UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on Choice

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Tara Sullivan

Committee in charge:

Professor Valerie Ramey, Chair
Professor Titan Alon
Professor Prashant Bharadwaj
Professor Leo Porter
Professor Sally Sadoff

2022

The Dissertation of Tara Sullivan is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

DEDICATION

To my parents, who always reminded me, even at my lowest moments, that I was far too stubborn to quit my degree.

And to William, who made the tortellini and did the dishes without even being asked, like the true hero he is.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the guidance of my committee. I would like to specifically thank my chair, Valerie Ramey, for always pushing my work forward.

I've had many friends and colleagues support me throughout this degree. Special acknowledgement is reserved for Alyssa, Becky, Chelsea, Katherine, and Sophie. I never could have completed this degree without your personal and professional support.

Chapter 1 of this dissertation utilizes data from the NCES Beginning Postsecondary Students Longitudinal Study (BPS). I was only able to access these data because of Julian Betts' IES data lab. I am deeply grateful to Julian for taking the time to help me access the resources I needed. The dissertation author was a primary author of this chapter. It is currently being prepared for submission for publication of the material.

Chapter 2 presents a theoretical model that builds on work by Titan Alon and his coauthor, Daniel Fershtman. In addition to serving on my committee, Titan generously shared his own research and resources with me, providing key early guidance to shape my project. The dissertation author was a primary author of this chapter. It is currently being prepared for submission for publication of the material.

Chapter 3 is a reprint of the material published in the Environmental Defense Fund Economic Discussion Paper Series in June 2021. This work would not have been possible without financial support from the Environmental Defense Fund Pre-Doctoral Internship. I am particularly thankful for the help of my summer advisor, Jeremy Proville, along with the Chief Economist at EDF, Suzi Kerr, and the staff of the Office of the Chief Economist for enabling such a productive experience. The dissertation author was a primary author of this chapter.

VITA

2012	Bachelor of Arts, Economics and International Studies, Boston College
2012–2015	Research Analyst, Federal Reserve Bank of New York
2015–2022	Teaching Assistant, Department of Economics University of California San Diego
2019	Pre-Doctoral Intern, Environmental Defense Fund
2021	Intern, Wayfair
2022	Doctor of Philosophy, Economics, University of California San Diego

ABSTRACT OF THE DISSERTATION

Essays on Choice

by

Tara Sullivan

Doctor of Philosophy in Economics

University of California San Diego, 2022

Professor Valerie Ramey, Chair

Chapter 1 measures gender differences in college major choice and major switching behavior using a nationally representative sample of beginning postsecondary students. I find that switchers overall have worse graduation outcomes and take longer to earn their degrees than non-switchers. GPA appears to be a driver of switching: for students initially enrolled in a given field, the average GPA of female and male switchers is lower than the overall average for women and men, respectively. However, for most fields, the difference between the average GPA of female switchers and the average overall GPA is small. This suggests that women might interpret the signal sent by GPA in different ways than men. I then show that, conditional on GPA and initial STEM enrollment, women are more likely to change majors than men.

Chapter 2 builds a model of group-based beliefs and human capital specialization. Individuals belonging to a particular group choose to work or study in heterogeneous fields. This maps to the scenario where group type corresponds to gender, and the specialization decision is an individual's college major choice. Agents form initial beliefs about their unknown abilities based on existing group outcomes, and update these beliefs as they study. Therefore, women may form their initial beliefs about their abilities in each field by considering how many women have specialized in that field in the past. These differences in initial beliefs can cause men and women to respond to signals about their ability in different ways, which can ultimately drive gender gaps in college major choice. This model explores a mechanism for why, conditional on GPA and initial STEM enrollment, women are more likely to change major than men. If women form their initial beliefs about their abilities based on existing group outcomes, then women in historically underrepresented fields (like in STEM) might be more likely to switch than men at a given GPA level.

Chapter 3 explores the literature on technological adoption in agricultural climate change adaptation. Specifically, this chapter reviews two recent approaches to studying climate change adaptation in agriculture: panel data methods and spatial general equilibrium models.

Chapter 1

Gender differences in college major choice and switching behavior: an empirical assessment

1.1 Introduction

Over the past fifty years, the gender gap in postsecondary degree attainment has reversed, as seen in figure 1.1.[1] However, significant gender gaps persist within particular areas of study. Consider figure 1.2a, which plots the ratio of women to men completing bachelor's degrees

1. The gender convergence in overall college degree attainment is a well-studied phenomenon, as documented in Becker, Hubbard, and Murphy (2010). See also the review of gender wage gap research in Blau and Kahn (2017), which details the literature on the increase in overall degree attainment amongst women.

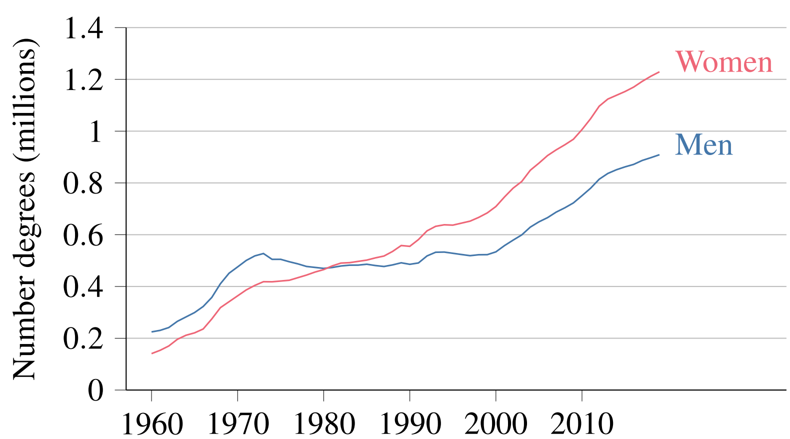

Figure 1.1. Number of Bachelor's Degrees awarded in US 4-year colleges. Source: IPEDS; Snyder (2013).

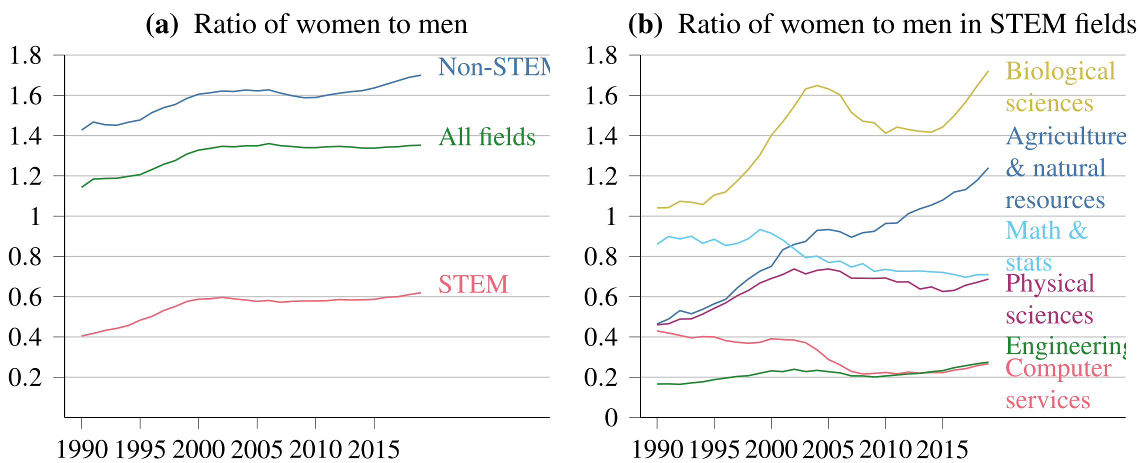

Figure 1.2. Ratio of women to men completing bachelor's degrees in U.S. 4-year colleges from 1990 - 2019. Figure (a) plots the ratio in STEM and non-STEM fields. Figure (b) plots the ratio in in various STEM sub-fields. Source: IPEDS.

in U.S. four-year colleges and universities from 1990-2019 in STEM and non-STEM fields. Although STEM degree holders have the highest average and median earnings post-graduation (Wiswall and Zafar 2018), the gender ratio in STEM has barely increased over the past twenty years. The gender ratios in STEM sub-fields are plotted in figure 1.2b. The current gender ratio in STEM is driven by women majoring in the biological sciences and, to a lesser degree, in the (much smaller) field of agricultural and natural resources. These "life sciences" as a whole pay much less than the physical sciences, engineering, math, or computer sciences (Altonji, Arcidiacono, and Maurel 2016). In general, women systematically choose majors with lower potential wages, a fact that has meaningful implications for the college-educated gender wage gap (Sloane, Hurst, and Black 2020).

An extensive literature attempts to explain the determinants of college major choice, a subset of which evaluates the factors that lead to the continued underrepresentation of women in certain fields (Patnaik, Wiswall, and Zafar 2020). A key hypothesis in both the theoretical and empirical literature revolves around dynamic learning processes.[2] Students may begin college

2. A brief review of the sequential uncertainty literature can be found in section 1.2; also see Arcidiacono (2004), Arcidiacono et al. (2016), and Stinebrickner and Stinebrickner (2014a, 2014b).

unsure about, say, their ability in or preference for a STEM field, and switch majors in response to learning about these factors while in school.[3] However, in spite of the rich literature built upon Altonji's (1993) sequential learning model, few papers summarize college major switching by gender; overall, it's unclear if men and women switch majors at different rates in particular fields.

In this paper, I empirically estimate the determinants of college major switching. Using a nationally representative, longitudinal survey of first-time beginning postsecondary students, I'm able assess students' college major switching behaviors beyond a single institution. I find that, conditional on GPA, women initially enrolled in STEM fields are more likely to change their majors than men. This suggests that men and women may respond to signals about performance in different ways, and provides an empirical motivation for the model developed in chapter 2 of this dissertation.

After briefly reviewing the literature on the determinants of college major choice in section 1.2, I outline my dataset and sample of interest in section 1.3. A key advantage of this resource is that it allows me to summarize outcomes for a nationally representative sample of first-time beginning postsecondary students pursuing bachelor's degrees. This dataset is an improvement over those commonly used in the literature because it includes students from a wide range of academic institutions. The results are therefore externally valid and include the 22% of students who begin their postsecondary studies in an associate's degree programs and intend to continue on to bachelor's.

Section 1.4 precisely defines major switching in this analysis. Specifically, switchers are students who change their major between the first wave of data collection (at the end of the first year of postsecondary study), and after the second wave of data collection (three years after students begin their postsecondary studies). To motivate switching as an outcome variable, I summarize graduation outcomes and time to degree for switchers and non-switchers in section

3. For this reason, major switching and switching costs play a key role in theoretical models of dynamic major choice, including Altonji, Arcidiacono, and Maurel (2016) and Arcidiacono et al. (2016).

1.5. Overall, students who change majors are less likely to graduate, and take longer to complete their degrees. At this point, gender differences in switching behaviors also emerge: female switchers initially enrolled in STEM fields are more likely to change into non-STEM fields than male-switchers initially enrolled in STEM.

I investigate the relationship between switching, major choice, and GPA in section 1.6. Unsurprisingly, switchers in general have lower GPAs than non-switchers across majors, though interesting patterns emerge when considering the gender of switchers. For most majors, the mean GPA of female switchers is higher than the mean GPA of male switchers. Further, there is little difference between the mean GPA of female switchers and the overall mean GPA in each major. This suggests that women may be more likely to compare their GPAs to those of other female students, rather than to the average student, when considering whether to change majors. Overall, the results in section 1.6 are consistent with the idea that GPA provides an important signal about whether to switch majors, but that men and women might interpret these signals differently.

Finally, section 1.7 evaluates the likelihood of switching major across years, with a focus on differences by gender. I find that, conditional on GPA, women initially enrolled in STEM fields are more likely to change their majors than men. This provides further evidence of men and women interpreting the signal sent by GPA in different ways within STEM.

This chapter provides empirical motivation for chapter two of this dissertation. In that chapter, I build a model of human capital specialization, where individuals belonging to a particular group choose to work or study in heterogeneous fields. This easily maps to the scenario where group-type corresponds to gender, and the specialization decision acts as a theoretical generalization of college major choice. I assume agents form initial beliefs about their unknown abilities based on existing group outcomes, and update these beliefs as they study. Therefore, women may form their initial beliefs about their abilities in each field by considering how many women have specialized in that field in the past. These differences in initial beliefs can cause men and women to respond to signals about their ability in different ways, which can ultimately drive gender gaps in college major choice. This model is capable of describing the scenario where a woman receives a particular first-year GPA in engineering and changes fields, whereas a man receives the same first-year GPA and persists.

The empirical results outlined below align with the theoretical model in chapter 2 in a few key ways. In the model, college major switchers tend to take longer to graduate, an empirical fact established in section 1.5.[4] The model also implies that women might be more likely to switch out of fields where they have been historically underrepresented; section 1.7 shows that women in the data are more likely to change out of aggregated STEM majors than men. Finally, the theoretical model implies that men and women might respond to signals about performance in different ways, particularly in fields where women have been historically under-represented. This is supported by the major switching facts documented in sections 1.6 and 1.7.

1.2 College major choice literature

A large literature exists on the determinants of college major choice, a subset of which focuses on how these determinants differ by gender. An extensive review of the empirical literature can be found in Patnaik, Wiswall, and Zafar (2020).[5] Note that while expected wages are clearly an important determinant of college major choice (Arcidiacono, Hotz, and Kang 2012), the literature suggests that non-pecuniary factors play a key role (Arcidiacono et al. 2020; Wiswall and Zafar 2015). In particular, differences in expected wages across gender do not appear to explain the gender gap in major choice (Zafar 2013; Wiswall and Zafar 2018). Therefore, this review focuses on several relevant non-pecuniary determinants of college major choice: first, preferences and norms; second, the influence of social networks and role models; and finally, uncertainty and learning about abilities.

4. Note that the only negative outcome in the theoretical model is time to graduate. The theoretical model does not explicitly measure the drop-out decision (though that could be considered a relatively straightforward extension). Intuitively, the model implies a cost for field switching, and one might expect worse graduation outcomes for students who change majors; this is consistent with the overall attrition results in section 1.5.

5. Earlier reviews about college major choice and occupational choice can be found in Altonji, Blom, and Meghir (2012) and Altonji, Arcidiacono, and Maurel (2016).

Preferences and norms are broadly defined in the college major choice literature; the catchall of "tastes" includes both the short-term enjoyment of coursework (Zafar 2013), as well as longer-term occupational preferences, such as those over temporal flexibility (Wiswall and Zafar 2018).[6] Recently, research using subjective expectations panel data has focused on unpacking "the black box of tastes" (Patnaik, Wiswall, and Zafar 2020). This evidence suggests that preferences and norms play a central role in pre-labor market specialization decisions (Wiswall and Zafar 2015; Arcidiacono et al. 2020).

Empirical evidence also suggests that social networks and role models can be important determinants of educational decisions. Role models can be cultural (Riley 2019), or drawn from one's own network (Lim and Meer 2020). Experimental work by Porter and Serra (2020) finds that the presence of same-gender role models has a significant impact on the number of women choosing to major in economics. Relatedly, both peer effects (De Giorgi, Pellizzari, and Redaelli 2010; Ost 2010) and family effects (Neilson et al. 2020; Patnaik, Wiswall, and Zafar 2020) appear to influence decisions about college major.

Abilities are also a key driver into different college majors. One can begin thinking about the role of ability in college major choice in terms of a simple Roy model (Roy 1951). If students begin college with known, idiosyncratic abilities that are best-suited for different majors, they should sort into the field best suited towards their particular skills. To that end, ability sorting, both at a general and field-specific level, appears important in the literature (Arcidiacono 2004).

As noted in Altonji's (1993) seminal work, this simple Roy model framework does not hold when students begin college uncertain about either their preferences or abilities. To that end, the sequential uncertainty literature uses theoretical and empirical models to understand how students learn about their own ability and preferences while in college (Arcidiacono 2004; Arcidiacono et al. 2016).[7] Empirical investigations of learning dynamics in college major choice

6. That this long-term taste influences college major choice is consistent with research suggesting that occupational temporal flexibility may be key in explaining the lifetime gender wage gap (Goldin 2014; Kleven, Landais, and Søgaard 2019). See Wiswall and Zafar (2018) for further discussion.

7. Note that the model developed in Altonji (1993) highlights the role of learning about both abilities and preferences. I will attempt to highlight those papers that emphasize both learning about preferences and learning

often at least partially estimate college major switching as one of several outcome variables; these studies are closely related to the empirical results in this paper. Arcidiacono (2004) presents a key early theoretical and empirical contribution to this sub-field. Similar to the results below, he uses longitudinal survey data (the NLS72), and presents descriptive statistics about college switching at the aggregated level (i.e. in the natural sciences, business, social science and humanities, and education). He then models the joint process of choosing college major and college choice, with major switching into a different (aggregate) field being a college major choice. However, gender differences in switching behaviors are not discussed in detail. His results are consistent with my own, and suggest that learning about ability is an important determinant of major switches; when students perform poorly, they may choose to switch majors. A focus in Arcidiacono (2004), and, later, in Arcidiacono et al. (2016), is using structural estimation to tease out the ability sorting. My results show that, if these hypotheses are appropriate, they may be important to investigate differentially by gender. Related work in Stinebrickner and Stinebrickner (2014a) estimate major choice at a single institution. They document that students often switch out of science majors, after being overly optimistic about their abilities, but do not focus on differential results for men and women. Wiswall and Zafar (2015) also focus on survey results from a single institution, but using an information experiment that allows them to identify counterfactual beliefs in the college major decision. Similar to other papers above, they do not focus on differences by gender.[8]

Related work considers whether men and women respond to signals about performance in similar ways (Stinebrickner and Stinebrickner 2014b). Several papers consider whether gender gaps in college major choice can be attributed to differential responses to grades (Rask and Tiefenthaler 2008; Ost 2010). Recent experimental evidence from Owen (2020) isolates the role of beliefs in college specialization decisions. Her results suggest that beliefs about academic performance partially drive the gender gap in college major decisions, providing evidence of

about abilities, as these models are sometimes closely related.

8. A more detailed analysis of gender differences in preferences for certain job characteristics is explored in Wiswall and Zafar (2018). However, this analysis does not focus on major switching, but rather occupational choice.

a potential mechanism for my results.[9] Overall, these papers are all focused on behavior at a single institution, and occasionally within a single major. My analysis allows me to compare student behavior across postsecondary institutions types.

Finally, a note about disentangling the effects of skills versus preferences. Empirical analyses of gender differences in college major choice often attempt to separate the effects of field-specific preferences and field-specific skills. This is perhaps most convincingly attempted within the subjective expectations literature. This subset of the college major choice literature uses subjective expectations data to elicit unobservable information on college major choice (Zafar 2013; Stinebrickner and Stinebrickner 2014a).[10] Survey questions are used to elicit beliefs about counterfactual major choices. This framework can also be used to specifically explain differences in college major choice by gender (Wiswall and Zafar 2018). However, one disadvantage of this approach is that it is often limited to the analysis of a single institution, limiting their external validity (Patnaik, Wiswall, and Zafar 2020); analyses of first-time beginning students at multiple institutions are less common.

1.3 BPS data

Overall, dynamic learning about abilities appears to play an important role in college major choice, and may explain why men and women choose to specialize in different fields. However, evidence on learning in college is often limited to analyses of a single institution, often within a single major. While this type of filter is beneficial for precise causal estimates, it hinders our ability to evaluate the population of students beginning their post-secondary studies for the first time.[11] In particular, 33% of students in the dataset begin their post-secondary studies enrolled in associate's degree programs, with the intention to transfer to bachelor's degrees at a

9. Her results isolate the mechanism whereby other types of informational interventions, such as those proffered in Porter and Serra (2020), Bayer, Bhanot, and Lozano (2019), and Li (2018), impact differences in college major choice by gender.

