change in the two calibrated models. We se the subsidy at 25% of the average fixed cost. The output is reported in Table 2.7.

Across all moments studied, the model without firm sales is more responsive to the operating subsidy than the benchmark economy. To understand these results, it is important to remember that the operating subsidy helps smaller, less productive firms remain active. Absent owner changes, the scope for growth of these firms is reduced, and thus their continuation values are more sensitive to an operating subsidy. By contrast,

Moment	Sales	No Sales
Total Output	-6.45%	-9.95%
Avg. TFP	-2.69%	-5.15%
Avg. Firm size	-7.03%	-8.63%
Exit Rate	-5.88%	-8.69%

Table 2.7: Aggregate Response to Operating Subsidy

when option value of sales is present, the continuation value of these small firms is larger, and thus less responsive to the subsidy. As a result, owner changes dampen the aggregate response of the economy to an operating subsidy.

2.6 Conclusion

In this paper we study the role of ownership changes, i.e. firm sales, for firm dynamics. Using a detailed, administrative dataset, we are able to track the ownership of firms over time and infer the sale of firms from changes in reported equity holdings. We show that firm sales are widespread. The sales rate of firms is larger than the exit rate for all employer firms, and grows in firm size. In order to study the economic significance of firm sales, we perform an event study analysis, matching sold firms with similar, non-sold firms pre-sale. We find that sales increase profitability on average but decrease survival in the first three years. We consider a number of alternative explanations and argue that the most plausible interpretation of the data is that ownership changes on average improve firm performance but can be risky – occasionally, new owners fail.

Guided by this evidence, we incorporate risky firm sales into a structural model of firm dynamics. We calibrate the model to salient features of the firm size distribution, growth dynamics and sales. We then use the model to run counterfactual experiments that help use quantify the role of firm sales. In the model, sales operate through two mechanisms. First, the option value to improve owner matches allows smaller firms to survive, even if these sales are never realized. Second, realized sales directly impact firm profitability and growth. At the firm level, we find that 13% of entrants survive exclusively due to the option value of sale. Among small firms, 18% of mean log employment growth is attributable to realized sales. Aggregating these impacts up, we find that firm sales operate primarily at the tails of the distribution, accounting for 85% of the dispersion in the top quartile. Thus sales contribute significantly to the size of the largest firms. Finally, we study the macroeconomic implications of incorporating the firm sales margin.

We compare a calibrated model without firm sales to our benchmark economy, and in both cases implement an identical subsidy to operating firms. We find that our benchmark model is substantially less responsive to these policies. This is because most of the impact of the subsidy happens among smaller firms. In the benchmark model, much of the continuation value of small firms derives from the option value of future sales, thus reducing the impact of a subsidy on total firm value and hence continuation decisions.

References

Becker, Sascha O. and Hans K. Hvide (2022). "Entrepreneur Death and Startup Performance". *Review of Finance*, 26 (1), 163–185.

Bhandari, Anmol, Paolo Martellini, and Ellen McGrattan (2022). "A Theory of Business Transfers". *University of Minnesota Working Paper.*

Brandt, Loren et al. (2022). "Serial Entrepreneurship in China". *Working Paper.*

Brüggemann, Bettina, Zachary L. Mahone, and Chris Muris (2024). "Business and Owner Contributions to TFP". *Working Paper.*

Choi, Joonkyu et al. (Forthcoming). "Early Joiners and Startup Performance". *Review of Economics and Statistics.*

Ewens, Michael and Matt Marx (2018). "Founder Replacement and Startup Performance". *The Review of Financial Studies*, 31 (4), 1532–1565.

Gaillard, Alexandre and Sumudu Kankanamge (2020). "Buying and Selling Entrepreneurial Assets". *Working Paper.*

Guntin, Rafael and Federico Kochen (2023). "Financial Frictions and the Market for Firms". *Working Paper.*

Hopenhayn, Hugo A. and Richard Rogerson (1993). "Job Turnover and Policy Evaluation: A General Equilibrium Analysis". *Journal of Political Economy*, 101 (5), 915–938.

Jaimovich, Nir, Stephen J. Terry, and Nicolas Vincent (2023). "The Empirical Distribution of Firm Dynamics and Its Macro Implications". *Working Paper.*

Jovanovic, Boyan (1982). "Selection and the Evolution of Industry". *Econometrica*, 50 (3), 649–670.

Mahone, Zachary (2023). "Business Ownership and the Secondary Market". *Working Paper*.

Pugsley, Benjamin, Petr Sedlacek, and Vincent Sterk (2021). "The Nature of Firm Growth". *American Economic Review*, 111 (2), 547–79.

Smith, Matthew et al. (2019). "Capitalists in the Twenty-First Century". *The Quarterly Journal of Economics*, 134 (4), 1675–1745.

Chapter 3

College Majors, Occupations, and Entrepreneurship

College Majors, Occupations, and Entrepreneurship

Thomas Palmer
McMaster University
palmet8@mcmaster.ca

Abstract

This paper examines the relationship between the college majors and entrepreneurship in the United States. Using data from the 1997 National Longitudinal Survey of Youth, I show that STEM graduates are systematically more likely to pursue STEM-related employment but less likely to pursue entrepreneurship than non-STEM graduates. To quantify the importance of major selection for entrepreneurship, I develop a model that links college major selection to post-graduation occupational outcomes and entrepreneurship decisions. The model, calibrated to the NLSY97 micro-data, successfully replicates the observed patterns in major selection, occupational sorting, and entrepreneurship rates. Counterfactual experiments reveal that reducing STEM tuition by 50 percent would nearly double the share of STEM majors, but decrease overall entrepreneurial activity. Conversely, reducing barriers to entrepreneurship increases entrepreneurial activity without any corresponding impact on STEM enrolment. These findings provide valuable insights into how educational trends influence entrepreneurial opportunities and raise novel implications for policies related to education and entrepreneurship support.

Keywords: *College Majors, Entrepreneurship, Self-Employment, Business Ownership, STEM.*

3.1 Introduction

The United States has experienced a dramatic shift toward Science, Technology, Engineering, and Mathematics (STEM) fields in recent decades. According to the National Center for Education Statistics, the share of STEM Bachelor's Degrees conferred by post-secondary institutions increased from 24.83 percent in 1980 to 34.72 percent in 2021—a substantial increase of nearly 10 percentage points, or 40 percent. The large rise in STEM majors has been attributed to a variety of factors, including technological change, evolving labour market demands, and targeted education policies.[1]

While the drivers of this shift have been well studied, their broader implications remain unclear. This paper aims to fill this gap by investigating the relationship between STEM education and aggregate entrepreneurship among recent college graduates. As policymakers and educators continue to emphasize STEM education, understanding its implications for career choices and labour market dynamics becomes increasingly important. Indeed, STEM graduates are often associated with innovation and technological progress, but their propensity for entrepreneurship—a key driver of economic growth—remains understudied. Accordingly, this paper has two broad objectives. First, at a micro-level, I examine how majoring in STEM fields influences the likelihood of engaging in entrepreneurship in the early stages of one's career. Second, at a macro-level, I analyze how these individual-level choices impact the aggregate composition of occupations and entrepreneurship.

To achieve these objectives, I proceed in three parts. First, I use data from the 1997 National Longitudinal Survey of Youth (NLSY97) to empirically assess the individual-level impact of majoring in STEM on post-graduate career outcomes. I find that college majors strongly predict post-graduate occupations: STEM majors are 53 percentage points more likely to work in STEM-related occupations after graduation. Moreover, I show that occupations vary significantly in rates of entrepreneurship: individuals employed in STEM-related occupations are 8 percentage points less likely to pursue entrepreneurship following graduation.

In the second part of the paper, I build a quantitative model that explicitly links college major choices to post-graduation career outcomes, including both occupations and entrepreneurship. The model is designed to replicate the micro-level facts from the NLSY97 and provide the necessary structure to analyze their macroeconomic implications through counterfactual simulations. The key components of the model include an educational choice, occupational sorting, and an entrepreneurship decision. Individuals choose between STEM and non-STEM majors based on their initial wealth, modelled as a binary distribution of low and high wealth; preference shocks; and major-specific tuition costs. Upon graduating, individuals are sorted into STEM and non-STEM occupations.

[1]For example, Deming and Noray (2020) study the impact of technological change on changing skill demands, with an emphasis on STEM-related careers. Liu, Sun, and Winters (2019) show that the proportion of STEM majors increased during the Great Recession. Bottia et al. (2018) show that students entering colleges from high-schools with dedicated STEM programs are more likely to major in STEM in college.

The probability of entering a STEM (non-STEM) occupation is higher for STEM (non-STEM) majors, which captures the alignment between field of study and career paths documented in the NLSY97. Finally, individuals decide between paid employment and entrepreneurship based on their current occupation, individual entrepreneurial talent, and occupation-specific entrepreneurial thresholds. Overall, the model captures the key mechanism uncovered from the empirical analysis: majors lead to certain occupations, which in turn differ in their propensities for entrepreneurship.