10. See reviews in Patnaik, Wiswall, and Zafar (2020) and Altonji, Arcidiacono, and Maurel (2016) for details.

11. Students beginning their post-secondary studies for the first time are referred to as "first-time beginning students" in NCES data.

later date. This common degree pathway is filtered out of single-institution analyses.

To that end, I turn to the Beginning Postsecondary Students Longitudinal Study (BPS), a nationally representative, longitudinal survey of first-time beginning students administered by the National Center for Educational Statistics (NCES). Section 1.3.1 reviews the BPS data, and reviews its advantages and drawbacks. Section 1.3.2 then defines our population of interest, and outlines the filters applied to the data. The precise definition of students' majors and related summary statistics are presented in section 1.3.3.

1.3.1 Overview of data

The 2012/17 Beginning Postsecondary Students Longitudinal Study (BPS:12/17), and associated Postsecondary Education Transcript Study (BPS:12 PETS) are derived from a nationally representative sample of postsecondary students who began their postsecondary studies for the first time in 2011-12.[12] Students were then surveyed three times over six academic years: in 2011-12, in 2014, and in 2017 (henceforth, I will often refer to these time periods as waves 1, 2, and 3, respectively). Data from student interviews and administrative data comprise the BPS:12/17 data. BPS:12 PETS collects data on postsecondary transcripts from all institutions attended by BPS:12/17 cohort members.[13] Sample members across surveys were weighted to account for sub-sampling, non-response, and student eligibility. This analysis primarily uses the BPS:12 PETS panel study weight.[14]

After data collection, the NCES constructed a set of derived variables to facilitate analysis. For example, transcript data on grades and credits are normalized, so they can be comparable across students. Additionally, students' majors are classified according to the 2010 Classification of Instructional Programs (CIP) code, or CIP code. Henceforth, I will generally use "field" and "major" interchangeably to refer to the BPS classification of the student's major.

12. Specifically, the student universe for the BPS:12/17 consists of all first-time beginner (FTB) students at post-secondary institutions eligible for inclusion in the 2011-12 National Postsecondary Student Aid Study (NPSAS:12). For additional data documentation on BPS:12/17, see Bryan, Cooney, and Elliott (2019).

13. For additional data documentation on BPS:12 PETS, please see Bryan et al. (2020).

14. Further details about survey weights can be found in A.1.1.

An advantage of the BPS study is that it encompasses a wide array of student experiences. Institution-specific analysis, such as those used in Zafar (2013) or Rask and Tiefenthaler (2008), are preferable for precise summaries of student performance. Samples are often limited to focus on students who graduate, students who do not transfer, etc. A single-university sample reduces endogeneity concerns stemming from institution-specific unobservables.[15] However, this type of approach comes at a cost. Often, the institutions studied are elite research institutions or private institutions, which are not representative of most degree-seeking students. Thus, the institutional focus neglects wide arrays of student experiences in favor of ameliorating endogeneity concerns.

1.3.2 Filters applied to data

This analysis aims to summarize major switching in students pursuing bachelor's degree. Thus, the population of interest is students beginning to pursue a bachelor's degree for the first time, making BPS data an appropriate resource. However, because BPS data captures so many types of educational experience, it's necessary to apply a number of filters to our data to ensure we are focused on the appropriate sample. These can be divided into two filters: one applied to the first wave of data collected, and one applied to the second.

Wave 1 filter

Wave 1 of data collection occurs during the 2011-12 school year, which is students' first year of postsecondary study. Because the BPS survey is meant to encompass a range of postsecondary experiences, not all students are enrolled in bachelor's degrees, or even have the intention of attaining one. Summary statistics concerning institution and degree types by gender for the entirety of the BPS data can be seen in table 1.1. Second, not all students beginning their post-secondary studies are pursuing bachelor's degrees. Over 40% of students beginning their post-secondary studies in 2011-12 are enrolled in associate's degree programs. Further,

15. For example, universities may have major enrollment caps or GPA restrictions that constrain students' major choice (though recent evidence presented in Bleemer and Mehta (2021) suggest this effect may be concentrated at public universities).

Table 1.1. Weighted proportions in BPS data (percentages). Top of table reports proportion of men and women by degree program in all BPS:12/17 data. Bottom of table reports overall fraction of men and women in the sample used for the remainder of this analysis, namely the sample of first-time beginning students in 2011-12 either (1) enrolled in a bachelor's degree program, or (2) enrolled in an associate's degree program, with the intention of continuing on to a bachelor's degree program within the next 5 years. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Undergraduate degree program 2011-12			
Certificate	6.00	11.99	9.46
	(0.88)	(0.79)	(0.58)
Bachelor's degree	46.70	45.26	45.87
	(1.87)	(1.26)	(1.28)
Not in a degree program or others	1.77	1.91	1.85
	(0.46)	(0.51)	(0.35)
Associate's degree no Bachelor's intentions in 2012	11.07	9.04	9.90
	(1.04)	(0.80)	(0.69)
Associate's degree with Bachelor's intentions in 5 years	34.46	31.80	32.93
	(1.92)	(1.21)	(1.31)

the majority of students enrolled in associate's degree programs intend to continue on to a bachelor's degree program in the next five years.[16] This presents a key advantage of the BPS data; I'm able to consider those with bachelor's degree intentions who are currently enrolled in associate's degree programs. Transferring from associate's degree programs is a common pathway to a bachelor's degree. Consider, for example, the California State University (CSU) system, the largest four-year public university system in the United States ("The CSU Fact Book 2021" 2021). California Community Colleges offer an "Associate Degree for Transfer" which guarantees admission to California State campuses as a junior. Overall, 80,000 California community college students transfer to a University of California (UC) or CSU campus each year.[17] Because my sample of interest is those aiming to complete bachelor's degrees, it's valuable that the BPS includes both those enrolled in bachelor's degree program in addition to

16. Specifically, respondents who are either enrolled in an associate's degree program or undergraduate classes only are asked in 2011-12: "Do you plan to continue on to a bachelor's degree program within five years from now?"

17. See the California Community College website for details: https://www.cccco.edu/Students/Transfer.

those in associate's degree programs.

To that end, my wave 1 filter includes first-time beginning students who are either (a) enrolled in a bachelor's degree program, or (b) enrolled in an associate's degree program and expect to attain a bachelor's degree within the next five years. Clearly it is possible to filter on other dimensions of the data; for example, table A.1 lists the institution type for students. The majority of students (about 70%) are completing their degrees at public institutions, and relatively few students are enrolled in private, for-profit institutions (about 7.6% of the data). In general, I will keep all institution types in my sample, though occasionally it will be helpful to consider how results change when I also filter on institution type.

Wave 2 filter

The second wave of BPS data collection began in June 2014. For students continuously enrolled in an undergraduate degree program, this would correspond to approximately their junior year. My wave 2 filter is consistent with the wave 1 filter outlined in section 1.3.2, and includes all students who, by the second wave of data collection, were either (a) last enrolled in a bachelor's degree program, or (b) last enrolled in an associate's degree program, with the intention of attaining a bachelor's degree in two years. Note that the wave 2 filter is applied after applying the wave 1 filter, and filters out approximately 20% of the sample.[18] Table A.2 lists the undergraduate degree programs and enrollment patterns in wave 2, after the wave 1 filter has been applied. It's worth noting that overall enrollment patterns are quite varied at this point in students' postsecondary careers, and the detailed outcomes for those originally enrolled in bachelor's degrees and associate's degrees in 2011-12 can be found in tables A.3 and A.4, respectively. Because the specific route towards a degree can be so varied, I will not focus on enrollment patterns, nor the attainment of degrees other than bachelor's degrees.

18. To be precise: the wave 2 filter is only applied after the wave 1 filter, meaning that the wave 2 filter is applied to the sample of beginning postsecondary students either (1) enrolled in a bachelor's degree program in wave 1, or (2) enrolled in an associate's degree program in wave 1, with the intention of continuing on to a bachelor's degree in 5 years. Approximately 20% of that sample has changed from "pursuing" a bachelor's degree in some form, meaning that they were last enrolled in an undergraduate certificate diploma, or associate's degree program without intentions for their bachelor's degree, or not in a degree program.

Table 1.2. Gender composition of sample. Fraction male and female in BPS data after applying (1) the wave 1 filter, and (2) the wave 1 and wave 2 filters. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Fraction Male	Fraction female
Wave 1 filter	0.43	0.57
	(0.01)	(0.01)
Wave 1 & wave 2 filters	0.42	0.58
	(0.01)	(0.01)

Table 1.2 presents the gender composition of the sample after applying (1) the wave 1 filter, and (2) the wave 1 and wave 2 filters. Notably, the gender composition does not meaningfully change across waves of data collection.

1.3.3 Definition of major/field of study

After constructing my sample, I can now turn to students' choice of major/field of study (note that I will be using "major" and "field of study" interchangeably throughout this analysis). As discussed in section 1.3.1, BPS data classifies students' majors according to the 2010 CIP Codes.[19] Table 1.3 presents the percentage of men and women enrolled in each field, as well as the overall gender ratio.[20] We might be concerned that students who declare a major by the end of their first year are not representative of the overall student body. However, conditional probability models discussed in appendix A.1.3 show that no observable characteristics predict first-year major declaration.

The ratio of women to men choosing to study each field is in the rightmost column. The overall ratio of women to men is 1.35. Men appear to be overrepresented in computer and information sciences, engineering, and mathematics. Women are overrepresented in education, health care fields, and psychology. Table A.6 aggregates these fields up to STEM and non-STEM fields.[21] In general, men are overrepresented across STEM fields.

19. I specifically use the variable "majors23" in the BPS data. A more disaggregated taxonomy of student majors is available for the first wave of BPS data collection, in 2011-12, though unfortunately not available in subsequent waves.

20. Bootstrap standard errors for the proportions in table 1.3 can be found in the appendix table A.5

21. The specific variable used ("stemmaj" in BPS data) indicates whether the student's major field of study in

Table 1.3. Initial undergraduate major. Undergraduate field of study from 2011-12 (wave 1 of data collection). The first two columns show the percentage of male and female students in the sample in each undergraduate major in 2011-12. The final column Shows the ratio of women to men in each field. BPS:12 PETS Panel study weight.

	Male	Female	Ratio
Undecided or Undeclared	5.38%	5.23%	1.31
Computer and information sciences	6.69%	0.65%	0.13
Engineering and engineering technology	13.03%	2.00%	0.21
Biological and physical science, science tech	8.82%	8.77%	1.34
Mathematics	0.86%	0.44%	0.68
Agriculture and natural resources	1.26%	1.33%	1.42
General studies and other	10.87%	8.45%	1.05
Social sciences	4.01%	3.85%	1.29
Psychology	2.45%	7.58%	4.16
Humanities	5.15%	8.05%	2.11
History	1.40%	0.49%	0.47
Personal and consumer services	2.67%	2.17%	1.09
Manufacturing, construction, repair, transportation	1.19%	0.55%	0.62
Military technology and protective services	5.26%	3.29%	0.84
Health care fields	6.94%	17.00%	3.3
Business	15.76%	11.20%	0.96
Education	2.98%	8.49%	3.83
Architecture	0.76%	0.38%	0.67
Communications	2.07%	3.70%	2.41
Public administration and human services	0.33%	2.77%	11.46
Design and applied arts	1.44%	2.05%	1.92
Law and legal studies	0.49%	0.88%	2.41
Theology and religious vocations	0.20%	0.68%	4.58
All	100.00%	100.00%	1.35

1.4 Definition of field switching

Now that I have established the sample of interest, I can formally define field switching. I introduce this definition in section 1.4.1. Potential alternative definitions are introduced and briefly reviewed in section 1.4.2.

1.4.1 Primary definition of field switching

In this paper, I define a "switcher" as a student whose undergraduate field of study changes between wave 1 of data collection (the 2011-12 school year) and wave 2 (2014). For students continuously enrolled in college, this corresponds to a student's major choice between their first and third year of postsecondary education. Because "field" and "major" are used interchangeably in this analysis, I similarly use "field switchers" and "major switchers" interchangeably.

In the sections that follow, I pay close attention to gender differences in major switching for STEM majors and non-STEM majors. STEM and non-STEM majors represent collections of the majors outlined in table 1.3.[22] However, it's worth emphasizing that my primary outcome variable focuses on changing majors, not changing into or out of aggregated STEM majors as a whole. A student who changes major from engineering to computer science may be labeled as a "switcher," while remaining continuously enrolled in STEM. This definition differs somewhat from related literature, in particular Arcidiacono (2004), which focuses on the decision to change into or out of aggregated STEM and non-STEM fields. An intuitive reason for employing my definition of field switching simply has to do with the nature of college major requirements. A student who changes majors within STEM still has to take on additional course requirements, and possibly learn different fundamentals; the time to degree overall may be longer for switchers

2011-12 has a focus on science, technology, engineering, and mathematics (STEM) fields. This more clearly indicates a STEM major than the CIP codes used to classify majors in table 1.3. For example, several majors within General Studies and Other are interdisciplinary STEM fields.

22. Specifically, STEM fields include: Computer and information sciences; engineering and engineering technology; biological and physical sciences and sciences technology; mathematics and statistics; and agricultural and natural resources. All other fields are labeled as "non-STEM." When undecided or undeclared is included in the sample, it will be labeled as non-STEM, though it is often dropped.

rather than non-switchers. Additional empirical motivation for utilizing my definition of field switching is developed in section 1.5. Nevertheless, I do review some summary statistics for changing into or out of aggregated STEM fields in section A.1.4.

Finally, I want to note that I am looking at switching at a very particular point in time. I label a student as a "switcher" when they change their major between the first and second waves of data collection. As will be briefly reviewed in section 1.4.2, it is relatively common for students to change their major within their first year of study (the 2011-12 school year). However, as the BPS data only collects information on initial major choice at the end of the 2011-12 school year, this behavior will not be captured in my primary definition of switching. It is also possible for students to change major after the second wave of data collection, and before graduation. This is a particularly important point when considering the graduation outcomes of students who switch majors in section 1.5; a student in my dataset may (1) initially be enrolled in a non-STEM field, (2) be labeled as a "non-switcher" according to my definition, but (3) subsequently change their major to a STEM field after wave 2 and graduate. Overall, this is quite not common,[23] but means that one should take care when interpreting the results in section 1.5, specifically tables 1.4 and 1.5.

1.4.2 Alternative definition of field switching

It's possible to employ alternative definitions of field switching in my analysis. For example, the BPS survey directly asks students about major switching at two points: first, during wave 1 of data collection, at the end of the 2011-12 school year; and second, during wave 2 of data collection, in June 2014. I classify students who claim to have changed majors in their survey responses as "self-reported" switchers. This is distinct from the "observable" field switching

23. As noted in table 1.4, 4% of male non-switchers initially in a non-STEM field eventually graduate with a STEM bachelor's degree (and thus switch fields after wave 2 of data collection). For the other groups: 2% of female non-switchers in non-STEM graduate with a STEM bachelor's degree; 4% of male non-switchers initially in STEM graduate with a non-STEM bachelor's degree; and 9% of female non-switchers initially in STEM graduate with a non-STEM bachelor's degree. As a result of these small sample sizes, the estimates for time-to-degree in table 1.5 for these sub-groups are particularly noisy and should be interpreted with care.

I utilize as my primary definition of field switching. In theory, self-reported major switching could capture switching that is unobservable in BPS data. For example, unobservable switching might occur if students change fields within their first year of study, before the initial major is measured in BPS data. Additionally, self-reported major switching could capture switching at a more disaggregated level than available in my data. My available major categories, listed in table 1.3, aggregate certain fields. Changing majors within these aggregations is not observable in BPS data.[24] In practice, self-reported switching seems to be either (a) driven by institutional factors, such as whether an institution is private or public; or (b) impacted by students under-reporting their own switching behavior. I briefly review why below, with more information being available in the appendix sections A.1.5 and A.1.6

Section A.1.5 measures the prevalence of self-reported major switching during a students' first year of postsecondary education. Overall, self-reported major switching during a student's first year is somewhat common; almost 20% of the sample formally change major during the 2011-12 school year. The determinants of self-reported switching within a student's first year are estimated using conditional probability models in table A.12. Overall, few variables are strongly predictive of self-reported switching within the first year, except for one: students enrolled in a public institution (compared to a private one) appear more likely to report switching their major during their first year. This can perhaps be explained by public institutions tending to rely on GPA restrictions for certain popular majors, causing students who fall below a threshold to switch (Bleemer and Mehta 2021). An implication of this result is that institutional factors may play a particularly important role in major switching within a students first year. A goal of this analysis is to evaluate a potential signaling story: do differences in college major choice represent men and women responding differentially to signals about performance? If major switching within a student's first year is strongly determined by institutional factors like GPA cutoffs, then a cross-institutional analysis might not be capable of capturing this behavior.

24. For example, a student may be a physics student in wave 1, and a biology student in wave 2, but that behavior would be unobservable in the BPS data.

Section A.1.6 discuses self-reported major switching across survey waves.[25] In theory, this variable should capture more information than the observable switching. In practice, students appear inclined to under-report their own switching behavior, as documented in table A.13. Investigating why students under-report their own switching behavior may be an interesting area for future research, but will not be addressed here.

1.5 Switching and graduation outcomes

Theoretical models generally assume major switching carries some sort cost (Altonji, Arcidiacono, and Maurel 2016), but this is often simply assumed and not evaluated empirically. To that end, this section establishes some basic empirical facts about switching behaviors and two types of graduation outcomes: whether or not a degree is completed, and the amount of time it takes for a student to attain a degree. Overall, I find that switchers are less likely to attain their bachelor's degrees than non-switchers, and on average take longer to complete their degrees. I also find preliminary evidence that female switchers initially enrolled in STEM fields are more likely to change into non-STEM fields than their male counterparts.

Table 1.4 measures the fraction of students attaining a bachelor's degree within six years across a number of groupings. Panel A presents the fraction of male and female switchers and non-switchers attaining degrees within six years. Sixty percent of students attain a bachelor's degree within six years, with this number being higher for women than for men. Students who switch majors are less likely to attain their bachelor's than students who do not switch majors. While 53% of female switchers attain their bachelor's within six years, only 45% of male switchers do.

Panels B and C of table 1.4 present the same results as panel A for students initially enrolled in STEM and non-STEM majors during wave 1 of data collection. The graduation outcomes are further disaggregated to reflect graduating with STEM or non-STEM degrees. The

25. Specifically, the variable "majchgba14" indicates the number of times the respondent formally changed his or her bachelor's degree major at any institution through June 2014, according to the student's first follow-up interview.

Table 1.4. Field switching and degree attainment. Shows the fraction of students attaining their bachelors' degree within six years by gender, by initial major type, and by whether they switched majors between 2011-12 and 2014. BPS:12 PETS data are weighted by BPS:12 PETS Panel study weight.

	Do not switch	Switch	All
A. Field switching and degree attainment			
All students			
Do not attain BA	0.32	0.50	0.40
Attain BA	0.68	0.50	0.60
Male students			
Do not attain BA	0.36	0.55	0.45
Attain BA	0.64	0.45	0.55
Female students			
Do not attain BA	0.29	0.47	0.37
Attain BA	0.71	0.53	0.63
B. Students initially in non-STEM fields			
Male students initially in non-STEM			
Graduated with STEM bachelor's	0.04	0.07	0.06
Graduated with non-STEM bachelor's	0.59	0.36	0.48
Did not graduate with bachelor's	0.36	0.57	0.46
Female students initially in non-STEM			
Graduated with STEM bachelor's	0.02	0.05	0.03
Graduated with non-STEM bachelor's	0.68	0.45	0.58
Did not graduate with bachelor's	0.30	0.50	0.39
C. Students initially in STEM fields			
Male students initially in STEM			
Graduated with STEM bachelor's	0.59	0.20	0.46
Graduated with non-STEM bachelor's	0.04	0.30	0.13
Did not graduate with bachelor's	0.37	0.50	0.41
Female students initially in STEM			
Graduated with STEM bachelor's	0.72	0.08	0.42
Graduated with non-STEM bachelor's	0.09	0.62	0.33
Did not graduate with bachelor's	0.19	0.31	0.25

overall graduation results for initial non-STEM students in panel B appear generally consistent with the results in panel A. Switchers initially enrolled in non-STEM fields are less likely to graduate than non-switchers, and women are more likely to graduate than men. 54% of male students initially in non-STEM fields graduate within six years (compared to the 55% rate for all male students), and 61% of female students initially in non-STEM fields (compared to 63% for all female students). Similarly, 50% of female switchers in non-STEM graduate within six years (compared to 53% of all female switchers), while 43% of non-STEM male switchers do (compared to 45%).