In the final part of the paper, I calibrate the model to match salient moments related to majors, occupations, entrepreneurship, and wealth from the NLSY97 data. Then, I use the calibrated model as a laboratory to quantify the impact of STEM education on aggregate entrepreneurship rates through two counterfactual policy experiments. First, I consider education financing policies designed to encourage participation in STEM education. I find that lowering STEM tuition by 50 percent increases the share of STEM majors in the economy from 24.93 percent to 44.72 percent; interestingly, however, the increasing share of STEM majors coincides with a reduction in overall entrepreneurship from 9.65 percent to 9.06 percent. Second, I assess the impact of entrepreneurship support programs, which I model as lowering entrepreneurship thresholds for all occupations. I find that reducing these thresholds by 50 percent more than doubles the overall share of entrepreneurs, but has no corresponding effect on the distribution of majors or employees across STEM and non-STEM fields.

The remainder of the paper is organized as follows. In Section 3.2, I discuss the related literature on college majors and entrepreneurship. In Section 3.3, I describe the data in detail and present the motivating facts. In Section 3.4, I describe the model and define the equilibrium concept. In Section 3.5, I discuss the calibration strategy and report the parameter estimates. In Section 3.6, I perform the quantitative analysis and discuss its results and implications. Section 3.7 concludes the paper.

3.2 Related Literature

This paper contributes to several strands of the literature. By examining the relationship between college majors, occupations, and entrepreneurship, it bridges gaps between these areas of research and provides novel insights into their connections.

First, this paper relates to the extensive empirical literature on college major choice and its labour market implications. Seminal work by Arcidiacono (2004) and Wiswall and Zafar (2015) has established that students tend to choose college majors based on expected future earnings, individual abilities, and preferences. More recently, Kirkeboen, Leuven, and Mogstad (2016) use a regression discontinuity design to show that different fields of study generate substantially different payoffs, with STEM fields generally offering higher returns. Deming and Noray (2020) show that the earnings premium for STEM majors declines faster over time compare to other majors, which they attribute to rapid technological change. Within this literature, the most closely related paper is Huang (2023), which examines the macroeconomic implications of financial aid policies for college major selection. Using the same NLSY97 micro-data, Huang (2023) shows that students from low parental income backgrounds are more likely to choose STEM, health, and education majors, which feature relatively higher initial earnings, lower earnings growth, and lower earnings risk. I extend this line of research by considering post-graduation employment opportunities—and, in particular, entrepreneurship—as an additional motivation and outcome of major selection.

Second, this paper contributes to the growing body of work on the determinants of entrepreneurship. While much of this literature has focused on factors such as personal characteristics (Levine and Rubinstein, 2017; Poschke, 2013; Vilalta-Bufi, Kucel, and Giusti, 2018), access to capital (Evans and Jovanovic, 1989; Quadrini, 2000; Cagetti and De Nardi, 2006), and regulation (Braunerhjelm, Desai, and Eklund, 2015; Kong and Qin, 2021), relatively less attention has been paid to the role of specific educational pathways. I thus complement this work by evaluation how specialization in STEM fields specifically affects the propensity for entrepreneurship.

Third, this paper contributes to the large quantitative literature attempting to explain the decline in entrepreneurship in the United States since the 1980s. In particular, this paper is most closely related to the recent work of Salgado (2020), Jiang and Sohail (2023), and Kozeniauskas (2022), which all emphasize the impact of technological change in driving the decline in entrepreneurship, particularly among college graduates. Salgado (2020) shows how skill-biased technical change in the spirit of Krusell et al. (2000) can account for a significant share of the fall in entrepreneurship and firm creation since the mid-1980s. Jiang and Sohail (2023) emphasize that new technologies have raised the opportunities costs of entrepreneurship, especially for new graduates, by altering the relative wage structure of college versus non-college workers. Kozeniauskas (2022) argues that, while skill-biased technical change can explain a large share of the relative decline in entrepreneurship among college graduates, rising entry costs are the key driving factor. I complement this literature by focusing on the impact of college major composition, particularly the increase in STEM graduates, on entrepreneurship. While I do not directly examine its role in explaining the historical decline in entrepreneurship, it is worth noting that the U.S. has experienced a dramatic increase in the share of STEM graduates concurrent with the fall in entrepreneurship rates.

Finally, this paper contributes to the literature on occupational and industry-level variation in entrepreneurship. Many recent studies have shown that entrepreneurial propensity varies widely across sectors and is influenced by factors such as workplace characteristics. For example, Delgado, Porter, and Stern (2010) documents significant variation in business start-up rates across industries, with sectors like construction and professional services having higher rates of entrepreneurship relative to manufacturing or wholesale. Elfenbein, Hamilton, and Zenger (2010) show that employees of small firms are more likely to become entrepreneurs than those of large firms. In the context of STEM fields, Braguinsky, Klepper, and Ohyama (2012) and Hsu, Roberts, and Eesley (2007) have examined how technical education affects entrepreneurship in science and engineering. I add to this literature by explicitly linking college major selection to entrepreneurship rates, offering additional insights into how educational pathways contribute to occupation-level variation in entrepreneurial activity.

3.3 Data

For the empirical analysis of this paper, I leverage data from the 1997 cohort of the National Longitudinal Survey of Youth (NLSY97). This dataset, maintained by the U.S. Bureau of Labour Statistics, offers a unique opportunity to examine the relationship between majors, occupations, and entrepreneurship among a recent generation of college graduates. Below, I describe the dataset in more detail, outline its unique benefits for this paper, and detail how I construct the relevant variables for the empirical analysis.

3.3.1 National Longitudinal Survey of Youth 1997 (NLSY97)

The 1997 National Longitudinal Survey of Youth (NLSY97) is a panel survey that follows a cohort of 8,984 youths who were born between 1980 and 1984. The first interview took place in 1997, with annual interviews conducted through 2011 and biennial interviews thereafter. The NLSY97 contains comprehensive information on education and labour market outcomes for a nationally-representative sample of youth. For the main analysis, I restrict attention to the subset of 1,560 individuals in the core sample who have graduated from a four-year college at any point during the survey.

The NLSY97 offers several key strengths for my analysis. First, it contains rich information on individuals' educational histories. In particular, the NLSY97 collects post-secondary data at the college, term, and year level and contains supplementary data pulled from college transcripts. This level of granularity allows for a nuanced examination of how specific educational pathways relate to subsequent entrepreneurial activities, including changes in college majors over the course of degrees.

Second, the NLSY97 includes comprehensive information on the sources and value of education financing. Specifically, the survey asks questions about traditional financing options, such as merit-based scholarships and needs-based loans, as well as non-standard sources of financing, such as family transfers and loans. As discussed by Abbott et al. (2019), the substitution among public and private sources of education financing is particularly important to consider when analyzing student loan programs. This detailed financial information enables a more thorough investigation of how different financing strategies might influence both educational choices and subsequent entrepreneurial decisions.

Third, the NLSY97 provides recent data, allowing for an examination of how modern technologies have altered education and career choices. This cohort came of age during a period of rapid technological advancement and digitalization, which has significantly impacted educational pathways and the job market. The data from NLSY97 thus offers insights into how these technological changes have influenced the relationship between college majors, career paths, and entrepreneurship in the contemporary context. This recency allows for a more up-to-date analysis of the factors shaping entrepreneurial activities among younger generations, capturing the effects of technological innovations on both educational decisions and subsequent career outcomes.

Finally, the survey is longitudinal in nature. As such, the NLSY97 offers substantial life-cycle coverage, following respondents from their teenage years into their late 30s. This allows for the observation of educational choices, initial career paths, and early entrepreneurial ventures, providing value insights into the early-to-mid career impacts of major selection on entrepreneurship.

3.3.2 Main Variables

In this section, I describe the main variables of interest, which include college majors, occupations, and entrepreneurship.

3.3.2.1 College Majors

The NLSY97 provides rich information to infer individuals' college majors. Since I am ultimately interested in assessing the impact of college majors on entrepreneurship, I restrict attention to respondents who have completed a four-year undergraduate degree, such as a Bachelor of Arts or Bachelor of Science, at some point during the survey and provide valid college major information. This approach simplifies the process of inferring majors by concentrating on a more homogeneous group of degree holders.

The NLSY97 asks individuals about their college experience at the college, semester, and academic year frequency. Accordingly, it is possible for survey respondents to report several different majors at the same college within the same academic year. To make this analysis tractable, I proceed in three steps. First, I aggregate the data to the college-year level by taking the mode across semesters. Second, I aggregate the data to the year level by taking the mode across colleges. Finally, I determine the overall major by taking the mode across years.

After assigning each respondent to a single college major in both datasets, I standardize the set of majors into a comparable grouping of STEM and non-STEM fields. The distribution of majors is reported in Table 3.1 below.

3.3.2.2 Occupations

To classify individuals by their post-graduate occupations, I follow a procedure similar to the one used for inferring majors. Specifically, I select the mode of the occupation associated with an individual's main job across all years of the survey for each individual. Aggregating occupations enables consistency with college majors and measures of business ownership, which are only determined at one point in time. In addition, using the modal occupation provides a stable representation of each respondent's primary occupation over time.

I then group these occupations into two broad categories: "STEM" and "non-STEM." This binary classification is designed to be consistent with the structure of the quantitative model developed in the next section. Table 3.2 reports the list of occupations contained within each group, along with the shares.