By disaggregating degree type, panel B shows that the majority of graduating male and female switchers initially enrolled in a non-STEM major remain in non-STEM fields (84% and 90%, respectively). Note that the definition of switcher and non-switcher in table 1.4 reflects switching majors between waves 1 and 2 of the data; therefore, in panel B, it's possible for a non-STEM major to be labeled as a non-switcher, but ultimately graduate with a STEM major. This is somewhat rare; only 4% of male non-switchers in non-STEM fields and 2% of women subsequently change majors into a STEM field and graduate.

Table 1.4 panel C lists the graduation outcomes for students initially enrolled in STEM fields. Similar to panels B and C, switching appears to be associated with worse graduation outcomes, and women are more likely to graduate than men. Here we see some marked differences from panels B and C. First, women initially enrolled in STEM are far more likely to graduate than men initially enrolled in STEM. Switchers initially in STEM are more likely to graduate in a non-STEM field than in a different STEM field; this is different than what is seen in panel B for non-switchers from non-STEM, who tend to change major within non-STEM. However, this difference is even more stark for women. While 20% of the men initially in STEM who switch majors still graduate in a STEM field, only 8% of women do. Most women initially enrolled in STEM who switch fields switch to a non-STEM major and graduate (62%). This corresponds to approximately 89% of the female switchers initially in STEM who graduate. Only 60% of the men initially enrolled in STEM who switch fields and graduate do so by switching to

Table 1.5. Field switching and time to degree. Table contains the average months enrolled before attaining a bachelor's degree, by initial major type, final degree type, and gender. BPS:12 PETS data are weighted by BPS:12 PETS Panel study weight.

	Male		Female	
	No switch	Switch	No switch	Switch
Initial STEM major				
Graduated with STEM bachelor's	43.7	46.6	42.3	43.4
Graduated with non-STEM bachelor's		45.3		44.6
Initial non-STEM major				
Graduated with STEM bachelor's		46.8		45.3
Graduated with non-STEM bachelor's	42.2	46.0	42.5	44.5

a non-STEM major.

Overall, table 1.4 illustrates worse graduation outcomes for switchers than non-switchers. Graduating students who switched out of a non-STEM field likely still graduate in a non-STEM field, whereas graduating students who switched out of a STEM field likely switched into a non-STEM field. This is especially true for women initially in STEM fields, who are far more likely to switch into a non-STEM field and graduate than men.

Table 1.5 considers a different type of graduation outcome: the average number of months enrolled before attaining a degree. For clarity, the relatively small group of students who appear to switch majors after wave 2 have been dropped from the table.[26] The results show that students who switch fields on average take longer to complete their degrees. This holds for both men and women initially enrolled in STEM and non-STEM fields.

That switchers have worse graduation outcomes should not imply that major switching is inherently bad for students. Instead, the results in tables 1.4 and 1.5 are consistent with the generally accepted hypothesis that college major switching carries some cost. Additionally, the

26. Recall from table 1.4 that 4% of male non-switchers initially in non-STEM fields and 2% of female non-switchers subsequently change majors into a STEM field and graduate. These types of male and female non-switchers are enrolled for an average of 42.7 and 41.3 months before graduation, respectively. Only 4% of male switchers and 9% of female non-switchers initially enrolled in STEM fields graduate with a non-STEM bachelor's. Male non-switchers who graduate in STEM are enrolled for an average of 41.7 months, whereas female non-switchers are enrolled for an average of 48 months. The female non-switchers who subsequently change into STEM are the only non-switchers who take longer to graduate than switchers. This may be an interesting sub-group to investigate, but will not be the purpose of this analysis.

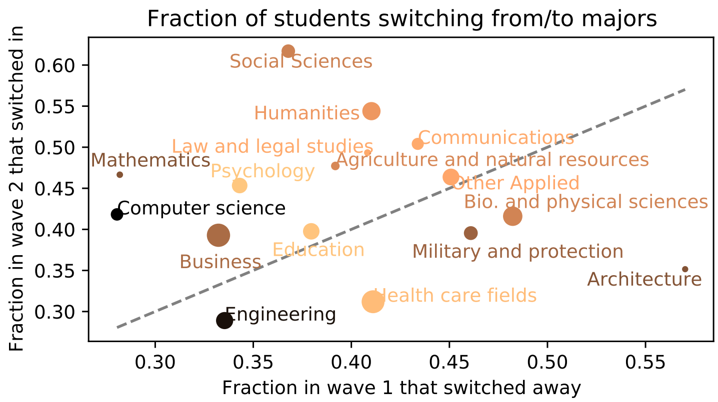

Figure 1.3. Comparison of the overall fraction of students switching out of a major during wave 1 of the BPS data collection to the fraction of students that switched in by wave 2. The color of the markers indicate overall gender ratio in that field during wave 1; darker markers indicate male-dominated fields. The size of the markers indicate the overall size of the field in the data during wave 1; fields with more students are larger.

results in table 1.5 specifically are consistent with the theoretical model outlined in chapter 2 of this dissertation. Women appear more likely to switch out of STEM, the field where they have been historically underrepresented. It's unclear, however, if this is shared across women in STEM, or if this behavior is driven by, say, women's worse performance in STEM fields. Therefore, in the following section, I investigate the relationship between switching behavior by field, first-year performance, and gender.

1.6 Switching, gender, and GPA

Section 1.5 shows that female switchers in STEM are more likely to graduate in non-STEM majors than their male counterparts. This suggests that men and women in different fields might be switching in different ways. This section further investigates that claim by documenting facts about observed switching patterns within specific fields by gender and GPA.[27] Across most fields, women switch majors at higher GPAs than men. The difference between the average GPA

27. For clarity, the fields "General Studies and Other," "Undecided or undeclared," and "Manufacturing, construction, repair, and transportation" are dropped from plots in this section. Note that "undecided and undeclared" is an ill-defined field. "General studies and other" may contain rather disparate majors within it; see A.1.7 for details. "Manufacturing, construction, repair, and transportation" is an extremely small, male-dominated field.

of female switchers and the average overall GPA in a field is small, suggesting that women may not interpret the signal sent by GPA in the same way that men do.

To start, the x-axis of figure 1.3 presents the fraction of students enrolled in a major during wave 1 of survey collection (at the end of the 2011-12 school year) who subsequently switch away by the end of wave 2 (in 2014). Similarly, the y-axis shows the fraction of students in each major in wave 2 that switched in. The size of the markers indicate the overall size of the field in the data during wave 1; fields with more students are larger, whereas fields with fewer students are smaller. The color of the markers indicate the gender ratios within fields; fields with fewer women are darker, whereas female-dominated fields are lighter.

The purpose of figure 1.3 is to determine whether there is a clear pattern in the fields that students switch from or to. For example, the social sciences, humanities, and communications all have relatively high fractions of students who switched in by wave 2. While the motivations behind these switches are not observable in the figure, these higher fractions are at minimum consistent with these fields having fewer institutional barriers to switching (such as course or GPA requirements). Next, consider the male-dominated STEM majors in the data: math, computer science, and engineering. These fields have slightly lower fractions of students who switch away, though no obvious patterns in switching patterns emerge in this figure.

I now consider whether switching patterns within fields differ by gender. To that end, figure 1.4a compares the fraction of male and female students who switch away from each field after wave 1 of data collection. Proximity to the 45 degree line implies that men and women are switching away at similar rates. Overall, many fields are somewhat close to the 45 degree line, indicating that similar fractions of men and women switch away after their first year. There are some exceptions: for example, relatively more women switch away from computer science than men. Additionally, there are a number of applied fields where men switch away more frequently than women, including the large, female-dominated fields of psychology, education, and health care. Figure 1.4b compares the average 2011-12 GPAs of men who switch majors to the average GPA of female switchers. A key takeaway from this figure is that the mean GPA of women who

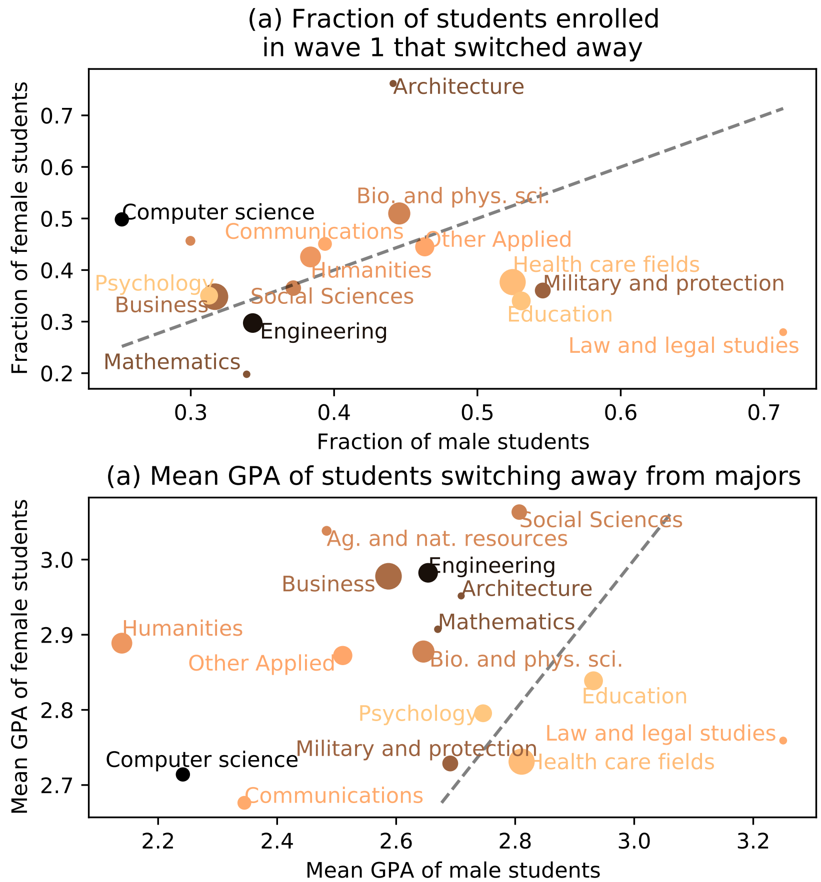

Figure 1.4. Figure (a) compares the fraction of men who switch away from each field after wave 1 of BPS:12/17 data collection to the fraction of female students that switch away. Figure (b) compares the mean GPA in 2011-12 of men who switch away from each field after wave 1 of BPS:12/17 data collection to the mean GPA of female students that switch away. The color of the markers indicate overall gender ratio in that field during wave 1; darker markers indicate male-dominated fields.

switch majors is, in general, higher than the mean GPA of male switchers. The fields with the highest share of women in the data (psychology, education, and health care fields) are exceptions to this rule; male and female switchers have similar GPAs. All of the STEM fields in the data lie well above the 45 degree line, as do all the male-dominated fields.

To better understand these patterns, figure 1.5 plots the the same information as figure 1.4b as solid circles, alongside the average GPA in each field as diamonds; these points are connected by dotted lines.[28] Diamonds in figure 1.5a represent the mean GPA of all male

28. In Figure 1.5, several additional fields are dropped for clarity: this includes Law and legal studies, Military

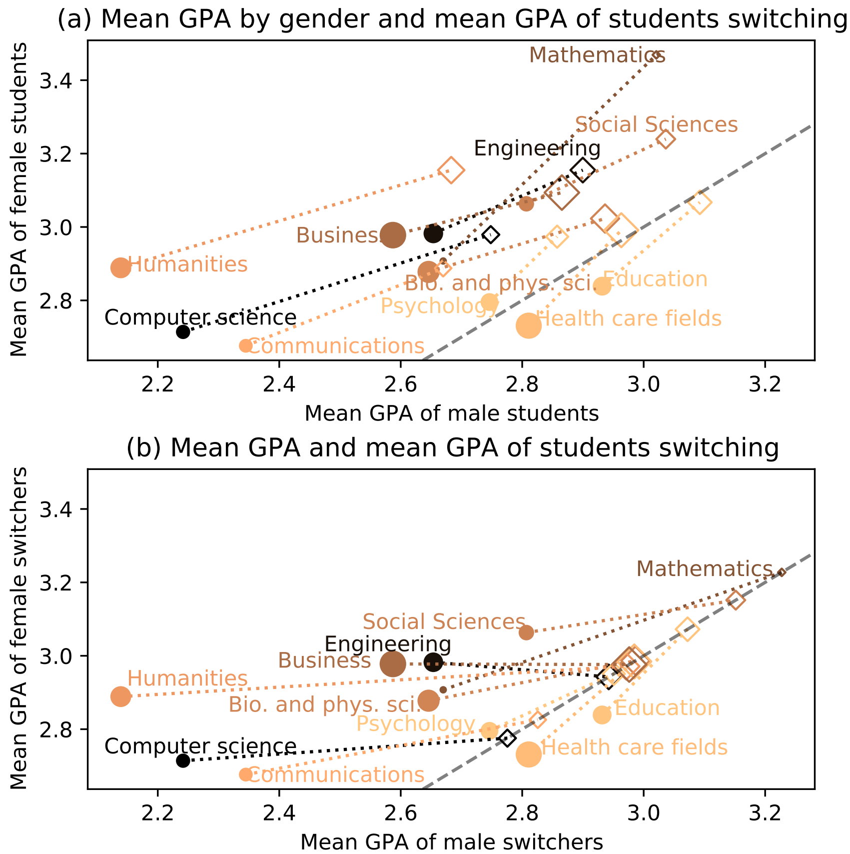

Figure 1.5. Circles in figures (a) and (b) represent the mean 2011-12 GPA of male and female students who subsequently switch fields. Diamonds in figure (a) represent the mean GPA of all male students and all female students on the x- and y-axes, respectively. Diamonds in figure (b) represent the overall mean GPA for all students. The color of the markers indicate overall gender ratio in that field during wave 1; darker markers indicate male-dominated fields.

students and all female students on the x- and y-axes, respectively. Diamonds in 1.5b represent the overall mean GPA for all students enrolled in a field during wave 1 of data collection (and therefore are plotted along the 45 degree line). One generally expects the dotted lines in figure 1.5 to have positive slopes; students choosing to switch away from a major would likely have lower GPAs compared to the average. If switching was fully preference based, or if GPA did not matter as a signal of ability, then dots and diamonds would be close together or would not have a clear pattern. Instead, the upward slope implies that students are taking into account the signal sent by GPA when choosing to change majors. The upward slopes in figure 1.5a imply that male

and protective services, Architecture, Agriculture and natural resources, and Other Applied.

and female switchers have lower GPAs than the average male and female student in each field. However, when one compares figure 1.5a to figure 1.5b, many lines noticeably flatten, due to the gender composition of each major. The small slopes indicate that the difference between the average GPA of female switchers and the average GPA in each field is small. The negative slope in engineering in figure 1.5b implies that the average GPA of women who change majors from engineering is actually higher than the overall average GPA in engineering.

The positive slopes in figure 1.5a suggest that switchers are taking the signal sent by GPA into account. However, the noticeable difference in slopes in figure 1.5b are consistent with men and women not interpreting these signals in the same way; women switchers may be comparing their GPA to the average woman in each field, not the overall average. This would be consistent with the theoretical model developed in chapter 2 of this dissertation. Women might look to information from their group type when assessing their beliefs about their abilities within fields. In the theoretical model developed in chapter 2, women use the outcomes of previous female cohorts to form their initial beliefs about their abilities. These differences in initial beliefs can result in men and women receiving the same signal behaving in different ways: women might switch away, while men persist. The relationship between gender, GPA, and switching patterns documented in figure 1.5 are consistent with this underlying mechanism.

1.7 Determinants of switching

The results presented in section 1.6 show that women switch out of fields at higher GPAs than men. Further, the results in section 1.5 suggest that men and women initially enrolled in STEM might have meaningfully different switching patterns. At this point, it's unclear if other factors besides gender, initial major choice, and first-year GPA enrollment are driving switching patterns. It's possible for institutional factors, preparedness for college, or familial and financial characteristics to be important drivers of major switching.

To that end, this section estimates the determinants of observed major switching across

Table 1.6. Probability of switching fields. Logit models (odds ratio reported), where the dependent variable indicates whether respondents' switch majors between the 2011-12 and 2014 BPS:12 PETS survey waves. Column (1) reports the results for the base sample. Column (2) reports results only for students last enrolled in a bachelor's. Column (3) reports results only for students enrolled in a bachelor's degree program in 2014. Column (4) reports results for students in the base sample who are not enrolled in private, for-profit institutions degree programs in 2014. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported. BPS:12 PETS Panel study weight used. Standard errors in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	(1)	(2)	(3)	(4)
Female	0.839	0.826	0.838	0.842
	(0.105)	(0.134)	(0.135)	(0.120)
STEM field 2011-12	0.488***	0.535***	0.561**	0.490***
	(0.084)	(0.099)	(0.105)	(0.088)
Female X STEM	2.380***	2.326***	2.141***	2.302***
	(0.523)	(0.523)	(0.464)	(0.538)
2011-2012 GPA	0.609***	0.584***	0.579***	0.614***
	(0.042)	(0.048)	(0.050)	(0.043)
Sample restrictions	None	BA only	Enroll in BA only	No for-profit
Percent correctly predicted	0.608	0.598	0.594	0.608
Goodness of fit F-stat	0.993	0.534	0.441	0.681
Goodness of fit p-value	0.447	0.849	0.911	0.725

survey waves using logit models, where the dependent variable indicates whether a student changes their listed major across the two survey waves.[29] I interact gender with initial STEM enrollment to evaluate whether women in STEM are more likely to change majors.[30] I control for first year GPA, as well as a host of institutional and personal characteristics.[31] Partial results

29. Students who are undeclared during the first survey wave are dropped. Students who are undeclared in the second wave only are considered field-switchers.

30. Due to survey design, it's not possible to interact gender with more granular aggregations of major choice and still construct bootstrap replicate standard errors.

31. Institutional characteristics include the following dummy variables: initial enrollment in a bachelor's degree program, initial enrollment in a public institution, initial enrollment in a 4-year institution, and the log of fall enrollment at the initial institution. Personal characteristics, family characteristics, and financial factors include the following dummy variables for each student: is an underrepresented minority; is first generation to go to college; has a mother with a bachelor's degree or higher; has a father with a bachelor's degree or higher; received a pell grant their first year; family received any federal benefits in 2011-12; submitted any financial aid application in 2011-12; reported zero adjusted gross income for their family on their financial aid application in 2011-12. The log of the families adjusted gross income in 2011-12 is also included. High school preparedness is measured by the SAT derived composite score, and a high school GPA dummy variable. The SAT derived composite score records either (1) the sum of the SAT verbal and math scores; or (2) the ACT composite score converted to an SAT combined

are presented in table 1.6; full results can be found in the appendix in table A.15. The baseline results are reported in column (1). Because students in this sample have a wide variety of enrollment and persistence patterns, columns (2) and (3) of table 1.6 restrict the sample. Column (2) only considers students who were last enrolled in a bachelor's degree program; this filters out students enrolled in associate's degree programs with bachelor's degree intentions. Column (3) further restricts the sample to only students currently enrolled in a bachelor's degree program, eliminating those who may not currently be enrolled in a degree program. Column (4) keeps only students enrolled in public or private non-profit institutions; students enrolled in private, for-profit institutions are dropped.

GPA during the 2011-12 school year appears to be an important determinant of field switching; students with higher GPAs are associated with a lower probability of switching majors, a result that is consistent with figure 1.5. Overall, students enrolled in STEM majors during the 2011-12 school year appear less likely to switch major than students enrolled in non-STEM majors when controlling for other factors. A key takeaway from table 1.6 is that women initially enrolled in STEM majors are more likely to switch majors than men initially enrolled in STEM majors, a result that holds across all samples.

Partial results for alternative specifications are presented in table 1.7, with full results available in the appendix table A.16. The baseline results are repeated in table 1.7 column (1) for reference. Ideally, I could include major GPA instead of overall GPA in this specification, but unfortunately that variable is not available in the BPS data. Instead, I do a robustness check using first year STEM GPA instead of overall GPA, restricting the sample to only include students with STEM GPAs. The results of this specification are reported in column (2). This variable is not ideal for a couple reasons: first, dropping students without a STEM GPA introduces selection into our sample; second, there are far more zeros in this variable than in the overall GPA variable. It's relatively common to find students who enrolled in only one or two STEM course and failed,

score. High school GPA is self-reported, and a dummy is used to indicate whether students self-report having an A or B average in high school.

Table 1.7. Probability of switching fields. Logit models (odds ratio reported), where the dependent variable indicates whether respondents' switch majors between the 2011-12 and 2014 BPS:12 PETS survey waves. Column (1) reports the results for the base sample. Column (2) reports results using student's 2011-12 STEM GPA, where students without STEM GPAs have been droped from this sample. Column (3) reports results for students who took some STEM credits during the 2011-12 school year. Column (4) reports results when interacting GPA, initial STEM GPA, and gender. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported. BPS:12 PETS Panel study weight used. Standard errors in parentheses; * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

	(1)	(2)	(3)	(4)
Female	0.839	0.780*	0.865	1.220
	(0.105)	(0.092)	(0.100)	(0.518)
STEM field 2011-12	0.488***	0.495***	0.505***	0.737
	(0.084)	(0.085)	(0.093)	(0.480)
Female X STEM	2.380***	2.055***	2.073**	2.193
	(0.523)	(0.426)	(0.468)	(2.464)
2011-2012 GPA	0.609***		0.582***	0.670***
	(0.042)		(0.046)	(0.070)
2011-12 STEM GPA		0.675***		
		(0.036)		
Female X GPA				0.879
				(0.134)
STEM X GPA				0.865
				(0.190)
Female X STEM X GPA				1.038
				(0.386)
Sample restrictions	None	STEM GPA	STEM cred	None
Percent correctly predicted	0.608	0.602	0.615	0.605
Goodness of fit F-stat	0.993	0.508	0.685	1.005
Goodness of fit p-value	0.447	0.867	0.722	0.438

receiving a zero GPA. Nevertheless, it provides a helpful check on the overall specification, particularly if we are interested in the patterns observed for STEM students. Similarly, table 1.7 column (3) presents the results where the sample is limited to students who have enrolled in some STEM credits. For both STEM GPA and STEM enrollments, the results are largely similar to the results in the baseline; women initially in STEM are more likely to change majors than men, controlling for first year GPA.

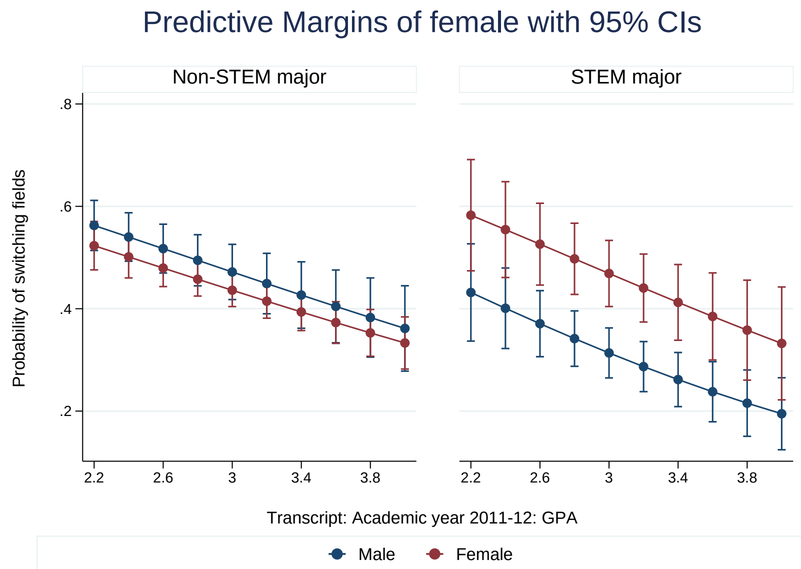

Figure 1.6. Predicted probability of men and women switching field between survey waves for students originally enrolled in STEM and non-STEM fields, calculated over first-year GPA levels. Based on the sample of students either (a) pursuing a bachelor's degree in 2014, or (b) enrolled in an associate's degree program, with the expectation of continuing on to a bachelor's degree program within two years.

Finally, table 1.7 column (4) further interacts gender and initial STEM enrollment with GPA. The predicted probabilities for switching fields for men and women in STEM and non-STEM fields are then plotted in figure 1.6. The coefficients for the interaction terms for women and STEM, and for women, STEM, and GPA are no longer significant. However, figure 1.6 still suggests meaningful differences in the way men and women respond to signals about performance. In non-STEM majors, the gender difference in field-switching across majors is small and statistically insignificant. However, there is a noticeable difference in the likelihood that men and women in STEM majors change fields. Although confidence bands overlap at low and high GPAs, this difference is significant between 2.6 and 3.4. We know from figure 1.5 that the mean GPA of female switchers across all fields falls within this range. The results in figure 1.7 provide meaningful evidence that women and men initially enrolled in STEM may be responding to signals sent about their performance in different ways. Thus, conditional on GPA, women initially enrolled in STEM appear more likely to change their majors than men.

1.8 Conclusion and discussion

In this paper, I provide the first nationwide longitudinal evidence of differential major switching by gender. Women are more likely to switch out of an initial STEM major. This relationship persists when I condition on GPA. Female switchers have higher GPAs than male switchers, and in the case of engineering, female switchers have GPAs that are higher than the overall average in that field. Given the demonstrated early interest in STEM, my evidence is suggestive that female students are more sensitive to early grades as a signal of their ability and likelihood of success in a STEM major, and require higher grades to remain in STEM fields. It is also suggestive that female students in a field compare their GPAs to those of other female students, rather than to the average student in that field, indicating the importance of same-gender role models. Regardless of the mechanism, this paper shows that part of the gender gap in major choice is driven by differential switching behavior, and efforts to retain early STEM students may reduce the gender gap.

These results are consistent with the theoretical model outlined in chapter 2 of this dissertation. I am able to establish some empirical facts about major switching that are otherwise not documented in the literature; for example, switchers take longer to graduate than non-switchers. This intuitive result is consistent with what we see in the theoretical model. The model also suggests that men and women receiving the same signal about performance may respond in different ways, assuming they form their initial priors about their abilities differently. This is consistent with what we see in the data, which shows that women are more likely to change their majors in STEM than men, controlling for other factors like GPA. Of course, other mechanisms, such as gendered differences in ex ante preference accuracy, may also play a role. Nevertheless, the empirical results documented in this chapter, alongside the theoretical framework outlined in chapter 2, document a mechanism otherwise absent in the college major choice literature.

While I cannot identify the mechanisms that cause female students to switch out of STEM majors, documenting the phenomenon in a nationwide data set is a crucial first step.

Ideally, these results can be examined in more detail in the future. In particular, major switching should be studied at a more granular level than simply STEM and non-STEM. My analysis was somewhat limited by survey design and the availability of computational resources. Future work could ideally take into account the switching behavior of men and women for more detailed fields. Alternatively, one could build on this analysis by considering intersectionality constraints or results outside the gender binary.

Additionally, I hope that some of the college choice literature focuses on the experiences of "non-traditional" students. Most analyses focus solely on the behaviors of students at elite, 4-year institutions, already enrolled in bachelor's degrees. However, students following this "traditional" route are, in general, not representative of the average degree seeking student. There is a wide variety of educational experiences that are lost when we solely focus on this population, and future work should take those experiences into account.

Appendix

A.1.1 Survey weights

The student interview cross-sectional weight (WTA000) includes the approximately 22,530 students who were first-time beginning students in 2011-12 and BPS:12/17 interview respondents. The student interview panel weight (WTB000) includes the approximately approximately 19,840 students who were BPS:12/14 and BPS:12/17 interview respondents (Bryan, Cooney, and Elliott 2019).

The BPS:12 PETS respondent weight is given by WTC000. These are the approximately 25,900 students for whom a transcript was received for at least one institution. The PETS panel weight is given by WTD000 and applies to students that were BPS:12/14 respondents, BPS:12/17 respondents, and PETS respondents, approximately 15,350 students (Bryan et al. 2020). The PETS panel weight is the weight primarily used in this analysis.

Information about the weighting procedure used in BPS:12 PETS can be found in Bryan et al. (2020), Chapter 4.[32]

A.1.2 Additional BPS summary statistics

Table A.1. Type of first institution 2011-12 (percentages). Control and level of first institution (IPEDS sector) 2011-12, after applying wave 1 filter. Bootstrap replicate standard errors for proportions are in parentheses. ‡ indicates that reporting standards are not met. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Public less-than-2-year	‡	‡	‡
Public 4-year	40.49	36.56	38.27
	(1.93)	(1.45)	(1.38)
Public 2-year	34.72	33.32	33.93
	(1.57)	(1.46)	(1.24)
Private nonprofit 4-year	17.72	21.30	19.75
	(0.70)	(0.66)	(0.39)
Private nonprofit 2-year	0.59	0.34	0.45
	(0.24)	(0.16)	(0.18)
Private for profit less-than-2-year	‡	‡	‡
Private for profit 4-year	5.65	6.52	6.14
	(0.98)	(0.33)	(0.49)
Private for profit 2-year	0.81	1.60	1.26
	(0.43)	(0.77)	(0.56)

32. Note that details on how Bootstrap replicate weights are calculated is outlined in Section 4.2 Bryan et al. (2020).

Table A.2. Persistence anywhere through June 2014 (percentages). Persistence anywhere through wave 2 of data collection (June 2014), after applying wave 1 filter. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Degree program when last enrolled in wave 2			
Undergraduate certificate/diploma	2.88	3.45	3.21
	(0.46)	(0.46)	(0.32)
Bachelor's degree	57.21	58.99	58.23
	(1.82)	(1.86)	(1.59)
Not in a degree program or others	2.82	3.61	3.27
	(0.41)	(0.51)	(0.36)
Associate's degree no Bachelor's intentions in 2 years	14.08	12.99	13.45
	(0.85)	(0.80)	(0.56)
Associate's degree with Bachelor's intentions in 2 years	23.01	20.96	21.83
	(1.93)	(1.69)	(1.50)
Persistence anywhere through wave 2			
Attained, still enrolled	5.26	6.68	6.06
	(0.45)	(0.57)	(0.37)
Attained, not enrolled	3.46	4.65	4.13
	(1.20)	(0.49)	(0.55)
No degree, still enrolled	64.31	67.76	66.26
	(1.56)	(1.46)	(1.14)
No degree, not enrolled	26.97	20.91	23.54
	(1.34)	(1.39)	(0.97)

Table A.3. Persistence for students initially enrolled in bachelor's degrees. Persistence anywhere through wave 2 of data collection (June 2014) for those enrolled in bachelor's degree programs in wave 1 (2011-12). ‡ represents less than 1%. Data are weighted by BPS:12 PETS Panel study weight.

Attained associate's degree; enrolled in associate's degree	‡
Attained associate's degree; enrolled in bachelor's degree	1.26%
Attained associate's degree; enrolled in not in a degree program or others	‡
Attained associate's degree; enrolled in undergraduate certificate/diploma	‡
Attained associate's degree; not enrolled; last enrolled in associate's degree	‡
Attained associate's degree; not enrolled; last enrolled in bachelor's degree	‡
Attained bachelor's degree; not enrolled	‡
Attained bachelor's degree; not enrolled; last enrolled in associate's degree	‡
Attained bachelor's degree; not enrolled; last enrolled in bachelor's degree	2.62%
Attained bachelor's degree; not enrolled; last enrolled in not in a degree program or others	‡
Attained bachelor's degree; not enrolled; last enrolled in undergraduate certificate/diploma	‡
Attained certificate; enrolled in associate's degree	‡
Attained certificate; enrolled in bachelor's degree	‡
Attained certificate; enrolled in not in a degree program or others	‡
Attained certificate; enrolled in undergraduate certificate/diploma	‡
Attained certificate; not enrolled	‡
Attained certificate; not enrolled; last enrolled in associate's degree	‡
Attained certificate; not enrolled; last enrolled in bachelor's degree	‡
Attained certificate; not enrolled; last enrolled in undergraduate certificate/diploma	‡
Enrolled in associate's degree; no degree	7.07%
Enrolled in bachelor's degree; no degree	65.18%
Enrolled in not in a degree program or others; no degree	3.05%
Enrolled in undergraduate certificate/diploma; no degree	1.09%
No degree; last enrolled in associate's degree	2.15%
No degree; last enrolled in bachelor's degree	7.47%
No degree; last enrolled in not in a degree program or others	‡
No degree; last enrolled in undergraduate certificate/diploma	‡
No degree; not enrolled	6.63%
All	100.00%

Table A.4. Persistence for students initially enrolled in associate's degrees. Persistence anywhere through wave 2 of data collection (June 2014) for those in associate's degree programs in wave 1 (2011-12) with the intention of continuing on to a bachelor's degree program in five years. ‡ represents less than 1%. Data are weighted by BPS:12 PETS Panel study weight.

Attained associate's degree; enrolled in associate's degree	5.57%
Attained associate's degree; enrolled in bachelor's degree	6.21%
Attained associate's degree; enrolled in not in a degree program or others	‡
Attained associate's degree; enrolled in undergraduate certificate/diploma	‡
Attained associate's degree; not enrolled	1.07%
Attained associate's degree; not enrolled; last enrolled in associate's degree	5.72%
Attained associate's degree; not enrolled; last enrolled in bachelor's degree	‡
Attained associate's degree; not enrolled; last enrolled in not in a degree program or others	‡
Attained associate's degree; not enrolled; last enrolled in undergraduate certificate/diploma	‡
Attained bachelor's degree; not enrolled; last enrolled in associate's degree	‡
Attained bachelor's degree; not enrolled; last enrolled in bachelor's degree	‡
Attained certificate; enrolled in associate's degree	1.38%
Attained certificate; enrolled in bachelor's degree	‡
Attained certificate; enrolled in not in a degree program or others	‡
Attained certificate; enrolled in undergraduate certificate/diploma	‡
Attained certificate; not enrolled	‡
Attained certificate; not enrolled; last enrolled in associate's degree	‡
Attained certificate; not enrolled; last enrolled in bachelor's degree	‡
Attained certificate; not enrolled; last enrolled in not in a degree program or others	‡
Attained certificate; not enrolled; last enrolled in undergraduate certificate/diploma	‡
Enrolled in associate's degree; no degree	28.74%
Enrolled in bachelor's degree; no degree	8.41%
Enrolled in not in a degree program or others; no degree	‡
Enrolled in undergraduate certificate/diploma; no degree	2.06%
No degree; last enrolled in associate's degree	19.54%
No degree; last enrolled in bachelor's degree	1.17%
No degree; last enrolled in not in a degree program or others	‡
No degree; last enrolled in undergraduate certificate/diploma	1.01%
No degree; not enrolled	13.48%
All	100.00%

Table A.5. Undergraduate field of study 2011-12 (percentages). BPS proportion of data by gender: Field of study: undergraduate (23 categories) 2011-12. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Undecided or Undeclared	5.38	5.23	5.29
	(0.89)	(0.55)	(0.53)
Computer and information sciences	6.69	0.65	3.23
	(1.44)	(0.13)	(0.67)
Engineering and engineering technology	13.03	2.00	6.70
	(1.05)	(0.27)	(0.49)
Biological and physical science, science tech	8.82	8.77	8.79
	(0.82)	(0.57)	(0.47)
Mathematics	0.86	0.44	0.62
	(0.19)	(0.17)	(0.12)
Agriculture and natural resources	1.26	1.33	1.30
	(0.28)	(0.24)	(0.19)
General studies and other	10.87	8.45	9.48
	(1.16)	(0.70)	(0.67)
Social sciences	4.01	3.85	3.92
	(0.69)	(0.41)	(0.37)
Psychology	2.45	7.58	5.39
	(0.40)	(0.84)	(0.50)
Humanities	5.15	8.05	6.81
	(0.59)	(0.80)	(0.51)
History	1.40	0.49	0.88
	(0.58)	(0.14)	(0.27)
Personal and consumer services	2.67	2.17	2.39
	(0.36)	(0.37)	(0.29)
Manufacturing, construction, repair, transportation	1.19	0.55	0.82
	(0.40)	(0.54)	(0.36)
Military technology and protective services	5.26	3.29	4.13
	(1.27)	(0.54)	(0.75)
Health care fields	6.94	17.00	12.71
	(0.96)	(1.17)	(0.71)
Business	15.76	11.20	13.15
	(1.03)	(0.63)	(0.49)
Education	2.98	8.49	6.14
	(0.38)	(0.70)	(0.46)
Architecture	0.76	0.38	0.54
	(0.19)	(0.20)	(0.15)
Communications	2.07	3.70	3.01
	(0.35)	(0.50)	(0.32)
Public administration and human services	0.33	2.77	1.73

Table A.5. Undergraduate field of study 2011-12 (percentages) (continued)

	Male	Female	All
	(0.13)	(0.69)	(0.40)
Design and applied arts	1.44	2.05	1.79
	(1.56)	(1.00)	(1.04)
Law and legal studies	0.49	0.88	0.72
	(0.20)	(0.22)	(0.15)
Theology and religious vocations	0.20	0.68	0.48
	(0.07)	(0.40)	(0.24)

BPS:12/17 and BPS:12 PETS proportion of data by gender: Field of study: undergraduate (23 categories) 2011-12. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

A.1.3 Formal major declaration

Many American schools do not require students to declare their major at the start of college (Patnaik, Wiswall, and Zafar 2020). If students are unlikely to declare their major, then switching behaviors may be unobservable. To that end, table A.6 summarizes the proportion of students that have declared their major during the first wave of data collection. Note that the following results only use the wave 1 filter, described in section 1.3.2. By the end of the 2011-12 school year, most students have declared a major, or at least decided on their field of study.

To better understand the potential determinants of declaring a major, table A.7 estimates conditional probability models, where the dependent variable indicates whether students have formally declared a major. Note again that the sample here is defined by applying the wave 1 filter only. The likelihood of declaring one's major is not significantly impacted by student characteristics, such as being a female, underrepresented minority, or first-generation student. The results suggest that SAT score may have a slight, negative impact on the likelihood of declaring one's major.[33] The F-statistics are significant at the 1% and 5% level for the logit and probit models, respectively. Declaring one's major by the end of first year appears relatively common across student and institutional characteristics, and doesn't appear to be a major

33. Note that the variable of interest, derived SAT, is derived as either the sum of SAT verbal and math scores or the ACT composite score converted to an estimated SAT combined score.