Table 3.1: Classification of Majors

Category	**Major**
STEM	Biology
(25.06%)	Computer Science
	Engineering
	Mathematics
	Physics
	Health
Non-STEM	Business
(74.94%)	Humanities
	Social Sciences
	Agriculture
	Pre-Law
	Education
	Communications

Table 3.2: Classification of Occupations

Category	**Occupation**
STEM	Computer and Mathematics
(25.93%)	Architecture and Engineering
	Life, Physical, and Social Sciences
	Healthcare
Non-STEM	Management
(74.07%)	Business and Finance
	Community and Social Services
	Legal
	Education
	Arts and Entertainment
	Healthcare Support

3.3.2.3 Entrepreneurship

Unlike many other survey data sets, such as the Panel Study of Entrepreneurial Dynamics (PSED) and Survey of Consumer Finances (SCF), the NLSY97 is not specifically designed to study patterns of entrepreneurship. Fortunately, however, it contains sufficient information to construct approximate measures of entrepreneurship from questions about self-employment and business ownership. Accordingly, I employ 4 separate definitions of entrepreneurship, with varying levels of strictness.

The first measure of entrepreneurship is self-employment. Under this definition, I classify an individual as an entrepreneur if they ever report being self-employed at any point during the survey period. Like occupations, I focus the individual's main job to determine self-employment.

The second measure of entrepreneurship is business ownership. In the NLSY97, information on business ownership is relatively more limited. However, at the ages of 25, 30, and 35, respondents are asked supplementary questions about their assets, which includes a question about business ownership. Under this definition in the NLSY97, then, I define entrepreneurs as individuals who ever report owning a business at ages 25, 30, or 35. This definition focuses on individuals who have taken the step of establishing a formal business entity, which may or may not be their primary source of income.

The third and fourth measures of entrepreneurship consists of the union and intersection of the previous two definitions. The union of self-employment and business ownership is the broadest measure, as it includes self-employed non-business owners, non-self-employed business owners, and self-employed business owners. On the other hand, the intersection of self-employment and business ownership is the strictest measure and identifies individuals who are deeply engaged in entrepreneurship as both an active worker and owner of their firm.

Ultimately, the use of multiple definitions of entrepreneurship allows for a nuanced examination of entrepreneurial activities, recognizing that entrepreneurship can take various forms and degrees of formality. It also provides the necessary structure to analyze whether certain college majors are more strongly associated with particular types of entrepreneurial activities (e.g., self-employment versus formal business ownership).

3.3.3 Stylized Facts

In this section, I use the set of constructed variables to establish a set of stylized facts, which relate college majors to occupations and entrepreneurship. The stylized facts provide the micro-foundations for the quantitative model developed in the following section.

Fact 1: College majors strongly predict post-graduate occupational choices.

The first key finding is that college majors are strong predictors of subsequent occupational choices. To formally quantify this relationship, I estimate a series of logistic regressions using data from the NLSY97. The regression model is specified as follows:

$$\Pr[\text{Occupation}_i = j|\text{Major}_i, X_i] = \frac{\exp\left(\beta_j'\text{Major}_i + \gamma_j' X_i\right)}{\sum_{k=1}^{J-1}\exp\left(\beta_k'\text{Major}_i + \gamma_k' X_i\right)}. \tag{3.1}$$

Here, Occupation_i represents individual i's modal occupation, Major_i denote their college major, and X_i is a vector of control variable including gender, race, average career earnings, an indicator for high-parental income, and test scores from the Armed Services Vocational Aptitude Battery (ASVAB). For average career earnings and test scores, I transform the raw variables into quartiles. Table 3.4 below reports the log-odds ratios (top panel) and average marginal effects (bottom panel) from these regressions. Column (1) excludes the vector X_i, while Column (2) includes it.

The results consistently show that individuals who graduate with a STEM major are significantly more likely to end up in STEM-related occupations upon graduating. More specifically,

1. Without controls (Column 1), STEM majors have a 59 percentage point greater likelihood of entering a STEM occupation compared to their non-STEM major counterparts.
2. With controls (Column 2), this effect remains strong at 53 percentage points.

These findings underscore the strong link between college major choice and subsequent career paths, particularly for STEM fields. The persistence of this relationship even after controlling for various individual characteristics suggests that the skills and knowledge

Table 3.3: Distribution of Entrepreneurs

Entrepreneur Definition	**Count**	**Share**
Self-Employed	343	22.06
Business Owners	264	21.06
Self-Employed or Business Owners	463	35.70
Self-Employed and Business Owners	144	9.52

Table 3.4: Occupational Choice Probabilities

Pr[STEM Occupation]	**(1)**	**(2)**
Log Odds Ratios		
STEM Major	3.17***	3.24***
	(0.24)	(0.29)
Average Marginal Effects		
STEM Major	0.59***	0.53***
	(0.04)	(0.04)
Additional Controls	–	✓
N	594	517
Pseudo R^2	0.327	0.41

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

acquired in STEM majors are highly-valued and directly applicable in STEM occupations.

Fact 2: STEM workers are less likely to become entrepreneurs.

The second fact connects occupations to entrepreneurship. Specifically, I find that individuals working in STEM-related (non-STEM) occupations are less (more) likely to pursue entrepreneurship after graduation. To reach this conclusion, I estimate the probability of engaging in various forms of entrepreneurship through a series of logistic regressions. Formally, the regressions are specified as follows:

$$\Pr[\text{Entrepreneur}_i = k | \text{Occupation}_i, X_i] = \frac{\exp\left(\delta_k' \text{Occupation}_i + \theta_k' X_i\right)}{1 + \exp\left(\delta_k' \text{Occupation}_i + \theta_k' X_i\right)}. \tag{3.2}$$

Here, the dependent variable Entrepreneur_i represents the various measures of entrepreneurship for individual i: self-employment, business ownership, either, or both. The main independent variable of interest is Occupation_i, which represents individual i's modal occupation. The vector X_i is a set of individual-level controls, including gender, race, test scores, parental income, and average career earnings.

Tables 3.5, 3.6, 3.7, and 3.8 below report the log-odds ratios (top panels) and average marginal effects (bottom panels) from these regressions for the NLSY97 cohort. In all cases, Column (1) represents the regression without the control vector X_i, while Column (2) includes the additional controls.

The results here consistently show that individuals in STEM-related occupations are less likely to engage in entrepreneurial activities. In particular, individuals in STEM occupations are associated with a 11 percentage point decrease in the probability of self-employment without controls, and a 17 percentage point decrease with controls (Table 3.5). Individuals in STEM occupations are associated with an 8 percentage point decrease in the probability of business ownership without controls, and an 11 percentage point decrease with controls (Table 3.6). Individuals in STEM occupations are associated with an 12 percentage point decrease in the probability of *either* self-employment *or* business ownership without controls, and a 17 percentage point decrease with controls (Table 3.7). Individuals in STEM occupations are associated with an 8 percentage point decrease in the probability of *both* forms of entrepreneurship without controls, and an 11 percentage point decrease with controls (Table 3.8).

Table 3.5: Probability of Self-Employment

Pr[Self-Employment]	**(1)**	**(2)**
Log Odds Ratios		
STEM Occupation	-0.67^{***}	-1.11^{***}
	(0.26)	(0.32)
Average Marginal Effects		
STEM Occupation	-0.11^{***}	-0.17^{***}
	(0.04)	(0.05)
Additional Controls	–	✓
N	594	517
Pseudo R^2	0.01	0.06

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.6: Probability of Business Ownership

Pr[Business Ownership]	**(1)**	**(2)**
Log Odds Ratios		
STEM Occupation	-0.42^{*}	-0.67^{**}
	(0.26)	(0.29)
Average Marginal Effects		
STEM Occupation	-0.08^{*}	-0.11^{**}
	(0.05)	(0.05)
Additional Controls	–	✓
N	499	444
Pseudo R^2	0.01	0.06

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.7: Probability of Self-Employment or Business Ownership

Pr[Self-Employment or Business Ownership]	**(1)**	**(2)**
Log Odds Ratios		
STEM Occupation	-0.51^{**}	-0.83^{***}
	(0.22)	(0.26)
Average Marginal Effects		
STEM Occupation	-0.12^{**}	-0.17^{***}
	(0.05)	(0.05)
Additional Controls	–	✓
N	512	453
Pseudo R^2	0.01	0.07

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.8: Probability of Self-Employment and Business Ownership

Pr[Self-Employment and Business Ownership]	**(1)**	**(2)**
Log Odds Ratios		
STEM Occupation	-0.89^{**}	-1.25^{***}
	(0.39)	(0.44)
Average Marginal Effects		
STEM Occupation	-0.08^{***}	-0.11^{***}
	(0.04)	(0.04)
Additional Controls	–	✓
N	581	508
Pseudo R^2	0.02	0.07

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

These findings suggest that the occupational environment for STEM graduates may be less conducive to entrepreneurial ventures, especially given the persistence of such effects after controlling for various individuals characteristics. Several potential explanations for this phenomenon include the higher opportunity costs associated with STEM occupation, as such occupations often offer higher salaries and more stable career paths; specialized skills, as the training received through STEM education may be more valuable to specific companies or jobs; or capital requirements, as STEM-related entrepreneurial ventures may require more capital, creating larger barrier for entry. Although my paper does not attempt to explain the underlying rationale, the finding is nevertheless consistent with recent studies of major choice such as Huang (2023).

3.3.4 Implications of the Stylized Facts

Together, the two stylized facts presented above highlight a crucial link between educational choices, career paths, and entrepreneurial outcomes among recent graduates. Ultimately, these findings have important implications for understanding the dynamics of entrepreneurship in the modern economy and motivate the development of a structural model to quantify their aggregate impacts.

More specifically, Fact 1 demonstrates a strong association between college majors and subsequent occupational choices. This suggests that the skills and knowledge acquired during college significantly shape individuals' career paths. The high probability of STEM majors entering STEM occupations on the order of 53 percentage points indicates a substantial degree of educational path dependence in career outcomes. In addition, Fact 2 reveals that STEM occupations are associated with lower rates of entrepreneurship across various measures. The consistent negative effects, ranging from 8 to 17 percentage points, suggest that the occupational environment of STEM fields may be less suited to entrepreneurial ventures.

Combining these facts suggests an indirect pathway through which educational choices influence entrepreneurship. Majoring in STEM fields encourages employment in STEM-related occupations, which in turn are associated with lower rates of entrepreneurship. This mechanism implies that the increasing emphasis on STEM education in recent decades may in fact be contributing to a decline in aggregate entrepreneurship. Accordingly, these findings generate important implications for education and economic policy. While promoting STEM education may drive innovation and productivity in established firms, it may simultaneously suppress entrepreneurship.

To formally quantify the impact of these relationships on overall entrepreneurship rates and to explore potential policy implications, a structural model is necessary. Such a model should capture the strong link between college major choices and subsequent occupations; incorporate the varying propensities for entrepreneurship across different occupations; account for individual characteristics and external factors that influence both education and career choices; and allow for the simulation of counterfactuals to assess the impact of changes in educational choices or policy interventions.

3.4 Model

In this section, I build a quantitative economic model designed to replicate the micro-level patterns documented in the NLSY data from Section 3.3 and study their aggregate implications. The model extends existing theories of occupational choice by introducing an explicit link between college majors and post-graduate career outcomes. By emphasizing the role of STEM education in shaping occupational selection and entrepreneurial activity, the model sheds new light into how early educational decisions influence long-term macroeconomic outcomes—such as entrepreneurship—and offers novel implications for policies related to education financing.

3.4.1 Environment

I consider a model with three distinct periods: education, employment, and consumption. The economy is populated by a large number of individuals who are heterogeneous in their initial wealth.

In the education phase, individuals select their college major, choosing between STEM and non-STEM as possible options. The college major decision depends on an individuals' initial wealth, tuition costs, and the expected utility associated with each major. I model the decision as a discrete choice problem with Extreme Value Type-1 shocks, which converts the college major policy function into a continuous probability.

During the employment phase, individuals are assigned to one of two occupations: STEM or non-STEM. The probability of entering each occupation is conditional on the chosen major, reflecting the impact of educational background on career outcomes, and exogenously determined. After occupation assignment, all individuals face a choice between pursuing paid employment or entrepreneurship. The decision to become an entrepreneur is made by comparing an individual's expected utility from entrepreneurship to the utility from paid employment, subject to occupation-specific entrepreneurial thresholds. The expected utility from entrepreneurship, in turn, depends on the individuals' realized entrepreneurial talent, which is distributed according to a known distribution.

Finally, during the consumption phase, individuals consume their entire income to maximize utility. An individual's income is determined by their wage income, tuition, and initial wealth. For paid employees, wages differ according to the individual's occupation. For entrepreneurs, income is directly proportional to their realized entrepreneurial talent and scaled by an occupation-specific factor. The model timeline is depicted graphically in Figure 3.1 below.

Ultimately, the model captures the key interactions between educational choices, occupational outcomes, and entrepreneurship decisions. It reflects how initial wealth influences educational decisions, how education affects occupational probabilities, and how the combination of education, occupation, and entrepreneurial talent determines income and, ultimately, consumption and utility.

FIGURE 3.1: Model Timeline

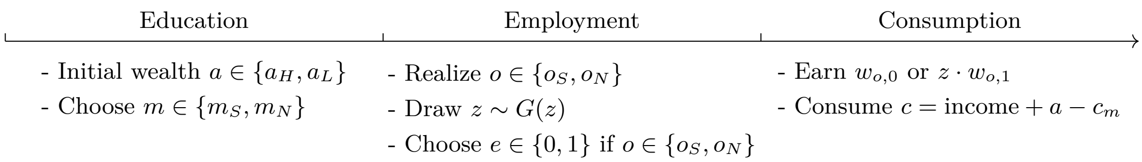

3.4.2 Education

The economy consists of a large number N of individuals who differ ex-ante in their initial wealth, a. I assume that initial wealth a can take one of two possible values: $a \in \{a_H, a_L\}$ with $a_H > a_L \geq 0$. The probability of wealth is given by a vector $p_{\text{wealth}} = [p_L, p_H]$, corresponding to low and high wealth, respectively.

Given initial wealth a, an individual begins their life by choosing their college major, m. I assume that there are two possible college majors: $m \in \{m_S, m_N\}$, where m_S represents STEM and m_N represents non-STEM. The majors are differentiated by their tuition rates c_m. For simplicity, I assume that the non-STEM major carries no cost, while the tuition associated with STEM majors is $c_S > 0$.

Individuals decide on their major by weighing the cost of education, represented by the tuition rate, against the expected utility from future earnings. I model the major selection process as a discrete choice problem with Extreme Value Type-1 shocks. In particular, the utility an individual with initial wealth a derives from choosing major $j \in \{S, N\}$ is given by:

$$U_j(a) = V_j(a) + \varepsilon_j,$$

where $V_j(a)$ is the deterministic component of utility and ε_j is an idiosyncratic preference shock following the Extreme Value Type-1 distribution. The deterministic utilities are defined as:

$$\begin{aligned} V_S(a) &= u(w_S + a - c_S) \\ V_N(a) &= u(w_N + a), \end{aligned}$$

where w_S and w_N represent the expected wages for STEM and non-STEM majors, respectively. Given the properties of the Extreme Value Type-1 shock, the probability of choosing a STEM major takes the following form:

$$P(m_S|a) = \frac{\exp(\zeta V_S(a))}{\exp(\zeta V_S(a)) + \exp(\zeta V_N(a))}, \tag{3.3}$$

where ζ is a scale parameter that governs the sensitivity of choices to differences in deterministic utility. A higher ζ implies that choices are more responsive to these differences, while a lower ζ suggests that choices are more random. This formulation captures how initial wealth, tuition costs, and expected future earnings influence the choice of majors, while also accounting for unobserved factors through the idiosyncratic shocks. The ζ parameter determines the relative importance of observed factors versus unobserved preferences in the decision-making process. The expected utility from consumption, $\mathbb{E}[u(c)]$, is implicitly considered in this formulation through the expected wages, w_S and w_N. Of course, expected wages will, in turn, depend on subsequent occupational choices, which are described next.

3.4.3 Employment

Following education, individuals enter the employment phase in one of two occupations: $o \in \{o_S, o_N\}$, representing STEM and non-STEM occupations, respectively. The probability of entering a specific occupation depends on the individuals' chosen majors, that is,

$$p_{o|m} \equiv \Pr[O = o|m], \tag{3.4}$$

where O is the random variable representing occupational choice.

Following the determination of occupations through these probabilities, all individuals receive a draw of entrepreneurial talent z from a known distribution $G(z)$. Then, given occupational assignments and entrepreneurial talent, individuals decide between paid employment and entrepreneurship. An individual chooses to become an entrepreneur if their utility from entrepreneurship exceeds their utility from paid employment and their entrepreneurial talent exceeds an occupation-specific threshold:

$$x = \begin{cases} 1 & \text{if } u(z \cdot w_{o,1} + a') > u(w_{o,0} + a') \text{ and } z > z_o^* \\ 0 & \text{otherwise} \end{cases} \tag{3.5}$$

where x is a binary variable equal to 1 for entrepreneurs and 0 otherwise; z is the individual's entrepreneurial talent drawn from a known distribution; $w_{o,1}$ is the wage multiplier for entrepreneurs in occupation o; $w_{o,0}$ is the wage for paid employment in occupation o; a' is the individual's wealth after paying for education; $u(\cdot)$ is the utility function; and z_o^* is an occupation-specific entrepreneurial talent threshold.

This formulation captures how educational backgrounds shape subsequent occupational outcomes, and how the combination of occupation, entrepreneurial talent, and initial wealth determines the choice between paid employment and entrepreneurship. Figure 3.2 below depicts the various linkages between initial wealth, majors, occupations, and job types permitted within the model.

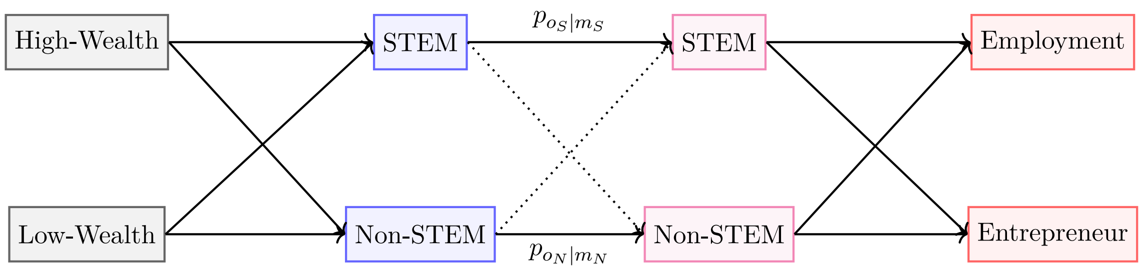

FIGURE 3.2: Model Structure

3.4.4 Consumption

All individuals consume the final good in the consumption phase after realizing their income, which is either entrepreneurial income z for entrepreneurs or wages $w_{m,o}$ for paid employees. Given periodic utility function $u(\cdot)$, the utility maximization problem for an individual with major m and occupation o can be stated as follows:

$$\max_{c} \; u(c) \tag{3.6}$$

subject to

$$c = \begin{cases} z \cdot w_{o,1} + a' & \text{if entrepreneur} \\ w_{o,0} + a' & \text{if paid employee.} \end{cases}$$

where c is consumption; z is entrepreneurial talent; $w_{o,1}$ is the wage multiplier for entrepreneurs in occupation o; $w_{o,0}$ is the wage for paid employment in occupation o; and a' is the individual's wealth after paying for education ($a' = a - c_S$ for STEM majors, $a' = a$ for non-STEM majors).