Table A.6. Initial major decisions by gender (percentages). The proportion of students who have formally declared their major field of study, and the proportion of students with a focus on STEM fields. sSample includes wave 1 filter only, i.e. students who are either (a) enrolled in a bachelor's degree in wave 1 of data collection (2011-12), or (b) enrolled in an associate's degree programs in wave 1 with the intention of continuing on to a bachelor's degree program in five years. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Formally declared major field of study 2012			
Formally declared	85.34	85.62	85.50
	(1.22)	(0.84)	(0.71)
Decided but not formally declared	9.25	9.28	9.26
	(0.97)	(0.74)	(0.59)
Undecided	5.42	5.10	5.24
	(0.73)	(0.54)	(0.48)
Major field of study with a focus on STEM fields 2011-12			
Math/Computer/Sciences/Engineering/Technologies	28.30	11.16	18.61
	(1.58)	(0.56)	(0.82)
Social/behavioral sciences	5.26	9.88	7.87
	(0.58)	(0.98)	(0.62)
Non-STEM field	61.03	73.86	68.28
	(1.70)	(1.03)	(1.02)
Undecided or not in a degree program	5.42	5.10	5.24
	(0.73)	(0.54)	(0.48)

Table A.7. Probability of declaring a major by wave 1. Conditional probability models where the dependent variable indicates whether respondents' have formally declared a major by the end of the 2011-12 school year. Based on the sample of students either (a) pursuing a bachelor's degree in 2011-12, or (b) enrolled in an associate's degree program, with the expectation of continuing on to a bachelor's degree program within five years. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported.

	Logit	Probit
Female	-0.010	-0.005
	(0.145)	(0.080)
Underrepresented minority	-0.075	-0.034
	(0.185)	(0.099)
First generation student	0.001	-0.002
	(0.196)	(0.104)
Enrolled in bachelor's	0.325	0.169
	(0.619)	(0.331)
2011-12 Pell grant recipient	0.138	0.074
	(0.141)	(0.076)
First-year STEM enrollment	0.082	0.045
	(0.200)	(0.113)
GPA	0.099	0.051
	(0.098)	(0.053)
Derived SAT	-0.001**	-0.001**
	(0.000)	(0.000)
Public institution	0.169	0.114
	(0.173)	(0.097)
4-year institution	-0.620	-0.320
	(0.689)	(0.369)
Log tuition and fees 2011-12	-0.181	-0.090
	(0.106)	(0.056)
Log fall enrollment 2011-12	-0.098	-0.062
	(0.077)	(0.043)
Doctorate offered at first institution	0.283	0.150
	(0.201)	(0.109)
Percent correctly predicted	0.489	0.486
Goodness of fit F-stat	2.628	1.940
Goodness of fit p-value	0.007	0.049

BPS:12 PETS Panel study weight used.

Standard errors in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A.8. Share of women in STEM. Table shows the share of women initially enrolled as STEM and non-STEM majors in BPS:12 PETS data. Data are weighted by BPS:12 PETS Panel study weight.

	Wave 1: 2011-12	Wave 2: 2014
Women initially in STEM	0.366	0.371
Women initially in non-STEM	0.632	0.625

Table A.9. Changing out of STEM/non-STEM. Table shows the share of initial STEM and non-STEM majors who change or switch fields, by gender in BPS:12 PETS data. Data are weighted by BPS:12 PETS Panel study weight.

Initial STEM major	**Stay in STEM**	**Change to non-STEM**
Male	0.74	0.26
Female	0.58	0.42
Initial non-STEM major	**Stay in non-STEM**	**Change into STEM**
Male	0.92	0.08
Female	0.94	0.06

empirical concern for this analysis.

A.1.4 Effect of switching on STEM enrollment

Section 1.5 motivates switching as an outcome variable by summarizing the graduation outcomes for switchers and non-switchers. Those results suggest that women change out of STEM majors at a higher rate than men. However, those results are conditional on graduation outcomes. To that end, this section briefly summarizes the gender ratio in STEM and non-STEM fields across time, and considers whether men and women change out of these fields at different rates. It's important to emphasize that this section is not looking at "switching," according to primary definition outlined in section 1.4.1. Instead, this section considers the overall composition of men and women in aggregated STEM and non-STEM fields at the two different waves of data collection. As a result, comparisons between the decompositions in tables A.9 and A.10 and, say, the switching patterns documented in section 1.6 should be done carefully.

Table A.8 shows the fraction of students that are female in STEM and non-STEM fields

Table A.10. Changing into STEM/non-STEM Table shows the share of wave 2 STEM and non-STEM majors by gender who were either enrolled in the the same major aggregation in wave 1 or changed in. Data are weighted by BPS:12 PETS Panel study weight.

Wave 2 STEM major	**STEM in wave 1**	**Changed in from non-STEM**
Male	0.81	0.19
Female	0.62	0.38
Wave 2 non-STEM major	**Non-STEM in wave 1**	**Changed in from STEM**
Male	0.88	0.12
Female	0.93	0.07

across waves 1 and 2 of data collection. Overall, the fractions remain largely unchanged over time. A similar majority of male and female students initially enrolled in non-STEM fields in wave 1 stay in a non-STEM field by wave 2. However, table A.9 suggests that STEM enrollment patterns between men and women are not the same. While 26% of men initially in STEM change into a non-STEM major, 42% of women do. Women appear more likely to change out of STEM fields and into non-STEM fields then men.

Finally, table A.10 lists the gender composition in STEM and non-STEM fields in wave 2. Non-STEM majors in wave 2 have generally persisted from wave 1, a pattern that holds for men and women. Overall persistence across waves is lower in STEM fields than in non-STEM fields. However, there is a marked difference in persistence for men and women in STEM: while 81% of male STEM majors in wave 2 persisted from wave 2, only 62% of the women in STEM in wave 2 were in STEM in wave 1. These results suggest that there may be some differences in switching behaviors for STEM and non-STEM fields, and for men and women.

It is worth contrasting the results from figure 1.3 to the results in tables A.8 and A.9. For example, table A.8 only considers the total fraction of STEM students who change into a non-STEM major, and vice versa. Those results suggest that changing from non-STEM into STEM is uncommon, which may lead one to think major changes are uncommon. Figure 1.3 suggests that major switching is indeed common in non-STEM fields. This is because only considering changes to and from aggregated STEM fields tends to undercount major switching overall.

Table A.11. Self-reported major switching during wave 1. Self-reported major switching during students' first year; proportions by gender. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
None	65.84	66.42	66.17
	(1.52)	(0.95)	(0.87)
One	15.84	15.24	15.50
	(1.10)	(0.64)	(0.62)
More than one	3.66	3.96	3.83
	(0.48)	(0.42)	(0.28)
No switch; undecided or undeclared major	5.42	5.10	5.24
	(0.73)	(0.54)	(0.48)
No switch; Decided but not formally declared major	9.25	9.28	9.26
	(0.97)	(0.74)	(0.59)

A.1.5 Switching within a student's first year

Students may formally switch their major within their first year. Because the BPS: 12/17 data are collected at the end of a students' first year, this behavior is not always directly observable in the data.[34] However, the survey does ask students about how many times they have formally switched their major during their first year.[35] Students' responses are summarized in table A.11. Switching majors during one's first year is relatively common; almost 20% of the sample formally change major during the 2011-12 school year. However, this breakdown is limited in two ways. First it only applies to students' formally declared major. If a student began the 2011-12 school year knowing their intended field of study without declaring it, and subsequently declared a different major, they are not listed as having switched their major. Second, original fields are unobserved. Therefore, we don't know if switching is more common in certain fields.

To better understand the potential determinants of major switching during a student's first

34. Specifically, these data are not observable in the BPS: 12 PETS derived variables. As noted in section 1.3.1, it is sometimes possible to construct alternative variables using the underlying restricted-use data. Because the BPS:12 /17 PETS data does include transcript information, it may be possible to construct an alternative measure of first-year switching behavior.

35. Specifically, the variable asks the student: "How many times have you formally changed your major at the first institution? "

year, table A.12 estimates conditional probability models, where the dependent variable indicates whether a student reports formally changing their major during their first year. Students enrolled in a public institution (compared to a private one) appear more likely to report switching their major during their first year. This can perhaps be explained by the fact that public institutions tend to rely on GPA restrictions for certain popular majors, causing students who fall below a threshold to switch majors (Bleemer and Mehta 2021). Overall, few variables are significant. These results suggest that major switching is common across a range of observable characteristics.

A.1.6 Self-reported switching across survey waves

There are different ways to define field switching across survey waves. I focus on two different definitions: (1) a student has switched between fields if their transcript-verified majors do not match across survey waves; and (2) a student self-reports switching majors in the second survey wave.[36] The proportion of students switching field according to either definition are reported in table A.13. The top half of the table simply compares the students reported in major from wave 1 of the data to the reported field in wave 2. Students who have "switched fields" are those whose majors do not match across the two waves of data. According to this measure, approximately half of students switch their major between the first and second wave of data.[37] This number is higher if one counts those switching from/to undeclared majors. An advantage of this definition is that it is verified by student transcripts. A disadvantage is that it does not capture switching during one's first year.

Alternatively, the bottom half of the table A.13 is based on a students' self-reported number of major changes. Students are asked in follow-up interviews about the number of times they have changed their major, and students who have "switched fields" in the table are those who have reported changing their major at least one time. In theory, the self-reported number of

36. Specifically, the variable "majchgba14" indicates the number of times the respondent formally changed his or her bachelor's degree major at any institution through June 2014, according to the student's first follow-up interview.

37. Note that this result is consistent with descriptive results from Arcidiacono (2004), who also define switching across survey waves using the NLSY79 data.

Table A.12. Probability of changing major during 2011-12 Conditional probability models, where the dependent variable indicates whether respondents' formally changing majors at least once during the 2011-12 school year, after applying wave 1 filter. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported.

	Logit	Probit
Female	0.054	0.030
	(0.131)	(0.077)
STEM	-0.280	-0.161
	(0.209)	(0.119)
Female X STEM	0.119	0.070
	(0.275)	(0.158)
Underrepresented minority	-0.009	-0.005
	(0.125)	(0.072)
First generation student	0.002	0.002
	(0.149)	(0.087)
2011-12 Pell grant recipient	0.025	0.015
	(0.100)	(0.058)
Enrolled in bachelor's	-0.501	-0.279
	(0.366)	(0.211)
GPA	-0.037	-0.022
	(0.080)	(0.047)
Derived SAT	0.001^{*}	0.000^{*}
	(0.000)	(0.000)
Public institution	0.554^{***}	0.316^{***}
	(0.143)	(0.082)
4-year institution	0.205	0.114
	(0.340)	(0.197)
Doctorate offered at first institution	0.234	0.132
	(0.165)	(0.092)
Selectivity of first institution	0.043	0.024
	(0.166)	(0.094)
Log fall enrollment 2011-12	-0.081	-0.045
	(0.056)	(0.032)
Constant	-1.402^{**}	-0.849^{**}
	(0.531)	(0.308)
Percent correctly predicted	0.524	0.522
Goodness of fit F-stat	0.327	0.597
Goodness of fit p-value	0.965	0.798

BPS:12 PETS Panel study weight used.

Standard errors in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table A.13. Observed and self-reported field switching across waves (percentages) BPS:12/17 and BPS:12 PETS proportion of data by gender: Field switch between waves 1 and 2; and Self-reported major changes anywhere through wave 2. Bootstrap replicate standard errors for proportions are in parentheses. Data are weighted by BPS:12 PETS Panel study weight.

	Male	Female	All
Field switch between waves 1 and 2			
Do not switch fields	35.65	37.19	36.53
	(2.40)	(1.28)	(1.19)
Switch fields	52.33	52.41	52.38
	(2.01)	(1.37)	(1.09)
Undeclared wave 1, declared wave 2	4.40	4.85	4.66
	(0.85)	(0.56)	(0.48)
Declared wave 1, undeclared wave 2	6.61	5.09	5.74
	(1.13)	(0.64)	(0.66)
Undeclared waves 1 and 2	0.98	0.38	0.64
	(0.37)	(0.10)	(0.17)
Missing wave 2	0.03	0.08	0.05
	(0.02)	(0.03)	(0.02)
Self-reported major changes anywhere through wave 2			
Do not switch fields	60.43	63.57	62.23
	(2.07)	(1.40)	(1.32)
Switch fields	26.35	25.16	25.67
	(2.29)	(1.26)	(1.46)
Undeclared wave 1, declared wave 2	4.40	4.85	4.66
	(0.85)	(0.56)	(0.48)
Declared wave 1, undeclared wave 2	6.61	5.09	5.74
	(1.13)	(0.64)	(0.66)
Undeclared waves 1 and 2	0.98	0.38	0.64
	(0.37)	(0.10)	(0.17)
Missing wave 2	0.03	0.08	0.05
	(0.02)	(0.03)	(0.02)
Missing self-reported major changes	1.21	0.87	1.01
	(0.27)	(0.24)	(0.16)

major changes could reflect the major-switching that is unobservable in the data (for example, field-switching that occurs before the first round of data is collected in 2011-12). In practice, the self-reported fraction of student's changing majors is lower than what we would assume based on reported student majors. Comparing the top and bottom halves of table A.13 suggests that student's are inclined to under-report the number of times they switch their major. For that reason, I'll focus on the observed field switching between waves 1 and 2 of the BPS:12/17 for the remainder of this analysis.

A.1.7 Comparison of IPEDS CIP labels and BPS major classifications

Section 1.6 compares switching behaviors across several fields of study. To make these figures more readable, I aggregated or dropped several major categories. Specifically, I built a concordance across the major categories in BPS data and CIP labels in IPEDS data. In theory, several variables in the BPS data aggregate fields of study into 23 categories that should correspond to the U.S. Department of Education's Classification of Instructional Programs, 2010 edition (CIP 2010), the major classification used in IPEDS data. In practice, some discrepancies between the field of study in BPS data and CIP codes exist, which will be detailed below.

The 19 major categories can be found in table A.14. Note that only the wave 1 filter has been applied to these data, which results in minor discrepancies from table 1.3. This concordance allows me to directly compare the ratios of women to men studying the fields in the BPS data to the gender ratios in degree completions in IPEDS data. Both ratios are plotted in figure A.1. Note that BPS and IPEDS data are substantively different; IPEDS reflects gender composition in degree attainment, while the gender ratios in 1.3 reflects initial major choice. Further, IPEDS data capture the universe of students completing bachelor's degrees only, while table A.14 focuses on a sample of first-time beginning students, which includes those pursuing associate's degrees with the intention of continuing on to a bachelor's. The existence of an "Undecided or undeclared" major category in BPS data also skews the data, and makes comparability between tables A.14 and IPEDS data difficult. Despite these caveats, the gender ratios are generally, suggesting patterns observed through 2011 in IPEDS data persist into the years covered by the BPS sample.[38] Fields with low gender ratios of degree completion in the 2011-2012 school

38. There are two exceptions to this general agreement across BPS and IPEDS data: first, fields with high female-to-male ratios, such as psychology and health care fields; while this could be attributed to more women switching into or away from these fields over the course of their studies, it could also easily be attributed to differences in scale. Because these fields are strongly female-dominated in both data sources, and because there is clear concordance across the IPEDS CIP labels and the BPS field classifications, this should not merit much concern. The "General

Table A.14. Alternative major categories. Undergraduate field of study in 2011-12 (15 categories; author derived variable). Ratio of women to men in rightmost column. Bootstrap replicate standard errors in parentheses. Data weighted by BPS:12 PETS Panel study weight.

	Male	Female	All	Ratio
Undecided or undeclared	5.42%	5.10%	5.24%	1.23
	(0.73%)	(0.54%)	(0.48%)	
Computer and information sciences	7.44%	0.82%	3.70%	0.14
	(1.31%)	(0.14%)	(0.60%)	
Engineering and engineering technology	12.05%	1.82%	6.26%	0.2
	(0.91%)	(0.24%)	(0.41%)	
Bio. and physical science, science tech	7.22%	7.58%	7.42%	1.36
	(0.61%)	(0.48%)	(0.38%)	
Mathematics	0.78%	0.35%	0.54%	0.59
	(0.18%)	(0.13%)	(0.10%)	
Agriculture and natural resources	1.43%	1.09%	1.24%	1.0
	(0.28%)	(0.17%)	(0.17%)	
General studies and other	11.94%	9.92%	10.80%	1.08
	(0.99%)	(0.77%)	(0.67%)	
Social Sciences	3.33%	3.16%	3.23%	1.23
	(0.49%)	(0.32%)	(0.29%)	
Psychology	2.18%	7.27%	5.06%	4.33
	(0.35%)	(0.98%)	(0.59%)	
Humanities	5.79%	7.29%	6.64%	1.63
	(0.58%)	(0.62%)	(0.42%)	
Manufacturing, transportation, etc.	2.57%	0.43%	1.36%	0.21
	(0.53%)	(0.40%)	(0.33%)	
Military and protective services	5.78%	3.47%	4.48%	0.78
	(1.01%)	(0.47%)	(0.61%)	
Health care fields	7.26%	19.75%	14.32%	3.53
	(1.08%)	(1.05%)	(0.71%)	
Business	15.81%	10.88%	13.02%	0.9
	(1.01%)	(0.57%)	(0.52%)	
Education	2.54%	7.95%	5.60%	4.05
	(0.32%)	(0.60%)	(0.37%)	
Architecture	0.63%	0.39%	0.49%	0.81
	(0.14%)	(0.16%)	(0.12%)	
Communications	2.10%	3.06%	2.64%	1.9
	(0.32%)	(0.40%)	(0.27%)	
Law and legal studies	0.47%	1.28%	0.92%	3.6
	(0.15%)	(0.41%)	(0.25%)	
Other Applied	5.27%	8.39%	7.03%	2.07
	(1.22%)	(0.86%)	(0.61%)	

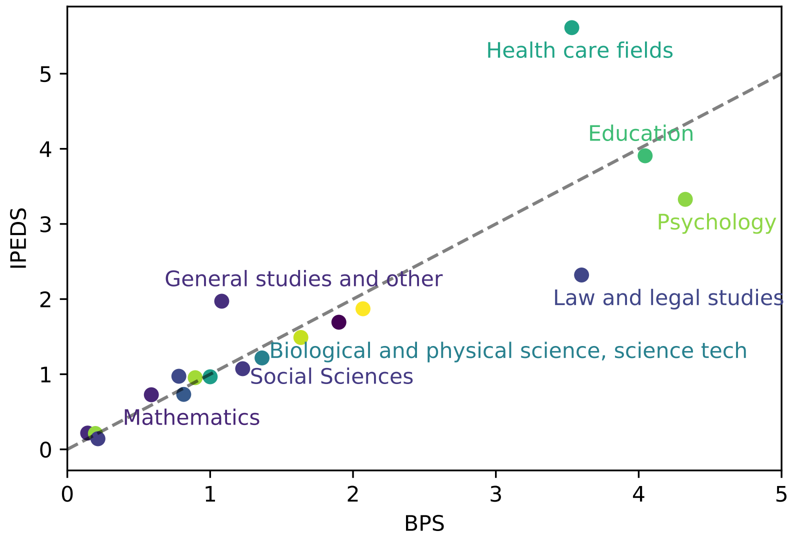

Figure A.1. Comparison of the ratio of women to men across two data sources. BPS:12/17 data calculates the gender ratio amongst a nationally representative sample of first-time beginning students pursuing either a bachelor's degree, or an associate's degree with intentions to complete a bachelor's degree. IPEDS data reflect the universe of bachelor's degree completions in 2011.

year (such as computer sciences, engineering, and mathematics) all have low gender ratios for students beginning their studies.