3.4.5 Equilibrium

An *equilibrium* in this model consists of a set of major choice probabilities $p(m_S|a)$ for each wealth level $a \in \{a_L, a_H\}$ and major $m \in \{m_S, m_N\}$; a set of occupational choice probabilities $p(o|m)$ for each major m and occupation $o \in \{o_S, o_N\}$; a set of entrepreneurship decision rules $x(z, o)$ for each level of entrepreneurial talent z and occupation o; a set of wages $w_{o,j}$ for each occupation o and job type $j \in 0, 1$ (where 0 represents paid employment and 1 represents entrepreneurship); and a distribution of entrepreneurial talent $G(z)$, such that:

1. The probability of majoring in STEM $P(m_S|a)$ is determined by:

$$P(m_S|a) = \frac{\exp(\zeta V_S(a))}{\exp(\zeta V_S(a)) + \exp(\zeta V_N(a))}$$

   where $V_S(a)$ and $V_N(a)$ are the deterministic utilities of choosing STEM and non-STEM majors, respectively.

2. The occupation choice probabilities $p(o|m)$ are given exogenously and satisfy:

$$\Pr(o_S|m_S) = p_{o_S|m_S}$$
$$\Pr(o_N|m_N) = p_{o_N|m_N}.$$

3. The entrepreneurship decision rules $x(z, o)$ are determined by:

$$x(z,o) = \begin{cases} 1 & \text{if } u(z \cdot w_{o,1} + a') > u(w_{o,0} + a') \text{ and } z > z_o^* \\ 0 & \text{otherwise.} \end{cases}$$

   where z_o^* is the occupation-specific entrepreneurial talent threshold.

4. Individual consumption is determined by:

$$c = \begin{cases} z \cdot w_{o,1} + a' & \text{if } x(z,o) = 1 \\ w_{o,0} + a' & \text{if } x(z,o) = 0 \end{cases}$$

5. The distribution of outcomes (college major choices, occupational choices, entrepreneurship decisions, incomes, and utilities) in the simulated population is consistent with the above rules and the exogenous distributions of initial wealth and entrepreneurial talent.

3.5 Calibration

To quantitatively assess the impact of college majors on entrepreneurship, I first calibrate the model to salient moments of the U.S. economy from the NLSY97 micro-data. The model features several parameters, which I either exogenously fix from public sources or internally calibrate to match moments. Note that, in taking the model to the data, I use the most conservative measure of entrepreneurship, that is, both self-employed and business owners.

3.5.1 Functional Forms

I assume that the periodic utility function is of the Constant Relative Risk Aversion (CRRA) type:

$$u(c) = \begin{cases} \frac{c^{1-\sigma}-1}{1-\sigma} & \text{if } \eta \geq 0 \text{ and } \eta \neq 1 \\ \ln(c) & \text{if } \eta = 1. \end{cases} \tag{3.7}$$

The model also features two distributions: the distribution of education taste shocks, ε, and the distribution of entrepreneurial talent, z. As mentioned previously, I assume that the taste shocks, ε, are drawn from an Extreme Value Type-1 distribution with scale parameter ζ. In addition, I assume that the distribution of entrepreneurial talent,

$G(z)$ is Pareto, that is,

$$z \sim \text{Pareto}(\alpha_z), \tag{3.8}$$

where $\alpha_z > 0$ is the shape parameter.

3.5.2 Externally-Calibrated Parameters

Table 3.9 below lists the parameters determined outside of the model.

TABLE 3.9: Parameters determined outside of the model

	Parameter	**Value**	**Source**
$w_{S,0}$	Wage (STEM worker)	22.49	NLSY97
$w_{S,1}$	Wage (STEM entrepreneur)	24.12	NLSY97
$w_{N,0}$	Wage (Non-STEM worker)	19.03	NLSY97
$w_{N,1}$	Wage (Non-STEM entrepreneur)	20.82	NLSY97
p_L	Share Low Wealth	54.23	NLSY97
p_H	Share High Wealth	45.77	NLSY97
c_N	Tuition (Non-College)	0.00	Normalization
σ	Risk aversion	2.00	Standard

As indicated, I exogenously fix wages outside of the model. Note that the NLSY97 implies both a STEM occupation and entrepreneurial premium on earnings: that is, individuals in STEM-related occupations earn more on average than individuals in non-STEM occupations; and, entrepreneurs tend to earn more than workers, on average. I normalize the cost of enrolling in non-STEM to zero. Finally, I determine the initial shares of low- and high-wealth using parental income data from the NLSY97. Specifically, I group respondents into two bins based on whether their average parents' income at the age of 18 was greater or less than the overall average of respondents' parents' income at age 18.

3.5.3 Internally-Calibrated Parameters

The model features seven parameters that are calibrated internally to match key moments from the NLSY97 data. This process ensures that the model closely replicates observed patterns in education, occupation, and entrepreneurial choices. Table 3.10 reports the parameters determined via internal calibration and their corresponding targets.

Table 3.10: Parameters determined jointly in equilibrium

Parameter		Value	Target	Data	Model
STEM Tuition	c_S	0.66	Share of STEM majors	25.06	24.93
Entrepreneurial Talent Distribution	α_z	2.82	Share of entrepreneurs	9.52	9.65
Prob. STEM Occupation	$p(o_S \mid m_S)$	0.65	Share STEM majors in STEM occ.	66.13	64.96
Prob. Non-STEM Occupation	$p(o_N \mid m_N)$	0.93	Share NS major in NS occ.	92.40	91.65
STEM Entrepreneur Threshold	z_S^*	2.82	Share of STEM occ. entrepreneurs	5.30	5.49
non-STEM Entrepreneur Threshold	z_N^*	2.20	Share of NS major entrepreneurs	81.25	80.89
Education Taste Shocks	ζ	8.51	Share of high-wealth entrepreneurs	10.34	9.62

The internal calibration process involves adjusting the key model parameters to match salient moments from the NLSY97 data. The STEM tuition cost, c_S, is calibrated to replicate the observed 25 percent share of STEM majors, while the shape parameter of the Pareto distribution for entrepreneurial talent, α_z is set to achieve the overall 9.52 percent entrepreneurship rate. To capture occupational sorting patterns accurately, I calibrate the probabilities of entering STEM and non-STEM occupations for respective majors, $p(o_S|m_S)$ and $p(o_N|m_N)$, to match the empirical shares of 66.13 percent and 92.40 percent, respectively. The occupation-specific entrepreneurial thresholds, z_S^* and z_N^*, are calibrated to match the 5.30 percent of STEM workers becoming entrepreneurs and the 81.25 percent of entrepreneurs originating from non-STEM majors. Lastly, the education taste shock parameter, ζ, is adjusted to match the 10.34 percent of high-wealth individuals choosing entrepreneurship.

The calibration process involves simulation the model with a population size of $N = 1,000,000$ individuals and adjusting the parameters iteratively until the model the model-generated moments closely match the targeted moments from the NLSY97 data. To formally minimize the objective function in this simulated method of moments approach, I employ a differential evolution algorithm. Overall, the calibrated model closely replicates the target moments.

3.5.4 Model Validation

Before turning to the quantitative analysis, I perform a final set of validation exercises using the calibrated model to assess its ability to replicate the facts from Section 1.3. Accordingly, I estimate the conditional probability of entering a STEM occupation by major, and the conditional probability of entrepreneurship by occupation, through a series of logistic regressions on the simulated data.

Table 3.11 and Table 3.12 report the results of these regressions, comparing the empirical estimates from the NLSY97 with those obtained from the simulated data.

TABLE 3.11: Empirical vs. Simulated Occupational Choice Probabilities

Pr[STEM Occupation]	**Data**	**Model**
Log Odds Ratios		
STEM Major	3.17^{***}	3.01^{***}
	(0.24)	(0.00)
Average Marginal Effects		
STEM Major	0.59^{***}	0.34^{***}
	(0.04)	(0.00)

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

TABLE 3.12: Empirical vs. Simulated Probability of Entrepreneurship

Pr[Entrepreneur]	**Data**	**Model**
Log Odds Ratios		
STEM Occupation	-0.89^{**}	-0.74^{***}
	(0.39)	(0.00)
Average Marginal Effects		
STEM Occupation	-0.08^{***}	-0.06^{***}
	(0.04)	(0.00)

Note: Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Importantly, the model performs well in qualitatively matching Facts 1 and 2 established earlier. For Fact 1, the simulated average marginal effect of majoring in STEM on entering a STEM-related occupation is 34 percentage points, compared to the 59 percentage points in the data. While the magnitude differs, the direction and significance of the effect are consistent. For Fact 2, the simulated average marginal effect of being in a STEM-related occupation on the probability of pursuing entrepreneurship is -6 percentage points, which closely mirrors the -8 percentage points observed in the data. The log-odds ratios from the simulated data also align reasonably well with their empirical counterparts, which further validates the model's ability to capture the underlying relationships. For instance, the log-odds ratio for STEM major in predicting STEM occupation is 3.01 in the model compared to 3.17 in the data, and for STEM occupation in predicting entrepreneurship, it is -0.74 in the model versus -0.89 in the data.