Discrepancies between BPS major categories and IPEDS CIP labels

For example, the BPS 12:/17 the variable "MAJORS23" lists the field of study that should map to CIP 2010 values, and the variable "MAJORS" provides a more detailed classification of students' field of study in 2011-12. However, when the detailed field of study equals "Family, consumer, and human sciences" (MAJORS == 13), BPS claims this variable may belong to a number of CIP categories, including either Personal and Consumer Services (MAJORS23 == 11) or Public administration and human services (MAJORS23 == 19). Yet "family and consumer sciences" is a 2010 CIP code of its own in IPEDS data, so one may expect it to either (a) exactly map to a "family and consumer science" value in the BPS classification, or (b) map to either the Personal and Consumer Services or Public administration and human services categories (but not

studies and other" category may differ slightly between the IPEDS and BPS data; details can be found in section A.1.7

both). The approach I adopt to is to aggregate fields with these types of discrepancies into the "Other Applied" Category, as that allows me to compare fields across IPEDS data and BPS data.

Additionally, the interdisciplinary CIP 2010 code (CIP == 30) does not clearly map to any given field in the BPS data. If we compare the fields where the detailed major in BPS 12:17 data is given as Multi/interdisciplinary study (MAJORS==21), then the corresponding CIP code may be in any number of fields (Biological and physical science [MAJORS23 == 3], General studies and other [MAJORS23 == 6], and Social Sciences [MAJORS23 == 7], among others). I assume that the CIP 2010 code for general and interdisciplinary studies maps to the General studies and other CIP code. This means that the general studies and other category will not be particularly comparable across IPEDS data and BPS data. Further, this may result in some discrepancies across the biological and physical sciences CIP code, and the social sciences.

A.1.8 Full logit results

Table A.15. Probability of switching fields. Logit models (odds ratio reported), where the dependent variable indicates whether respondents' switch majors between the 2011-12 and 2014 BPS:12 PETS survey waves. Column (1) reports the results for the base sample of students either (a) pursuing a bachelor's degree in 2014, or (b) enrolled in an an associate's degree program, with the expectation of continuing on to a bachelor's degree program within two years. Column (2) reports results only for students last enrolled in a bachelor's. Column (3) reports results only for students enrolled in a bachelor's degree program in 2014. Column (4) reports results for students in the base sample who are not enrolled in private, for-profit institutions degree programs in 2014. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported. BPS:12 PETS Panel study weight used. Standard errors in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	(1)	(2)	(3)	(4)
Female	0.839	0.826	0.838	0.842
	(0.105)	(0.134)	(0.135)	(0.120)
STEM field 2011-12	0.488***	0.535***	0.561**	0.490***
	(0.084)	(0.099)	(0.105)	(0.088)
Female X STEM	2.380***	2.326***	2.141***	2.302***
	(0.523)	(0.523)	(0.464)	(0.538)
2011-2012 GPA	0.609***	0.584***	0.579***	0.614***
	(0.042)	(0.048)	(0.050)	(0.043)
Enrolled in bachelor's	0.795	0.815	0.989	0.635
	(0.386)	(0.379)	(0.487)	(0.357)
Underrepresented minority	1.154	1.269	1.275	1.173
	(0.196)	(0.258)	(0.270)	(0.203)
First generation student	0.880	0.823	0.862	0.881
	(0.145)	(0.139)	(0.145)	(0.146)
Mother: bachelor's or higher	0.808*	0.776*	0.794*	0.818*
	(0.083)	(0.081)	(0.083)	(0.079)
Father: bachelor's or higher	1.151	1.213	1.175	1.152
	(0.122)	(0.127)	(0.129)	(0.122)
Public institution	1.358*	1.275	1.273	1.322*
	(0.186)	(0.167)	(0.173)	(0.186)
4-year institution	0.725	0.489	0.417	0.881
	(0.355)	(0.228)	(0.205)	(0.481)
Log fall enrollment 2011-12	1.018	1.044	1.035	1.003
	(0.055)	(0.059)	(0.056)	(0.060)
A or B GPA in high school	1.107	1.040	0.980	1.104
	(0.133)	(0.145)	(0.143)	(0.126)
SAT derived composite score	1.000	1.000	1.000	1.000
	(0.000)	(0.000)	(0.000)	(0.000)

Table A.15. Probability of switching fields (continued)

	(1)	(2)	(3)	(4)
2011-12 Pell grant recipient	0.849	0.774	0.781	0.879
	(0.108)	(0.112)	(0.117)	(0.115)
Received fed. benefit (2011-12)	1.120	1.124	1.068	1.130
	(0.144)	(0.144)	(0.143)	(0.153)
No financial aid app	1.575*	1.631*	1.650*	1.564*
	(0.301)	(0.373)	(0.376)	(0.305)
Log(AGI + 1)	0.961	0.907	0.914	0.976
	(0.054)	(0.059)	(0.063)	(0.057)
Zero AGI reported	0.488	0.322	0.348	0.580
	(0.297)	(0.233)	(0.270)	(0.356)
Student controls	Y	Y	Y	Y
Sample restrictions	None	BA only	Enroll in BA only	No for-profit
Percent correctly predicted	0.608	0.598	0.594	0.608
Goodness of fit F-stat	0.993	0.534	0.441	0.681
Goodness of fit p-value	0.447	0.849	0.911	0.725

Table A.16. Probability of switching fields. Logit models (odds ratio reported) where the dependent variable indicates whether respondents' switch majors between the 2011-12 and 2014 BPS:12 PETS survey waves. Column (1) reports the results for the base sample of students either (a) pursuing a bachelor's degree in 2014, or (b) enrolled in an an associate's degree program, with the expectation of continuing on to a bachelor's degree program within two years. Column (2) reports results using student's 2011-12 STEM GPA, where students without STEM GPAs have been droped from this sample. Column (3) reports results for students who took some STEM credits during the 2011-12 school year. Column (4) reports results when interacting GPA, initial STEM GPA, and gender. Probability threshold for percent correctly predicted based on a fraction of successes in the sample. F-adjusted mean residual test and associated p-values reported. BPS:12 PETS Panel study weight used. Standard errors in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

	(1)	(2)	(3)	(4)
Female	0.839	0.780*	0.865	1.220
	(0.105)	(0.092)	(0.100)	(0.518)
STEM field 2011-12	0.488***	0.495***	0.505***	0.737
	(0.084)	(0.085)	(0.093)	(0.480)
Female X STEM	2.380***	2.055***	2.073**	2.193
	(0.523)	(0.426)	(0.468)	(2.464)
2011-2012 GPA	0.609***		0.582***	0.670***
	(0.042)		(0.046)	(0.070)

Table A.16. Probability of switching fields (continued)

	(1)	(2)	(3)	(4)
Enrolled in bachelor's	0.795	0.693	0.767	0.786
	(0.386)	(0.314)	(0.422)	(0.378)
Underrepresented minority	1.154	1.047	1.068	1.151
	(0.196)	(0.140)	(0.147)	(0.191)
First generation student	0.880	0.976	0.911	0.883
	(0.145)	(0.177)	(0.162)	(0.143)
Mother: bachelor's or higher	0.808*	0.859	0.868	0.808*
	(0.083)	(0.095)	(0.103)	(0.083)
Father: bachelor's or higher	1.151	1.127	1.195	1.157
	(0.122)	(0.120)	(0.129)	(0.123)
Public institution	1.358*	1.215	1.157	1.352*
	(0.186)	(0.199)	(0.193)	(0.187)
4-year institution	0.725	0.783	0.684	0.721
	(0.355)	(0.365)	(0.390)	(0.350)
Log fall enrollment 2011-12	1.018	1.022	1.071	1.020
	(0.055)	(0.063)	(0.069)	(0.055)
A or B GPA in high school	1.107	1.127	1.120	1.108
	(0.133)	(0.139)	(0.144)	(0.134)
SAT derived composite score	1.000	1.000	1.000	1.000
	(0.000)	(0.000)	(0.000)	(0.000)
2011-12 Pell grant recipient	0.849	0.955	0.922	0.847
	(0.108)	(0.138)	(0.124)	(0.107)
Received fed. benefit (2011-12)	1.120	1.103	1.128	1.115
	(0.144)	(0.160)	(0.150)	(0.142)
No financial aid app	1.575*	1.360	1.435	1.566*
	(0.301)	(0.291)	(0.288)	(0.301)
Log(AGI + 1)	0.961	0.964	0.967	0.960
	(0.054)	(0.056)	(0.055)	(0.054)
Zero AGI reported	0.488	0.542	0.566	0.483
	(0.297)	(0.364)	(0.389)	(0.294)
2011-12 STEM GPA		0.675***		
		(0.036)		
Female X GPA				0.879
				(0.134)
STEM X GPA				0.865
				(0.190)
Female X STEM X GPA				1.038
				(0.386)
Sample restrictions	None	STEM GPA	STEM cred	None
Percent correctly predicted	0.608	0.602	0.615	0.605
Goodness of fit F-stat	0.993	0.508	0.685	1.005

Table A.16. Probability of switching fields (continued)

	(1)	(2)	(3)	(4)
Goodness of fit p-value	0.447	0.867	0.722	0.438

A.1.9 Notes on logit estimation

Consider the binary response model of the form:

$$\mathbb{P}(y = 1 \,|\mathbf{x}) = \Lambda(\mathbf{x}\beta) = \frac{\mathrm{e}^{\mathbf{x}\beta}}{1+\mathrm{e}^{\mathbf{x}\beta}}$$

where $\mathbf{x}_i$ is a $1 \times K$ vector with the first element equal to one, β is a $K \times 1$ vector, and $\Lambda(z)$ is the cumulative distribution function for the standard logistic distribution. Given data on $\mathbf{x}$, the model can be estimated using maximum likelihood estimation. For details, see Wooldridge (2010).

Two primary goodness-of-fit tests are utilized in this analysis after estimating probit models. The first is the percent correctly predicted. Because the proportion of successes in the data can be low for certain dependent variables, a prediction threshold of 0.5 may be inappropriate. Therefore, I use the fraction of successes in the sample to determine the prediction threshold.

The second goodness-of-fit test is the F-adjusted mean residual test, outlined in Archer and Lemeshow (2006). A common goodness-of-fit test for logistic regression models is given in Hosmer and Lemeshow (1980), which tests if the lack of fit is significant (potentially indicating problems with the model). Observations are grouped into quantile groups by expected probability, and a chi-squared test is to test whether the expected and observed frequencies across groups are the same. This test runs into difficulties when applied to survey data, because the sampling weights used effectively increase the observed and expected number of observations during maximum likelihood estimation. To address this issue, Archer and Lemeshow (2006) propose using the F-adjusted mean residual test statistic:

$$\hat{F} = \hat{X}^2(d-g+2)/(dg),$$

where d is the design degrees of freedoms, g denotes the number of quantile group, and $\hat{X}$ is the Wald test statistic for testing the g groups. This test statistic is approximately follows an F distribution with $g-1$ numerator degrees of freedom and $d-g+2$ denominator degrees of freedom. Given a small enough p-value $p = \mathbb{P}\left(F \geq \hat{F}\right)$, one can reject the null hypothesis that the observed and expected frequencies are generally in line. Therefore, rejecting the null may be considered evidence that the model is not a good fit for the data.

Chapter 2

Group-based beliefs and human capital specialization

2.1 Introduction

The reversal of the gender gap in postsecondary degree attainment over the past fifty years is related to a number of macroeconomic benefits, including a reduction in the gender wage gap (Blau and Kahn 2017) and increased aggregate economic productivity (Hsieh et al. 2019). However, the overall convergence masks the persistence of gender gaps within certain fields of study (Black et al. 2008). A large literature exists to explain the determinants of college major choice, a subset of which examines how and whether these motivations result in gender gaps in major decisions.[1] As reviewed in chapter 1 of this dissertation, empirical analysis has demonstrated that important factors include uncertainty and learning about abilities; the influence of preferences and societal norms; and the influence of social networks and role models. However, it is generally beyond the scope of empirical analysis to evaluate how these factors interact with one another. For example, how can the presence of role models influence learning about ability? Or how might norms dynamically change in the presence of uncertainty? A structural model is needed to synthesize these explanations and dynamically characterize the evolution of college major choice over time; in fields where women are historically under-represented, their beliefs about their ability may have higher variance than men. This can result in the persistence of

1. For a broad overview of the determinants of college major choice, see Patnaik, Wiswall, and Zafar (2020).

gender gaps in college major choices.

The goal of this chapter is to build and employ such a model. Specifically, I build on Alon and Fershtman (2019) to develop a model of group-based belief formation and gradual human capital specialization. I assume individuals belonging to a particular group choose to work or study in heterogeneous fields. Returns to education are stochastic, and underlying abilities are unknown. Agents form beliefs about their unknown abilities based on existing group outcomes, and update these beliefs as they study. I can use this model to analytically characterize the dynamics of students' belief distributions as their education proceeds.

This theoretical model is motivated by empirical results presented in chapter 1 of this dissertation, specifically with respect to major switching in STEM majors. STEM degree-holders have the highest average and median earnings (Patnaik, Wiswall, and Zafar 2020), yet women have been historically underrepresented; this allocation away from higher-earning majors has implications for the overall gender wage gap (Sloane, Hurst, and Black 2020). In chapter 1, I present evidence that, conditional on GPA, women initially enrolled in STEM fields are more likely to change major than men. The model developed below presents a potential mechanism for why this might be the case.

This chapter proceeds as follows: after briefly reviewing the literature, I outline the model and discuss the role of group-based beliefs in section 2.2. Section 2.3 presents notes on solving the model. Section 2.4 then presents simulations of agent behavior to highlight the important role that beliefs play in specialization decisions. I end this chapter in sections 2.5 and 2.6 by discussing potential extensions to this model. In section 2.5, I intuitively describe the connection between the theoretical model outlined here and theories of statistical discrimination. Section 2.6 establishes the minimum conditions needed to embed this model in a dynamic framework, allowing for the application to self-fulfilling prophecies and feedback loops. Section 2.7 concludes.

Literature

This paper builds upon the extensive literature on human capital formation (Becker 1962; Ben-Porath 1967; Mincer 1974; Rosen 1983). I expand on Alon and Fershtman's (2019) theoretical model of gradual specialization, a recent contribution to the literature. Their framework closely relates to two classical papers from this field: the seminal Mincer (1974) model of the returns to education, and the Roy (1951) model of occupational choice and skill heterogeneity. Alon and Fershtman's (2019) model, and by extension my own, can be viewed as a generalization of the Mincerian model of human capital accumulation to include a dynamic Roy model with unknown heterogeneous abilities and sequential learning.

Additionally, this project draws from the large literature on college major choice, a subset of which focuses on differences in college major and occupational choice by gender. The extensive empirical literature on the determinants of educational choices is reviewed in detail in Patnaik, Wiswall, and Zafar (2020).[2] This paper shares several theoretical commonalities with Arcidiacono et al. (2016), who build a dynamic model of school and work decisions, though with a focus on attrition. Sloane, Hurst, and Black (2020) provides a key empirical motivation for this paper; they use American Community Survey data to explore the importance of major choice in explaining labor market outcomes. Particularly relevant dimensions of this literature are reviewed in chapter 1 of this dissertation, however a few additional notes are warranted before proceeding.

The model discussed below details a mechanism through which gendered beliefs about performance may influence gender gaps in college major choice. This mechanism is supported by a number of recent experimental findings. For example, Porter and Serra (2020) finds that the presence of same-gender role models has a significant impact on the number of women choosing to major in economics. Recent experimental evidence from Owen (2020) isolates the role of belief in college specialization decisions. Her results suggest that beliefs about academic

2. Earlier reviews about college major choice and occupational choice can be found in Altonji, Blom, and Meghir (2012) and Altonji, Arcidiacono, and Maurel (2016).

performance partially drive the gender gap in college major decisions.[3]

A note on preferences, in particular with respect to the subjective expectations literature. The subjective expectations literature, which is comprehensively reviewed in Patnaik, Wiswall, and Zafar (2020) (and briefly discussed in chapter 1), has made considerable progress in "unpacking the black box of tastes," with respect to college major choice decisions. Preferences and norms are broadly defined in the college major choice literature; the catchall of "tastes" includes both the short-term enjoyment of coursework (Zafar 2013), as well as longer-term occupational preferences, such as those over temporal flexibility (Wiswall and Zafar 2018). This evidence suggests that preferences and norms play a central role in pre-labor market specialization decisions (Wiswall and Zafar 2015; Arcidiacono et al. 2020), in particular in explaining differences in major choice by gender (Wiswall and Zafar 2018). Within this realm of the literature, comprehensive surveys are carefully executed at single institutions to elicit details about students' preferences for studying different fields. While the model developed in this chapter could easily accommodate differences in preferences, the exposition here abstracts from student preferences entirely. This is meant to present an alternative mechanism to explain why we see these gender differences in college major choice that does not rely on gender differences in preferences at all. In fact, an over-reliance on residualized notions of preferences would mask the dynamics present in my model. My goal is not establish whether these gender differences college major choice are due to differences in preferences or differences in beliefs, but rather to present an alternative framework for explaining these gaps that does not rely on men and women simply preferring different choice sets.

Finally, the results of this paper can be closely tied to statistical discrimination literature. The theoretical connection between the model outlined in this paper and the theory of statistical discrimination will be discussed in detail in section 2.5, at which point relevant literature will be reviewed.

3. Her results isolate the mechanism whereby other types of informational interventions, such as those proffered in Porter and Serra (2020), Bayer, Bhanot, and Lozano (2019), and Li (2018), impact differences in college major choice by gender.

2.2 Model of human capital specialization

This section outlines my theory of group-based beliefs and human capital specialization. Please note that the agent's specialization problem and subsequent decision rule follows Alon and Fershtman's (2019) model of gradual human capital specialization; as such, I refer to the reader to their original paper for details.

2.2.1 Specialization decision

Assume infinitely lived agents in discrete time choose to work or study among J fields. Specifically, at each time period t, individuals decide to either study a single field j at a postsecondary institution, or to work in a field j, where $j \in \{1,\ldots,J\}$. The variable m_{jt} is equal to one if an agent matriculates and studies field j at time t; otherwise, m_{jt} equals zero. Likewise, ℓ_{jt} indicates whether an agent works in field j at time t. An individual therefore faces the following time constraint in each period t:

$$\sum_{j=1}^{J} \left(m_{jt} + \ell_{jt}\right) = 1.$$

Individuals aim to maximize their expected lifetime utility. Let U_j denote the expected per-period utility associated with working in field j, where U_j is a bounded function that is non-decreasing in field-specific wages and in field-specific human capital (w_j and h_{jt}, respectively).[4] The agent's expected lifetime payoff can be written as:

$$\sum_{t=0}^{\infty} \delta^t \sum_{j=1}^{J} U_j(w_j, h_{jt})\ell_{jt}, \tag{2.1}$$

where $\delta \in (0,1)$ is the discount rate.

Agents are initially endowed with some level of field-j human capital, h_{j0}, and can stochastically accumulate more field-j human capital by studying. Recall that if $m_{jt} = 1$, an agent matriculates in time period t to take a course in field j. Let s_{jt} indicate whether an agent

4. As noted in Alon and Fershtman (2019), per-period utility can be an object that evolves after graduation if we assume the summand is the expected average per-period payoff.

studying j at time t succeeds and passes that course; if the student succeeds, then $s_{jt} = 1$, otherwise $s_{jt} = 0$. I assume that whether a student passes or fails a particular field-j course is stochastic. Specifically, each student is endowed with some immutable probability of success in field-j courses, denoted θ_j. A student then passes any field-j course with probability θ_j, implying that s_{jt} is a Bernoulli random variable with parameter θ_j.[5] Agents only accumulates human capital when they pass courses. Therefore, an agent's field-specific human capital evolves according to:

$$h_{j,t+1} = h_{jt} + \nu_j s_{jt} m_{jt}, \qquad s_{jt} \sim \text{Bernoulli}(\theta_j), \tag{2.2}$$

where $\nu_j \geq 0$ is the human capital gain associated with passing the course.[6]

A student's probability of success in a field-j course, θ_j, is an ability parameter; students with high values of θ_j are more likely to pass any given class in field j, whereas students with low values of θ_j are more likely to fail. Students do not know their personal value of θ_j, but they have beliefs about what their value of θ_j might be. Their initial beliefs about their own ability are described by the distribution P_{j0}. As students take courses in field-j, they update their belief about what their value of θ_j may be, according to some updating rule Π_j:

$$P_{j,t+1} = \Pi_j(P_{jt}, s_{jt})$$

Overall, students have two incentives for studying a particular field: first, they can potentially accumulate j-specific human capital, which directly increases lifetime utility if they specialize in j. Second, studying j reveals information about their ability in that field, which is central to the specialization decision.