In summary, the alignment between the model and the data along various dimensions—not just in aggregate moments but also in key micro-level relationships—instils confidence that the calibrated model accurately reflects the underlying mechanisms linking major selection, occupations, and entrepreneurship and provides a strong foundation for the ensuing counterfactual experiments.

3.6 Quantitative Analysis

In this section, I use the calibrated model to perform a set of counterfactual experiments to assess the impact of policy interventions on educational choices and entrepreneurship. I consider two experiments in particular: (1) education financing policies that reduce the tuition associated with STEM majors and (2) entrepreneurship support programs that lower the thresholds for becoming an entrepreneur. The objective of these experiments is to quantify the effects of different policy interventions—albeit in a reduced-form way—and ultimately provides insights into their potential implications for educational choices and entrepreneurship in the present day.

3.6.1 Education Financing Policies

The first experiment examines the impact of reducing STEM tuition costs. This policy experiment is motivated by various recent initiatives aimed at encouraging more students to pursue STEM education, such as targeted scholarships or increased public funding for STEM programs. These policies are often discussed or implemented with the objective of addressing skill shortages in STEM fields and promoting innovation and growth more generally.

To perform this experiment, I take the calibrated model and gradually reduce the STEM tuition cost, c_S, from its baseline value. In particular, I consider reductions of 10 percent, 25 percent, and 50 percent from the original tuition cost. For each new scenario, I re-simulate the model and compute the new equilibrium outcomes. Table 3.13 below presents the results of this experiment.

TABLE 3.13: Education Financing Policy Experiment

Tuition Reduction	**STEM Majors**	**STEM Occ.**	**Entrepreneurs**
Baseline	24.93	22.46	9.65
10% Reduction	30.40	25.59	9.49
25% Reduction	37.03	29.33	9.28
50% Reduction	44.72	33.68	9.06

The results of this experiment yield several interesting insights into the relationship between tuition costs and educational choices. Specifically, as STEM tuition decreases, there is a substantial increase in the share of STEM majors in the economy. The share rises from a baseline of 24.93 percent to 30.40 percent with a 10 percent reduction in tuition, 37.03 percent with a 25 percent reduction, and 44.72 with a 50 percent reduction in tuition. The strong link between the cost and participation in STEM education suggests that financial aid policies can be an effective tool for encouraging more STEM graduates.

Corresponding to the rise in STEM majors, the economy also experiences a proportional increase in STEM occupations as the cost of STEM education decreases. In particular, the share of workers in STEM-related occupations starts at 22.46 percent at baseline and increases to over 33 percent with a 50 percent reduction in STEM tuition.

Interestingly, despite the increase in STEM majors and occupations, reductions in STEM tuition lead to a fall in the overall share of entrepreneurs across all three experiments. In the baseline, the share of entrepreneurs is 9.65 percent and it falls to 9.06 percent following a 50 percent reduction in STEM tuition. This finding suggests that there may be a potential trade-off between promoting STEM education and fostering entrepreneurship. It is possible that STEM graduates, faced with more attractive employment opportunities, may be less inclined to take the risk of pursuing entrepreneurial ventures.

These findings have important implications for policies related to education and entrepreneurship. The effectiveness of reducing tuition for increasing STEM participation in college supports its use as a policy tool. However, policymakers should be aware of the potential trade-off between a greater share of STEM workers and overall entrepreneurial activity. Moreover, while the model shows a clear short-term increase in STEM participation, it is remains unclear what the longer-term effects of such policies are. A larger STEM workforce could drive innovation and productivity growth, which would outweigh

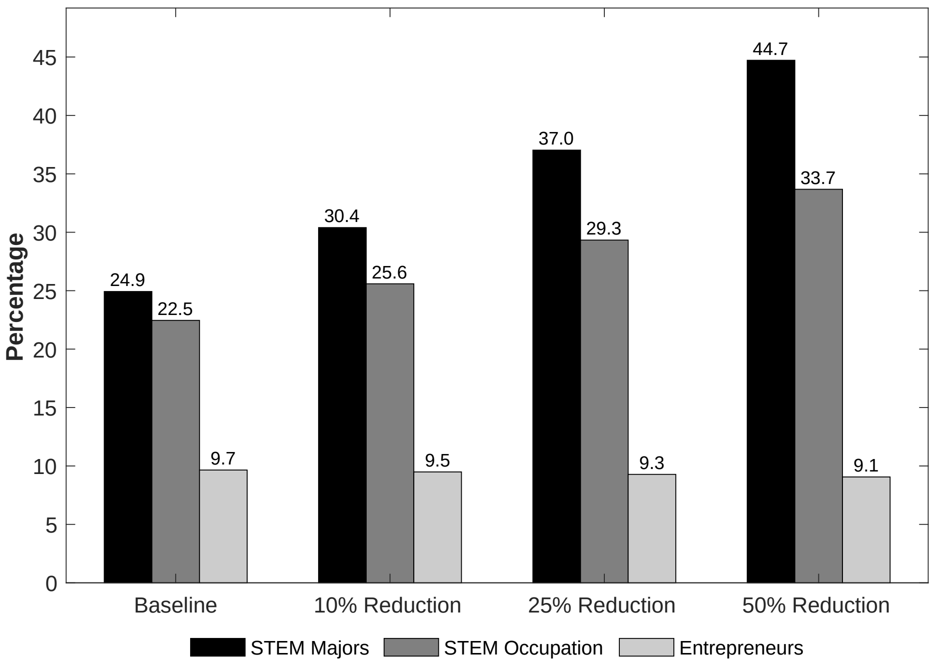

FIGURE 3.3: Education Financing Policy Experiment

the small decrease in entrepreneurship rates. It may also be possible that, conditional on becoming an entrepreneur, STEM-related ventures produce more profitable businesses. Such questions are left for future research.

3.6.2 Entrepreneur Support Programs

The second experiments explores the effects of policy interventions designed to encourage entrepreneurship by lowering the barriers to entry. These programs may include business incubators, mentorship initiatives, or reduced regulatory burdens for startups.

To implement this experiment in the model, I reduce the entrepreneurial thresholds, z_S^* and z_N^*, for both STEM and non-STEM occupations at equal magnitudes. Similar to the education financing experiment, I consider uniform reductions of 10 percent, 25 percent, and 50 percent from the baseline values. Given the new thresholds, I re-simulate the model and compute the implied equilibrium outcomes. Table 3.14 below presents the results of this experiment.

Table 3.14: Entrepreneurship Support Policy Experiment

Tuition Reduction	**STEM Majors**	**STEM Occ.**	**Entrepreneurs**
Baseline	24.93	22.46	9.65
10% Reduction	24.93	22.46	10.80
25% Reduction	24.93	22.46	13.03
50% Reduction	24.93	22.46	19.94

Lowering barriers to entrepreneurship unsurprisingly lead to an increase in entrepreneurial activity. In the baseline economy, the share of entrepreneurs stands at 9.65 percent. After lowering the thresholds by 10 percent, the share expands to 10.80 percent, 13.03 percent with a 25 percent reduction, and 19.94 percent with a 50 percent reduction. This increase suggests that policies aimed at reducing the barriers to entrepreneurs may be an effective way of encouraging more individuals to start a business.

Notably, the response to threshold reduction is highly non-linear, with a more pronounced increase in entrepreneurship rates when the thresholds reduce by 50 percent. This non-linear pattern implies that more substantial reduction in barriers to entry could have disproportionately larger effects on entrepreneurship rates. Yet, such policies appear to have no sizeable effect on STEM education or employment.

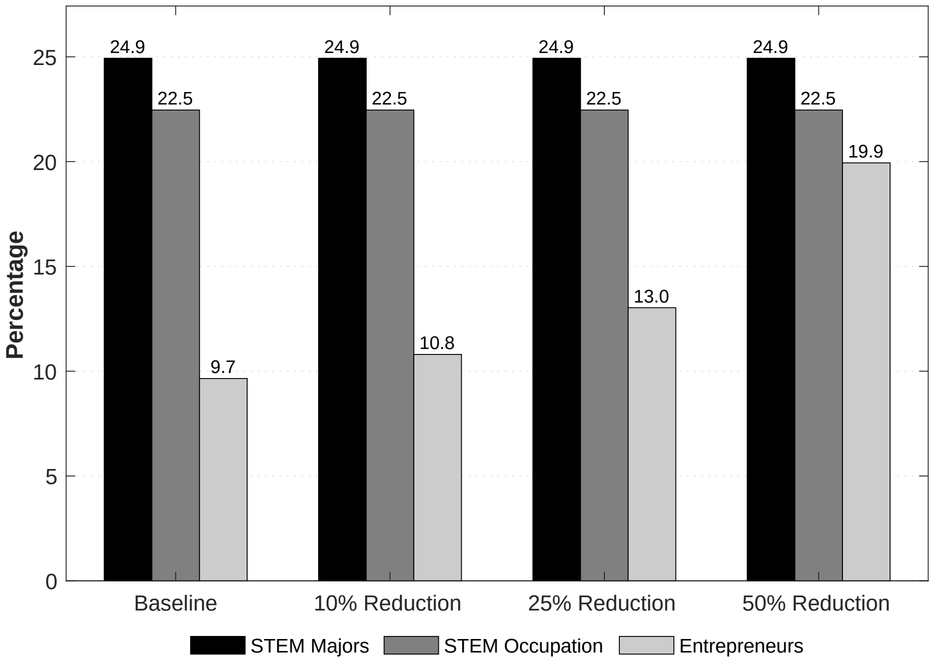

FIGURE 3.4: Entrepreneurship Support Program Experiment

The effectiveness of entrepreneurship support programs in encouraging entrepreneurship underscores their potential as a policy tool for fostering a more entrepreneurial economy. However, the lack of change in STEM versus non-STEM participation raises questions about whether targeted interventions could be more beneficial. For instance, if STEM entrepreneurship is deemed relatively more valuable for innovation and growth, policymakers might consider additional support specific to STEM graduates or workers.