From equation (2.1) we know that agents will only want to work in the field that yields the highest expected lifetime utility. The choice of which field to work in is an individual's

5. The variable s_{jt} is assumed to be independent and stationary over time.

6. The value ν_j can be considered to be the Mincerian returns for a particular course. The per-period expected accumulation of human capital additionally must be non-negative and bounded. Regularity conditions imposed in section 2.2.3 ensure that is the case.

specialization decision. Agents that plan to specialize in field j will study j to accumulate j-specific human capital, and will eventually endogenously enter the labor market as a field-j specialist. The decision of when to stop studying j and enter the labor market is an agent's *stopping problem*. Specifically, the stopping decision is the time when an agent expects to stop studying j and begin work as a j-specialist, ignoring the existence of other fields.[7]

2.2.2 Group-based beliefs

Assume each student has a group type, g. To simplify the exposition, consider two groups, men and women ($g \in \{m, f\}$). The distributions of underlying abilities, θ_j, are the same for men and women. However, initial beliefs about underlying abilities, P^g_{j0}, are different for the two groups.

To make this explicit, consider the following parameterization of a student's belief distribution. Assume the initial beliefs about θ_j follow a beta distribution with parameters $(\alpha^g_{j0}, \beta^g_{j0})$, implying $P^g_{j0} = \mathscr{B}(\alpha^g_{j0}, \beta^g_{j0})$. To understand why this is a reasonable assumption, recall that our unknown ability parameter, θ_j, is the probability that a student succeeds ($s^g_{jt} = 1$) or fails ($s^g_{jt} = 0$) a field-j course at time t. Because the realizations of s_{jt} are independent over time, the sequence of successes and failures for some number of total field-j courses taken is a binomial random variable. The beta distribution is a conjugate prior for the binomial distribution, and is thus a natural and tractable choice for modeling beliefs about θ_j (Casella and Berger 2002, pg. 325). Therefore, if we assume that a student updates their beliefs about θ_j using Bayes' rule, their posterior is also a Beta distribution:

$$P^g_{j,t+1} = \mathscr{B}(\alpha^g_{j,t+1}, \beta^g_{j,t+1}), \qquad (\alpha_{j,t+1}, \beta_{j,t+1}) = \begin{cases} (\alpha^g_{jt} + 1, \beta^g_{jt}) & \text{if } s^g_{jt} = 1 \\ (\alpha^g_{jt}, \beta^g_{jt} + 1) & \text{if } s^g_{jt} = 0 \end{cases} \tag{2.3}$$

7. The point that this decision is made ignoring the existence of other fields is important. This is because I assume that agents do not earn a wage while they are studying. Therefore, an agent's expected field-j stopping time directly impacts their expected field-j payoff.

To develop intuition about how this assumption influences specialization decisions, it's helpful to proceed with an illustrative, albeit somewhat contrived, parameterization. Let α^g_{j0} and β^g_{j0} denote the number of type g students who have succeeded and failed in field j at time 0, respectively. As an example, suppose a type g student is forming their initial beliefs about their probability of success in field j, and therefore asks five type g upperclassmen about their experiences in field j. If three of those type g upperclassmen passed the introductory course in field j, while two failed, then the student's initial belief parameters $(\alpha^g_{j0}, \beta^g_{j0})$ would equal $(3,2)$.

Using this parameterization, the observed group-g success rate, μ^g_{j0}, is given by:

$$\mu^g_{j0} = \frac{\alpha^g_{j0}}{\alpha^g_{j0} + \beta^g_{j0}}.$$

This average is based on a sample of $n^g_{j0} = \alpha^g_{j0} + \beta^g_{j0}$ type-g students. Note that the beta distribution parameters, α^g_{j0} and β^g_{j0}, can be expressed using the average success rate, μ^g_{j0}, and the sample size, n^g_{j0}:

$$\alpha^g_{j0} = \mu^g_{j0} n^g_{j0}, \qquad \beta^g_{j0} = (1 - \mu^g_{j0}) n^g_{j0}.$$

Therefore, an alternative parameterization of the prior is given by:

$$\mathscr{B}\left(\alpha^g_{j0}, \beta^g_{j0}\right) = \mathscr{B}\left(\mu^g_{j0} n^g_{j0}, (1 - \mu^g_{j0}) n^g_{j0}\right).$$

Assume the sample size of men is larger than that of women, but that the observed success rate is the same for the two genders:

$$n^m_{j0} > n^f_{j0}, \qquad \mu_{j0} = \mu^m_{j0} = \mu^w_{j0}.$$

Figure 2.1 provides a numerical example to illustrate how these assumptions affect the priors of men and women. Although women and men have the same probability of success in expectation, women have more initial uncertainty regarding their underlying abilities.

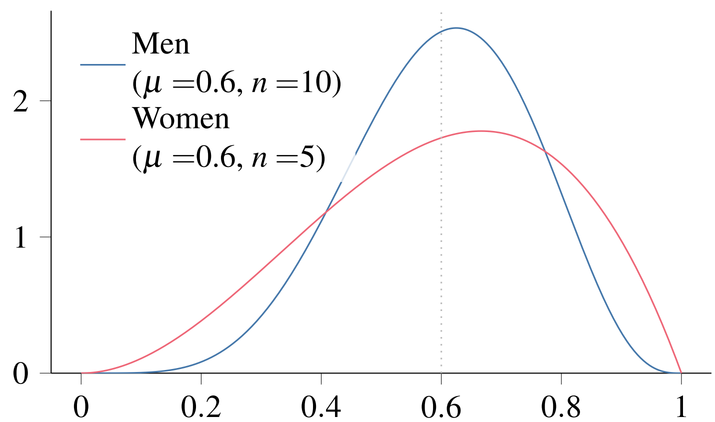

Figure 2.1. PDF of Beta distribution

In section 2.4, I discuss how differential initial beliefs influence specialization decisions. Before moving forward, it is helpful to highlight some shortcomings of the above illustrative example. First, I am being purposefully vague about what "success" means when agents form their initial priors, α_{j0}^{g} and β_{j0}^{g}. The example above, in which students solicit feedback from upperclassmen, is helpful for building intuition, and is consistent with the literature highlighting the importance of role-models in specialization decisions; see Porter and Serra (2020) for directly relevant empirical evidence. However, success could mean many things in this model. It could be the number of type g students who graduate into field j, the number of type g professors, the number of students attaining graduate degrees in field j, etc. More generally, while it is illustrative to use the parameters α_{j0}^{g} and β_{j0}^{g} to tally the total number of type g successes and failures, it is by no means necessary (or reasonable) to calibrate the model in this way.

2.2.3 Optimal policy

To summarize the individual's problem, let h_t^g, P_t^g, m_t^g, ℓ_t^g denote the $J \times 1$ vectors of field-specific human capital, beliefs, study decisions, and labor decisions, respectively. A policy $\pi : (h_t^g, P_t^g) \to (m_t^g, \ell_t^g)$ is optimal if it maximizes lifetime expected utility:

$$\mathbb{E}^{\pi}\left[\sum_{t=0}^{\infty} \delta^t \left(\sum_{j=1}^{J} U_j(h_{jt}^g, w_j)\ell_{jt}^g\right) \middle| h_0^g, P_0^g\right], \tag{2.4}$$

given the following time constraint:

$$\sum_{j=1}^{J}(m^g_{jt}+\ell^g_{jt})=1, \qquad m^g_{jt},\ell^g_{jt}\in\{0,1\},$$

subject to the human capital accumulation and belief transition laws:

$$h^g_{jt+1} = h^g_{jt}+\nu_j s^g_{jt} m^g_{jt}, \qquad s^g_{jt}\sim \text{Bernoulli}(\theta_j), \qquad \theta_j\sim P^g_{j0}\equiv \mathscr{B}(\alpha^g_{j0},\beta^g_{j0}),$$

$$P^g_{j,t+1} = \mathscr{B}(\alpha^g_{j,t+1},\beta^g_{j,t+1}), \qquad (\alpha^g_{j,t+1},\beta^g_{j,t+1}) = \begin{cases}(\alpha^g_{jt}+1,\beta^g_{jt}) & \text{if } m^g_{jt}=1 \text{ and } s^g_{jt}=1\\ (\alpha^g_{jt},\beta^g_{jt}+1) & \text{if } m^g_{jt}=1 \text{ and } s^g_{jt}=0\\ (\alpha^g_{jt},\beta^g_{jt}) & \text{if } m^g_{jt}=0\end{cases}\cdot$$

Alon and Fershtman (2019) characterize the optimal policy to the above problem. To apply their results, first assume the following initial condition:

$$h_{j0}\le \alpha^g_{j0}\nu_j. \tag{2.5}$$

This assumption will be discussed in detail in section 2.3.2. In brief, this condition assures that the stopping problem is monotonic, which in turn implies the optimality of the following policy. Let τ denote the optimal stopping rule defined over $\{s^g_{j1},s^g_{j2},\dots\}$. Define the field-$j$ index as the expected lifetime payoff an agent would receive if they commit to studying field j given their state (h^g_{jt},P^g_{jt}):

$$\mathscr{I}_{jt}(h^g_j,P^g_j)=\sup_{\tau\ge 0}\mathbb{E}^{\tau}\left[\sum_{t=0}^{\infty}\delta^t U_j(h^g_{jt},w_j)\ell^g_{jt}\,\middle|\,(h^g_{j0},P^g_{j0})=(h^g_j,P^g_j)\right] \tag{2.6}$$

Define the graduation region of field j as the states where an agent committed to studying field j

would choose to stop studying and enter the labor market:

$$\mathscr{G}_j(h_j^g, P_j^g) = \left\{ (h_j^g, P_j^g) \middle| \arg\max_{\tau \geq 0} \mathbb{E}^\tau \left[\sum_{t=0}^{\infty} \delta^t U_j(h_{jt}^g, w_j) \ell_{jt}^g \middle| (h_j, P_j^g) \right] = 0 \right\} \tag{2.7}$$

Then the following policy $\pi : (h_t^g, P_t^g) \to (m_t^g, \ell_t^g)$ is optimal:

1. At each $t \geq 0$, choose skill $j^* = \arg\max_{j \in J} \mathscr{I}_j$, breaking ties according to any rule
2. If $(h_{j^*}, P_{j^*}^g) \in \mathscr{G}_j$, then enter the labor market as a j^* specialist. Otherwise, study j^* for an additional period.

2.3 Notes on solving the model

The optimal policy outlined in 2.2.3 is characterized by two objects: the index (2.6), which summarizes the agent's expected lifetime payoff if they commit to one field and ignore all others; and the graduation region (2.7), which defines the states where an agent would choose to stop studying a particular field and enter the labor market. This section presents several tools for computing the model. Section 2.3.1 begins by introducing alternative notation to characterize state variables. In section 2.3.2, I discuss the intuition behind the initial condition assumption (2.5), and its implications for the optimal stopping problem. This motivates a tractable solution to the graduation region. Section 2.3.3 uses the results from 2.3.2 to derive a simplified version of the index.

2.3.1 Comment on state variables

State variables in this model are given by an agent's vector of human capital, h_{jt}, and their beliefs, P_{jt}. This section presents an alternative characterization of the agent's state variables that simplifies the analytical solution.

Define $\tilde{m}_{jt}$ as the total number of times a student has chosen to matriculate in field j by time t, and define $\tilde{s}_{jt}$ as the total number of times a student has passed their field j courses:

$$\tilde{m}_{jt} = \sum_{n=0}^{t-1} m_{jn}, \qquad \tilde{s}_{jt} = \sum_{n=0}^{t-1} s_{jn}. \tag{2.8}$$

The individual's state variables at time t are $(h_{jt}, \alpha_{jt}, \beta_{jt})$. Using some simple algebraic transformations,[8] we can now characterize the states at time t using $(\alpha_{j0}, \beta_{j0}, h_{j0}, \tilde{m}_{jt}, \tilde{s}_{jt})$. In words, the agent's state variables at time t are the initial belief parameters α_{j0} and β_{j0}, initial human capital h_{j0}, the endogenous number of field-j courses $\tilde{m}_{jt}$, and the stochastic number of times an agent passed their field-j courses, $\tilde{s}_{jt}$. Given the structure of the problem, there is no need to directly track the evolution of $(\alpha_{jt}, \beta_{jt}, h_{jt})$ over time, because (1) all information about the evolution of beliefs is captured by initial beliefs $(\alpha_{j0}, \beta_{j0})$, course choices ($\tilde{m}_{jt}$), and course outcomes ($\tilde{s}_{jt}$); and (2) all information about human capital evolution is characterized by initial human capital endowments (h_{j0}), course choices ($\tilde{m}_{jt}$), and course outcomes ($\tilde{s}_{jt}$).

2.3.2 Initial condition assumption and optimal stopping time

The optimal policy from section 2.2.3 is contingent on assuming equation (2.5), which states that $h_{j0} \leq \nu_j \alpha_{j0}$. Assuming that $h_{j0} \leq \nu_j \alpha_{j0}$ ensures that the stopping problem is monotonic, which in turn implies the optimality the policy outlined in section 2.2.3. As such, I will occasionally refer to (2.5) as the *monotonic initial condition* or the *monotonicity assumption*.

It may be helpful to briefly outline why the monotonic initial condition implies optimality of the above policy. Recall that the graduation index, (2.7), characterizes the states where an individual would stop studying field-j and enter the labor market as a field-j specialist, ignoring all other fields. Therefore, this object characterizes the field-specific stopping problem facing an individual. Under the monotonic initial condition (2.5), the stopping problem for a given field is monotone. Intuitively, monotonicity of the stopping problem means that an agent who wants to stop studying j at time t would also want to stop studying j at time $t+1$ if they continued on,

8. Specifically, note that (1) $\tilde{m}_{jt} + \alpha_{j0} + \beta_{j0} = \alpha_{jt} + \beta_{jt}$; (2) $\alpha_{jt} = \tilde{s}_{jt} + \alpha_{j0}$; and (3) $h_{jt} = \nu_j \tilde{s}_{jt} + h_{j0}$.

independent of stochastic outcomes. Therefore, an agent's stopping decision can be reduced to a comparison between their current expected lifetime payoff in j and their expected payoff the following period. In other words, monotonicity implies the optimality of a one-step-look-ahead comparison for the field-specific stopping problem.

Therefore, the monotonicity condition (2.5) ensures that an agent evaluating field j at time t will stop studying j if their expected lifetime payoff in the current period exceeds their expected lifetime payoff in $t+1$:

$$\frac{1}{1-\delta} w_j h_{jt} \geq \frac{\delta}{1-\delta} w_j \mathbb{E}_t \left[h_{j,t+1} \middle| h_{jt}, \alpha_{jt}, \beta_{jt} \right]$$

This equation can be simplified using the human capital accumulation function (2.2):

$$h_{jt} \geq \delta(h_{jt} + \nu_j \mathbb{E}_t \left[s_{jt} \middle| h_{jt}, \alpha_{jt}, \beta_{jt} \right]).$$

Recalling that the course outcome s_{jt} is a Bernoulli(θ_j) random variable, this can be written using an agent's beliefs about θ_j at time t:

$$\frac{1-\delta}{\delta} \geq \frac{\nu_j \alpha_{jt}}{h_{jt}(\alpha_{jt} + \beta_{jt})} \tag{2.9}$$

Using the definitions of $\tilde{m}_{jt}$ and $\tilde{s}_{jt}$ from equation (2.8), the stopping condition (2.9) becomes:[9]

$$\tilde{m}_{jt} \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{jt}}{h_{j0} + \nu_j \tilde{s}_{jt}} \right) - \alpha_{j0} - \beta_{j0} \tag{2.10}$$

9. To replicate this derivation, note $\nu_j \alpha_{jt} = h_{jt} - h_{j0} + \nu_j \alpha_{j0}$. Then (2.9) implies:

$$h_{jt} \left(\alpha_{jt} + \beta_{jt} - \frac{\delta}{1-\delta} \right) \geq \frac{\delta}{1-\delta} \left(\nu_j \alpha_{j0} - h_{j0} \right)$$

Using the fact that $\tilde{m}_{jt} + \alpha_{j0} + \beta_{j0} = \alpha_{jt} + \beta_{jt}$:

$$h_{jt} \tilde{m}_{jt} + h_{jt} \left(\alpha_{j0} + \beta_{j0} \right) \geq \frac{\delta}{1-\delta} \left(\nu_j \alpha_{j0} - h_{j0} + h_{jt} \right)$$

The simplified stopping condition under monotonicity (2.10) follows from the fact that $h_{jt} - h_{j0} = \nu_j \tilde{s}_{jt}$

Intuitively, this stopping condition says that an agent will stop studying a field j once their total number of completed field-j courses exceeds the right-hand-side inequality. The graduation index (2.7) can now be written to reflect the stopping condition (2.10):

$$\mathscr{G}_j = \left\{ \tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0} \middle| \tilde{m}_{jt} \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{jt}}{h_{j0} + \nu_j \tilde{s}_{jt}} \right) - \alpha_{j0} - \beta_{j0} \right\} \tag{2.11}$$

Before proceeding, it is helpful to discuss some properties of this object. An agent in this model begins with the initial states $(h_{j0}, \alpha_{j0}, \beta_{j0})$. They have taken zero courses in j, implying $\tilde{m}_{j0} = 0$, and thus have passed zero courses in j, meaning $\tilde{s}_{jt} = 0$. Under the initial monotonicity condition (2.5), the fraction $\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{jt}}{h_{j0} + \nu_j \tilde{s}_{jt}}$ is always greater than or equal to 1, and approaches 1 as $\tilde{s}_{jt}$ increases. This fact is important for bounding the number of periods an agent spends in school.

2.3.3 Simplified index

The goal of this section is to derive a simplified version of the index (2.6). Recall that the index $\mathscr{I}_j$ from equation (2.6) characterizes the expected lifetime payoffs associated with specializing in j, ignoring other fields. If the the simplified stopping condition under monotonicity holds at time t (i.e. $\left(\tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}\right) \in \mathscr{G}_j$), then the agent would expect to enter the labor market (ignoring other fields). Their expected lifetime payoff in j equals their expected lifetime earnings given their current levels of human capital:

$$\frac{1}{1-\delta} w_j h_{jt} = \frac{1}{1-\delta} w_j \left(h_{j0} + \nu_j \tilde{s}_{jt} \right)$$

If $\left(\tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}\right) \notin \mathscr{G}_j$, then the agent would plan on continuing their studies in j. Their expected lifetime payoff depends on how much human capital they expect to accumulate in j. To make this concrete, let m_j^* denote the total number of periods an agent expects to study field j before entering the labor market. Then the agent expects to be in school for $m_j^* - \tilde{m}_{jt}$ more periods. Because the agent is not earning an income while they are in school, their expected lifetime payoff will be discounted by $\delta^{m_j^* - \tilde{m}_{jt}}$. They then expect to enter the labor market at time $t + m_j^* - \tilde{m}_{jt}$ with some level of human capital, given by $h_{j,t+m_j^*-\tilde{m}_{jt}}$. The index when $(\tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}) \notin \mathscr{G}_j$ is then given by:

$$\frac{1}{1-\delta} w_j \mathbb{E}_t \left[\delta^{m_j^* - \tilde{m}_{jt}} h_{j,t+(m_j^* - \tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0} \right]$$

Therefore, the index (2.6) is characterized by:

$$\mathscr{I}_j = \begin{cases} \frac{w_j h_{jt}}{1-\delta} & \text{if } (\tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}) \in \mathscr{G}_j, \\ \frac{w_j}{1-\delta} \mathbb{E}_t \left[\delta^{m_j^* - \tilde{m}_{jt}} h_{j,t+(m_j^* - \tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0} \right] & \text{otherwise.} \end{cases} \tag{2.12}$$

Thus, key to solving the agent's problem is evaluating the index.