When considered alongside the results of the tuition reduction experiment, these findings indicate that a combination of education and entrepreneurship policies might be necessary to simultaneously increase STEM participation and overall entrepreneurship rates. Such a multi-faceted approach could help to balance the goals of building a strong STEM workforce and fostering a dynamic entrepreneurial and innovative ecosystem.

3.7 Conclusion

A large quantitative literature has explored the factors influencing entrepreneurship rates in the United States. In this paper, I use data from the 1997 National Longitudinal Survey of Youth to document that the distribution of college majors, particularly the prevalence of STEM majors, is a quantitatively important factor affecting entrepreneurial activity among college graduates through its impact on occupational choices.

To reach this conclusion, I start by documenting a new set of facts using the NLSY97 micro-data. In particular, I show that STEM majors are significantly more likely to pursue employment in STEM-related occupations upon graduation, with a 53 percentage point greater probability compared to non-STEM majors; and, that individuals in STEM-related occupations exhibit systematically lower rates of entrepreneurship compared to their non-STEM counterparts. These facts hold even after controlling for various observable characteristics, including gender, race, parental income, and test score, and for various measures of entrepreneurship, including self-employment and business ownership.

I use these empirical findings to motivate a quantitative model of educational choice, occupational sorting, and entrepreneurship. The model provides the necessary structure to analyze the relationship between college major choices and entrepreneurship among graduates at both the micro- and macro-levels. Consistent with the empirical evidence, I find that the choice of college major significantly influences the likelihood of entrepreneurship through its impact on occupational sorting upon graduation.

After calibrating the model to match salient moments from the NLSY97 micro-data, I use it as a laboratory to assess the impact of two counterfactual policy experiments: an education financing policy, which reduces the tuition associated with STEM education, and an entrepreneurship support program, which reduces the threshold required to pursue an entrepreneurial venture. I find that reducing STEM tuition costs significantly increase STEM participation in education and employment, with the share of STEM majors rising from 24.93 percent ot 44.72 percent under a 50 percent tuition reduction. However, the rise in STEM majors also coincides with a decrease in entrepreneurship, which falls from 9.65 percent in the benchmark to 9.06 percent under a 50 percent tuition reduction. Conversely, entrepreneurship support programs increase entrepreneurship rates across all fields of study without any corresponding impact on the distribution of STEM enrollees or employees. In particular, a 50 percent reduction in entrepreneurial thresholds leads to an increase in the overall share of entrepreneurs from 9.65 percent to 19.94 percent.

Overall, this paper provides new insights into the relationship between educational choices, occupational sorting, and entrepreneurship. Future research could explore the long-term economic impacts of increased STEM education versus entrepreneurship rates, as well as investigate the impact of the large rise in STEM graduates for the decline in entrepreneurship in the United States since the 1980s.

References

Arcidiacono, Peter (2004). "Ability Sorting and the Returns to College Major". *Journal of Econometrics*, 121 (1-2), 343–375.

Bottia, Martha Cecilia et al. (2018). "Boosting the Numbers of STEM majors? The Role of High Schools with a STEM Program". *Science Education*, 102 (1), 85–107.

Braguinsky, Serguey, Steven Klepper, and Atsushi Ohyama (2012). "High-Tech Entrepreneurship". *The Journal of Law and Economics*, 55 (4), 869–900.

Braunerhjelm, Pontus, Sameeksha Desai, and Johan Eklund (2015). "Regulation, Firm Dynamics, and Entrepreneurship". *European Journal of Law and Economics*, 40, 1–11.

Cagetti, Marco and Mariacristina De Nardi (2006). "Entrepreneurship, Frictions, and Wealth". *Journal of Political Economy*, 114 (5), 835–870.

Delgado, Mercedes, Michael Porter, and Scott Stern (2010). “Clusters and Entrepreneurship”. *Journal of Economic Geography*, 10 (4), 495–518.

Deming, David J and Kadeem Noray (2020). “Earnings Dynamics, Changing Job Skills, and STEM Careers”. *The Quarterly Journal of Economics*, 135 (4), 1965–2005.

Elfenbein, Daniel W, Barton H Hamilton, and Todd R Zenger (2010). “The Small Firm Effect and the Entrepreneurial Spawning of Scientists and Engineers”. *Management Science*, 56 (4), 659–681.

Evans, David and Boyan Jovanovic (1989). “An Estimated Model of Entrepreneurial Choice under Liquidity Constraints”. *Journal of Political Economy*, 97 (4), 808–827.

Hsu, David H, Edward B Roberts, and Charles E Eesley (2007). “Entrepreneurs from Technology-Based Universities: Evidence from MIT”. *Research policy*, 36 (5), 768–788.

Huang, Yueling (2023). “Rethinking College Financing: Wealth, College majors, and Macroeconomic Consequences”.

Jiang, Helu and Faisal Sohail (2023). “Skill-Biased Entrepreneurial Decline”. *Review of Economic Dynamics*, 48, 18–44.

Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad (2016). “Field of Study, Earnings, and Self-Selection”. *The Quarterly Journal of Economics*, 131 (3), 1057–1111.

Kong, Dongmin and Ni Qin (2021). “Does Environmental Regulation Shape Entrepreneurship?” *Environmental and Resource Economics*, 80 (1), 169–196.

Kozeniauskas, Nicholas (2022). “What’s Driving the Decline in Entrepreneurship?” *Working Paper*, 1–51.

Krusell, Per et al. (2000). “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis”. *Econometrica*, 68 (5), 1029–1053.

Levine, Ross and Yona Rubinstein (2017). “Smart and Illicit: Who Becomes an Entrepreneur and Do They Earn More?” *The Quarterly Journal of Economics*, 132 (2), 963–1018.

Liu, Shimeng, Weizeng Sun, and John V Winters (2019). “Up in STEM, Down in Business: Changing College Major Decisions with the Great Recession”. *Contemporary Economic Policy*, 37 (3), 476–491.

Poschke, Markus (2013). "Who Becomes an Entrepreneur? Labor Market Prospects and Occupational Choice". *Journal of Economic Dynamics and Control*, 37 (3), 693–710.

Quadrini, Vincenzo (2000). "Entrepreneurship, Saving, and Social Mobility". *Review of Economic Dynamics*, 3 (1), 1–40.

Salgado, Sergio (2020). "Technical Change and Entrepreneurship". *SSRN 3616568*.

Vilalta-Bufi, Montserrat, Aleksander Kucel, and Giovanni Giusti (2018). "Who Becomes an Entrepreneur? The Role of Ability, Education, and Sector Choice".

Wiswall, Matthew and Basit Zafar (2015). "Determinants of college major choice: Identification using an information experiment". *The Review of Economic Studies*, 82 (2), 791–824.

Conclusion

This thesis has explored three pressing economic issues: the impact of employer-provided training on wage inequality; the role of firm sales and ownership transfers on firm entry, exit, and growth dynamics; and the impact of college major selection on occupational sorting and entrepreneurship. Through a combination of empirical analysis of micro-level data on individuals and firms and quantitative macroeconomic modelling, this research has yielded several important insights that contribute to our existing understanding of human capital development, firm dynamics, and the pathways to entrepreneurship.

In Chapter 1, I quantitatively established that the rise in employer-provided training, driven by the introduction of new technologies during the 1980s, 1990, and early 2000s, attenuated the rise in the college wage premium by 63 percent. To do so, I started by documenting a significant rise in training participation: between 1980 and the early 2000s, training participation among the working age population increased by over 40 percent, with the largest gains accruing to non-college educated workers. The finding that training had, in fact, attenuated the increase in the college wage premium underscores the importance of workplace learning and skills development programs in an era of skill-biased technological change. This result suggests that policies aimed at encouraging employer-provision of training may prove useful in addressing wage inequality and skill gaps in the workforce moving forward. In future work, it would be interesting to explore how public-funded training support programs impact wage inequality. Specifically, investigating whether government-support training programs would be welfare-improving, and determining the optimal level of training support, could provide value insights for policymakers seeking to address wage inequality in an era of rapid technological change and job-displacing automation.

In Chapter 2, Bettina Brüggemann, Zachary Mahone, and I shed new light on the previously under-explored role of firm sales and ownership transfers in shaping firm entry, exit, and growth dynamics. We first established, using administrative matched employer-employee data from Canada, that firm sales occur frequently: approximately 1.5 percent of firms are sold annually in Canada, which is larger than the exit rate for full-time employer businesses. Then, embedding firm sales into a quantitative model of firm dynamics, we showed that 13 percent of new entrant firms survive exclusively due to the option value of sale, and that realized ownership changes account for 18 percent of average log employment growth among small firms. Together, these findings highlight the importance of accounting for the firm sales margin in designing policies to encourage and support entrepreneurship and business growth. Building on this work, future research could analyze how business cycles impact the option value of firm sales.

In Chapter 3, I quantified the impact of college major selection on post-graduation occupational sorting and entrepreneurship. To do so, I used micro-data from the 1997 National Longitudinal Survey of Youth to empirically establish that STEM graduates are more likely to pursue STEM-related employment opportunities but less likely to become entrepreneurs than non-STEM majors. I then integrated an explicit college major choice into a model of occupational sorting and entrepreneurship to assess the impact of education financing and entrepreneurship support policies. I found that reducing STEM tuition would increase STEM enrolment at the cost of reducing entrepreneurship, while reducing barriers to entrepreneurship would increase the share of entrepreneurs without affecting STEM enrolment. Together, these results underscore the complex trade-offs involved in simultaneously fostering STEM education and entrepreneurship. Indeed, developing policies to achieve both objectives will require a balanced approach that explicitly accounts for the interconnected nature of educational choices and subsequent career outcomes. An intriguing direction for future work is to investigate the extent to which the rise in STEM majors has contributed to the decline in entrepreneurship observed in the United States since the 1970s and 1980s. Such research could shed new light on the long-term macroeconomic implications of changes in the aggregate composition of college majors and its role in driving entrepreneurial activity.

The three papers in this thesis collectively highlight the complex relationships between human capital development, firm dynamics, and entrepreneurship in modern economies. By examining these interconnections through a combination of micro-data analysis and macroeconomic theory, this research provides new and comprehensive insights into how individual choices, firm behaviour, and policy interventions jointly influence aggregate outcomes. A key theme that emerges across the papers is the importance of adaptability in both individual skill development and policy design. The first paper demonstrates how job-related training can dampen wage inequality arising from technological change, which underscores the value of continuous skill upgrading. The second paper reveals the significant role of ownership transfers in firm dynamics, emphasizing the need to separately account for business and owner dynamics in designing policies to support growth. The third paper examines the link between educational choices and post-graduate career outcomes, which highlights the need for balanced policies that promote both specialized skills and entrepreneurship. Overall, this thesis provides a solid foundation for future research to explore the complex interplay between human capital, entrepreneurship, and firm dynamics. By continuing to investigate these relationships, researchers can help policymakers design evidence-based strategies to address the evolving challenges and opportunities of the modern economy.

Appendix A

Appendix to Chapter 1

A1 Data Sources

A1.1 Workplace and Employee Survey (WES)

The main data source for this paper is the Workplace and Employee Survey (WES). The WES is a matched employer-employee survey data set from Canada, which covers approximately 20,000 employees and 6,000 employers at an annual frequency from 1999 to 2006. In each year, the WES contains two components: a workplace component and an employee component. The target population for the workplace component consists of all business locations operating in Canada with paid employees in March of the survey year with the exception of employers operating in the Territories; crop or animal production; fishing, hunting, and trapping; private households, religious organizations, or public administration. Hence, each observation in the workplace component of the WES is an *establishment*. In the main text, I use the words "workplace", "establishment", "employer", and "firm" interchangably. The target population for the employee component of the WES consists of all employees working or on paid leave in March of the survey year who (1) are employed by an establishment in the workplace component and (2) receive a Canada Revenue Agency T-4 Supplementary form. Workers receiving T-4 slips from multiple different workplaces are counted as distinct observations in the employee component of the WES.

The sampling methodology of the WES is divided into two parts. First, a sample of employers is drawn from the Business Register at Statistics Canada. Second, employees from the participating workplaces are selected at random from lists provided by employers to the surveyors. The initial sample was drawn in 1999. Every two years, a subset of new employers is added to the workplace component from establishments added to the Business Register since the last survey occasion. Every two year, a new sample of employees from the participating workplace is also drawn. Hence, the employee component of the WES is fully refreshed every odd year, while the workplace component is only partially refreshed. I pool the cross-sectional data from the 1999, 2001, 2003, and 2005 surveys, and restrict attention to workers of age 25 to 64 years old.

The main variables of interest from the WES include: educational attainment, classroom and on-the-job training participation, and earnings on the employee side; and, firm size, the share of employees using computers, and classroom or on-the-job training provision on the workplace side. Details about how I construct the measures of skills, training, and technology used in the empirical analysis are contained in the main text.

A1.2 Additional Data Sources for Calibration

For the calibration of the initial steady state, I use two additional data sources. The first is publicly-available aggregate data from the OECD, which is available here. I use the OECD data to obtain a target for the average unemployment rate and share of high-skill (tertiary-educated) workers in 1980. The second source is Statistics Canada and Human Resources Development Canada (2001), which documents training participation rates by education from the past revisions of the Adult Education and Training Survey (AETS).

A2 Empirical Specifications

The regression specifications underlying Fact 1 of Section 1.3 in the main text are detailed below. In all cases, I estimate the standard errors by bootstrap using the bootstrap weights provided by Statistics Canada and 100 replications.

A2.1 Workplace-Level Productivity

In the first regression, I estimate the effect of technology on (log) revenue productivity according to the following specification:

$$\ln(\text{Productivity}_{j,t}) = \beta_0 + \beta_1 \text{HighTech}_{j,t} + \xi Z_{j,t} + \theta_t + \varepsilon_{j,t}. \tag{A.1}$$

where j indexes firms and t indexes time. The main covariate of interest is $\text{HighTech}_{j,t}$, which is an indicator equal to 1 if firm j is a high-tech firm and 0 if low-tech. I also include a vector $Z_{j,t}$ of time-varying workplace-level control variables, which includes industry, training provision, and firm size; a set of year-fixed effects θ_t; and an error term $\varepsilon_{j,t}$.

A2.2 Employee-Level Training Participation

In the second regression, I estimate the probability of training participation at the employee-level to evaluate whether employees of high-technology firms are relatively more or less likely to receive training. To this end, I estimate a logistic regression of the following form:

$$\Pr[\text{Train}_{i,t} = 1] = \beta_0 + \beta_1 \text{HighTech}_{j,t} + \beta_2 \text{HighSkill}_{i,t} + \delta X_{i,t} + \xi Z_{j,t} + \theta_t + \varepsilon_{i,t}, \quad \text{(A.2)}$$

where, again, i indexes individuals, j indexes firms, and t indexes time. The main covariate of interest is the indicator $\text{HighTech}_{j,t}$, which equals 1 if worker i's employer j is a high-technology firm. I also control for the worker's level of education, $\text{HighSkill}_{i,t}$, a set of time-varying employee-level covariates $X_{i,t}$, a set of time-varying workplace-level covariates $Z_{j,t}$, year fixed-effects θ_t, and an error term $\varepsilon_{i,t}$.

A2.3 Probability of High-Technology Employment

For the third regression, I estimate the probability of being employed by a high-technology firm. That is, for each individual i employed by workplace j in year t, I estimate the following logistic regression:

$$\Pr[\text{HighTech}_{j,t} = 1] = \beta_0 + \beta_1 \text{HighSkill}_{i,t} + \delta X_{i,t} + \xi Z_{j,t} + \theta_t + \varepsilon_{i,t}, \quad \text{(A.3)}$$

where $\text{HighTech}_{j,t}$ indicates whether employee i's workplace j is high-technology; $\text{HighSkill}_{i,t}$ is an indicator for whether employee i is a high-skill worker; $X_{i,t}$ is a vector of time-varying worker-level control variables, which includes training participation, occupation, age, experience, gender, CBA coverage, immigration status, and tenure; and $Z_{j,t}$ is a vector of time-varying workplace-level control variables, which includes industry, firm size, and productivity.

Appendix B

Appendix to Chapter 2

Figure A2.1: Firm Survival Upon Sale Conditional on Risk Type

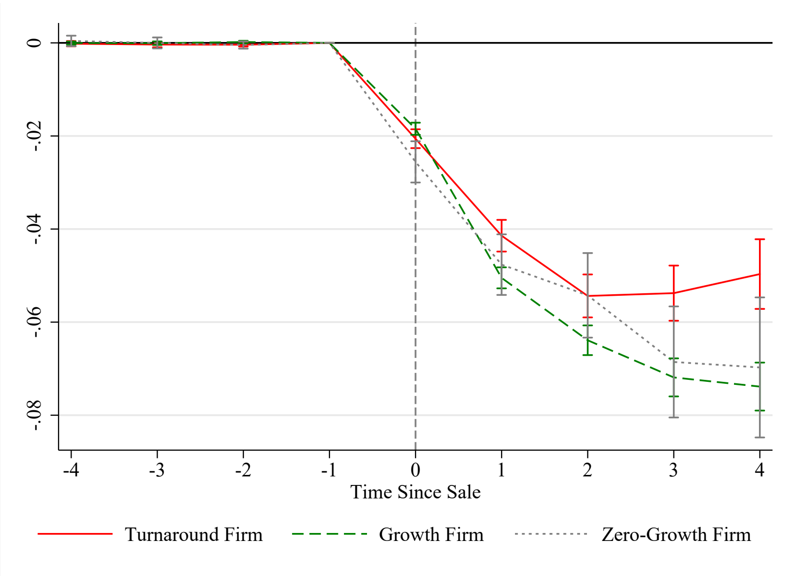

FIGURE A2.2: Firm Survival Upon Sale Conditional on Buyer Type