2.3.4 Solving the index

An analytical solution to the index is possible in certain contexts. As noted in Alon and Fershtman (2019), a key advantage of the model is its computability under the stronger initial condition assumption $h_{j0} = \alpha_{j0} \nu_j$. This assumption is not out of line with the human capital accumulation function (2.2), and may be reasonable for simulation exercises when all parameters of the problem are known. Indeed, this is the approach I take in section 2.4. However, this assumption can present both theoretical and empirical objections. The goal of this paper is to assess how beliefs impact specialization decisions. In the context of gender, this may involve considering whether a man and woman with similar initial human capital levels but different beliefs make different specialization choices; in the example outlined in section 2.2.2, this involves assessing whether men and women with similar h_{j0} make different specialization choices when $\alpha_{j0}^m > \alpha_{j0}^w$. However, assuming $h_{j0} = \alpha_{j0}^g \nu_j$ implies that women begin with lower levels of human capital than men in a particular field j, complicating this type of counterfactual analysis. More generally, assuming $h_{j0} = \alpha^g_{j0} v_j$ effectively eliminates the variable h_{j0}; it's no longer possible to control for initial human capital levels when agents make initial specialization choices.

However, it may be possible to maintain the the weaker monotonic weaker monotonic initial condition assumption, (2.5), and still produce an analytical solution to the model. To that end, section, A.2.1 presents a potential analytical solution to the model under the weaker condition, $h_{j0} \leq v_j \alpha_{j0}$; the relevant proofs and derivations are presented in section A.2.2.

2.4 Implications of the model

An agent's specialization decision is impacted by individual, group, and field characteristics. This section illustrates how these factors motivate an individual's behavior using a simplified version of the model. Section 2.4.1 develops a version of the model where an agent chooses between two completely symmetric fields. Simulations are explored in section 2.4.2 to illustrate how different factors influence decision making. In particular, I emphasize the role that beliefs play in an agent's specialization decision.

2.4.1 Choice between symmetric fields

Assume a student can choose to work or study in one of two fields, field X or field Y. Utility in field $j \in \{X,Y\}$ at time t is equal to income:

$$U_j(w_j, h^g_{jt})\ell^g_{jt} = w_j h^g_{jt} \ell^g_{jt} \tag{2.13}$$

Wages in fields X and Y are equal and are normalized to 1 ($w_X = w_Y = 1$), as are returns to successfully studying human capital ($v_X = v_Y = 1$). The student's underlying abilities in the two fields, θ_X and θ_Y, are both equal to 0.5. Therefore, the student has a 50% chance of passing any given field X or field Y course. Finally, I assume the student's beliefs about their own abilities in fields X and Y are equal to the uniform prior:[10]

$$P_{X,0} = \mathscr{B}(\alpha_{X,0}, \beta_{X,0}) = \mathscr{B}(1,1), \qquad P_{Y,0} = \mathscr{B}(\alpha_{Y,0}, \beta_{Y,0}) = \mathscr{B}(1,1),$$

For tractability, I modify the assumption (2.5) as follows:

$$h^g_{j0} = \nu_j \alpha^g_{j0}. \tag{2.14}$$

Equation (2.14) is consistent with the human capital accumulation function in equation (2.2). Further, this assumption ensures that the number of periods an agent studies in school is a deterministic function of initial beliefs, as discussed in section A.2.2.[11] The role of beliefs can be more clearly seen in these simulations because all agents who specialize in field j with the same initial beliefs will take the same number of courses in j.

The point at which an agent has "specialized" in a particular field j is not clearly defined in the model. I intuitively describe what specialization looks like in the simulations below, but it is helpful to provide some concrete definition of specialization. In the figures below, an agent has "specialized" in a field j if they would choose to continue to specialize in j if they failed all of their remaining courses in that field. Mathematically, this is represented by the following condition, letting m^*_j denote the number of courses an agent with beliefs $(\alpha_{j0}, \beta_{j0})$ would take in j before specializing:

$$\frac{1}{1-\delta} \delta^{m^*_j - \tilde{m}_{jt}} w_j h_{jt} > \mathscr{I}_k(\tilde{m}_{kt}, \tilde{s}_{kt}, \alpha_{k0}, \beta_{k0}, h_{k0}), \quad \forall k \neq j.$$

The left-hand side of this inequality is the agent's lifetime payoff of specializing in field j if they

10. Note that if $(\alpha, \beta) = (1,1)$, the beta distribution $\mathscr{B}(\alpha, \beta)$ equals the uniform distribution over $[0,1]$. This distribution can be seen graphically in figure 2.3a. Intuitively, this implies that the student thinks all values of θ_j are equally likely.

11. This is specifically addressed while finding lower and upper bounds for stopping times in section A.2.2. Assuming $h_{j0} = \nu_j \alpha_{j0}$ implies that all agents specializing in j with initial beliefs $(\alpha_{j0}, \beta_{j0})$ will take exactly $\left\lceil \frac{\delta}{1-\delta} \right\rceil - \alpha_{j0} - \beta_{j0}$ courses in j before entering the labor force.

fail all of their remaining courses in that field. This would imply they do not accumulate any more human capital in field j, so their lifetime payoff is based on their current levels of human capital, h_{jt}, discounted according to the number of periods they expect to study. The right-hand side of this inequality is the expected lifetime payoff associated will all other fields.

2.4.2 Simulations

Each subplot in figure 2.2 plots the fraction of simulated agents choosing to study field X or field Y at each time period t. Recall that agents studying field j at time t will either pass and successfully accumulate human capital ($s^g_{jt} = 1$) or they will fail ($s^g_{jt} = 0$), where $s^g_{jt} \sim \text{Bernoulli}(\theta_j)$. The student then updates their beliefs about their own underlying ability, θ_j. Line movements in figure 2.2 are caused by agents switching fields in response to updated beliefs. Eventually, students will specialize in one field and enter the labor market as a field-X or field-Y specialist. The line for any field j ends once any agent specializing in j stops studying and enters the labor market. Therefore, the length of the lines in figure 2.2 denote the minimum amount of time an agents spends studying before becoming a field-j specialist. Specialization in figure 2.2 is generally represented by a flattening of the curve; once a student has made their specialization decision, they no longer switch fields. To make this explicit, I use the definition of specialization from section 2.4.1, and define specialization as the point when an agent could fail all of their remaining field j courses, and would still choose to specialize in that field. The median point by which simulated agents have made their specialization decision is represented by the dashed vertical line in each plot.

The baseline scenarios in figures 2.2a and 2.2b illustrate these dynamics. Figure 2.2a plots the baseline scenario for 10,000 simulations; figure 2.2b plots the first 50 of these simulations. The first takeaway is that the agent's specialization decision in the baseline is effectively a coin flip. This is most clearly seen in figure 2.2a; at all time periods, approximately 50% of the agents are studying field X and and 50% are studying field Y. This should be expected, as fields X and Y are completely symmetric.

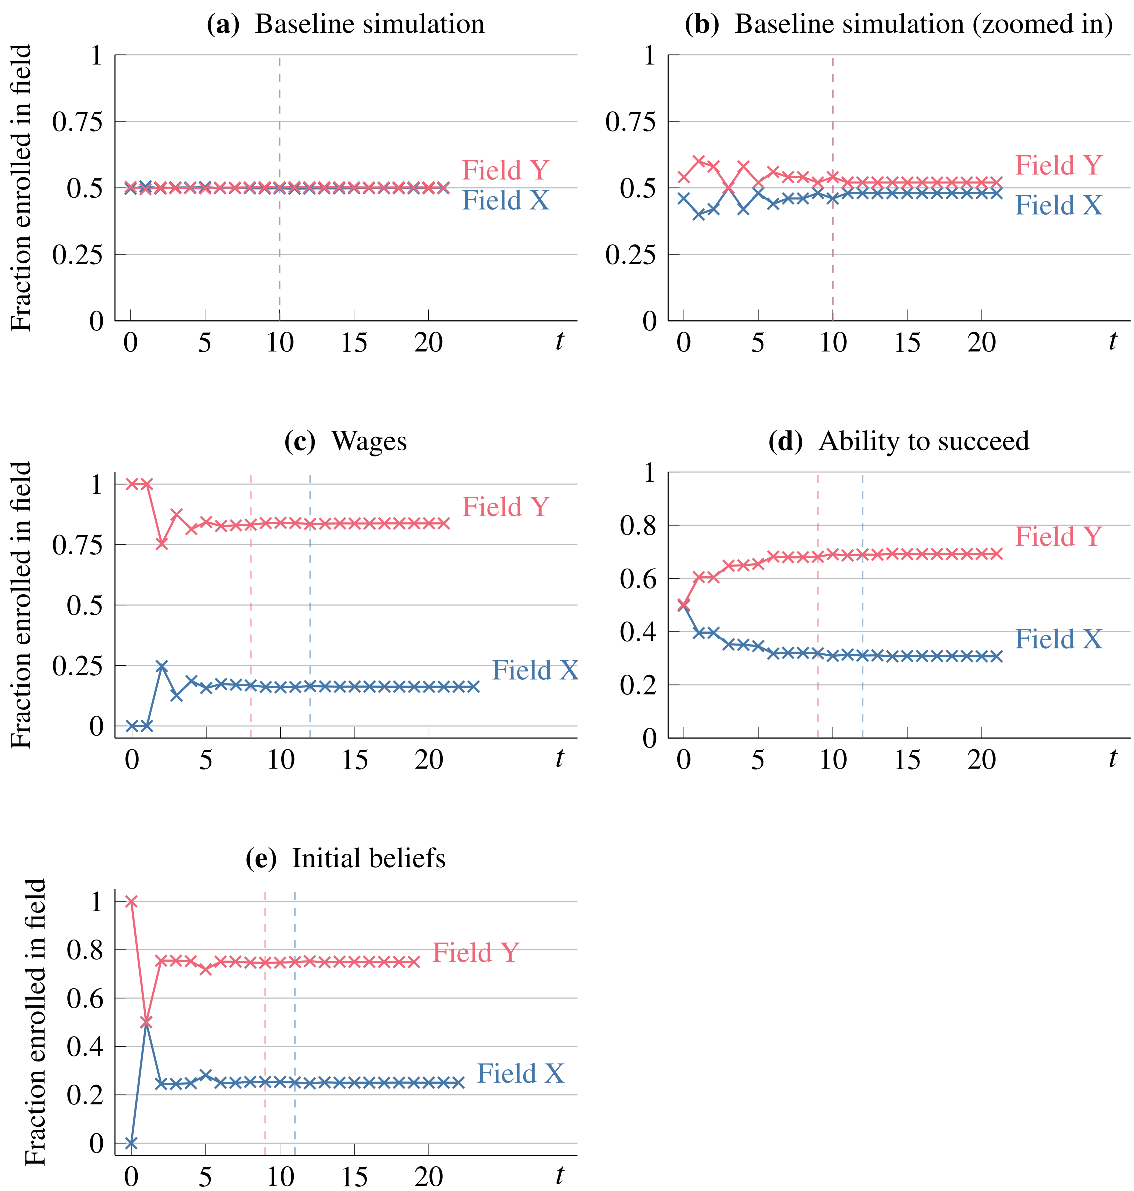

Figure 2.2. Simulations of simple version of model. Figure (a) presents the baseline for $N = 10,000$ simulations; figure (b) does the same for the first 50 simulations. The remaining figures have $N = 10,000$ simulations. Figure (c) repeats the simulations for $w_X = 1$ and $w_Y = 1.5$. Figure (d) repeats the simulations when $\theta_X = 0.4$ and $\theta_Y = 0.6$. Figure (e) repeats the simulations when $(\alpha_{X0}, \beta_{X0}) = (1,1)$ and $(\alpha_{Y0}, \beta_{Y0}) = (2,2)$.

Some of the more subtle decision dynamics can only be seen with fewer observations. Therefore, 2.2b zooms in on the first 50 of these simulations. Note that the fraction of students studying field X or field Y moves in early periods, but flattens out in later periods. This is

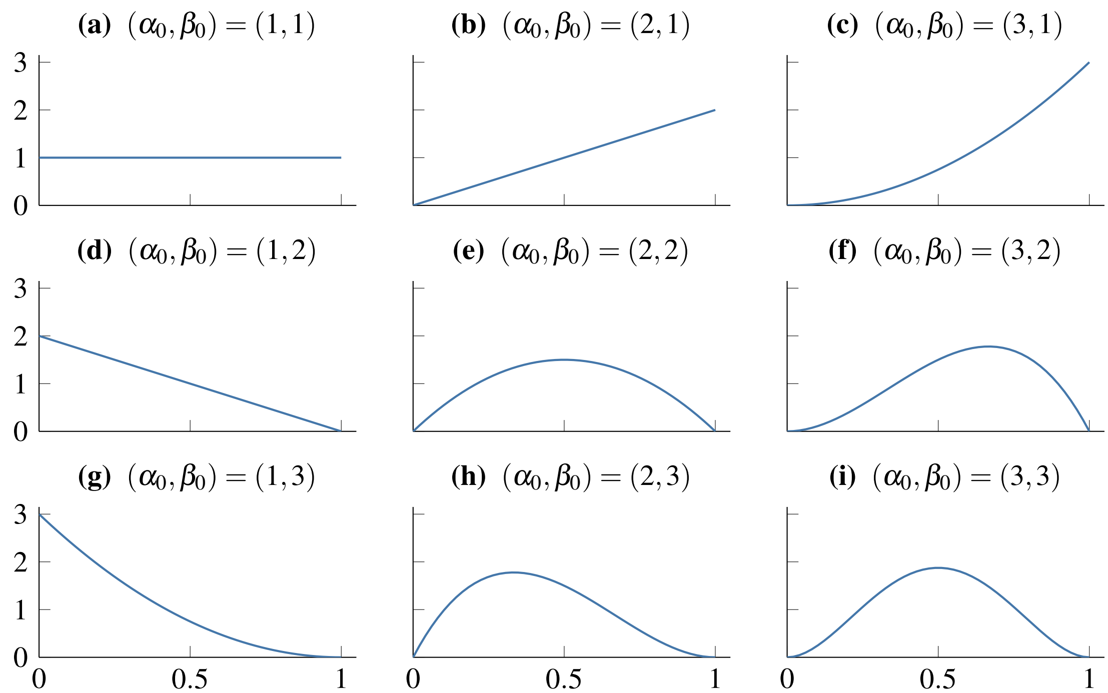

Figure 2.3. Evolution of the Beta distribution $\mathscr{B}(\alpha_0, \beta_0)$ for different values of (α_0, β_0).

because students at the beginning of their education will update their beliefs in response to course outcomes. These updated beliefs may cause students to switch fields, shifting the composition of simulated agents studying X or Y. In later periods, simulated agents have made their specialization decision and no longer switch fields. This specialization is represented by the flattening of the lines in figure 2.2b. As in 2.2a, approximately 50% of agents specialize in field X, and 50% specialize in field Y.

The remainder of figure 2.2 plots variations of the baseline for $N = 10,000$ simulations. In figure 2.2c, wages in field Y are 50% higher than wages in field X. All other variables are identical to the baseline scenario. Unsurprisingly, higher wages in Y drive specialization into that field. Because the expected lifetime payoff is so much higher, approximately 80% of agents choose to specialize in Y.

The field X line in figure 2.2c is longer than the field Y line, implying that agents who specialize in X spend more time in school. To understand why this happens, first note that all agents begin their education studying Y because of the higher relative wages. However, after two periods, a large fraction of agents switch from studying Y to studying X. This is due to agents (randomly) failing their first two courses in field Y, and switching into field X. The reason agents switch fields can be seen by the evolution of their belief distributions, shown in figure 2.3. The student's initial belief distribution is plotted in figure 2.3a. A student that fails their first course in field Y updates their beliefs about their underlying ability in Y to the distribution plotted in figure 2.3d; if they fail their second course in Y, they update their beliefs to 2.3g. As we can see from figure 2.3g, a student that fails their first two classes in Y will believe they likely have a lower ability in that field. As such, if they choose to specialize in Y, they would not expect to successfully accumulate much human capital over the course of their studies, implying a lower expected lifetime payoff. As a result, these agents switch to studying field X, in spite of the lower wages. This switching leads to more overall time in school; as mentioned above, equation (2.14) implies that the number of periods an agents spends studying field X or Y before becoming a specialist is a deterministic function of initial beliefs. Agents' initial beliefs about their abilities in X and Y are the same, and as such, agents specializing in either X or Y will study their chosen discipline for the same number of periods. However, because all agents spend their first two periods studying field Y, those who specialize in field X will study for a minimum of two more periods.

Figure 2.2d augments the baseline scenario so agents have a higher ability in field Y. Specifically, probability of success in any given field X course, θ_X, equals 0.4, whereas the probability of success in field Y is given by $\theta_Y = 0.6$. Unsurprisingly, a higher ability in field Y drives specialization into that field.

We now turn to the impact of differential priors on specialization dynamics, plotted in figure 2.2e. I assume simulated agents are initially more certain about their abilities in field Y relative to field X. Specifically, I assume a student's initial prior about their ability in Y is given by $P_{Y0} = \mathscr{B}(2,2)$; this corresponds to the distribution plotted in figure 2.3e. Their initial prior about their ability in X continues to equal the uniform distribution, $P_{X0} = \mathscr{B}(1,1)$. Note that agents have the same belief about their probability of success in X and Y in expectation. However, the variances of the initial distributions suggest that agents have more certainty about their underlying ability in field Y than in field X.

The first consequence of this assumption is that agents specializing in field X study for more periods than those specializing in field Y, as shown in figure 2.2e. As mentioned above, equation (2.14) implies that the number of periods an agent spends studying j before specializing in that field is a deterministic function of the agent's initial beliefs. Agents have more initial uncertainty about their abilities in X than in Y, and therefore they will study X for more periods before specializing in that field. The second takeaway from 2.2e is that all agents begin their education studying field Y. Agents know that if they become field Y specialists, they will finish their education earlier, and begin earning an income sooner. The prospect of ending their education earlier drives agents to initially study field Y.

The key takeaway from figure 2.2e is that increased initial certainty about field Y abilities causes more agents to specialize in field Y. [12] Although agents are equally likely to succeed in fields X and Y, and although the payoffs for specializing in these fields are the same, differential initial beliefs about underlying abilities drives the majority of simulated agents to specialize in field Y. Thus, initial beliefs play a key role in specialization decisions.

2.5 Connection to statistical discrimination

The model outlined in sections 2.2 and 2.3 is theoretically connected to the concept of statistical discrimination. To see this, it helps to briefly review statistical discrimination. Fundamentally, statistical discrimination is a theory whereby inequality results from rational agents forming expectations based on existing group characteristics. To better understand this definition, it helps to briefly review models of taste-based discrimination that preceded it. The canonical model of taste-based discrimination in labor markets, as formulated by Becker

12. It's worth emphasizing that this is not driven by risk aversion across fields; assuming linear utility in (2.13) ensures that agents are risk neutral across fields. Rather, concavity due to discounting ensures that agents are risk averse across time.

(1957), assumes prejudiced employers receive disutility from hiring employees belonging to a particular group.[13] A key implication of the Becker model is that discriminatory firms will be less profitable than firms that do not discriminate. Thus, long-run neoclassical analysis suggests that discrimination will be driven out of the marketplace, an implication that appears incongruous with the persistence of unexplained wage differentials between groups of workers.[14]

The theory of statistical discrimination, as first formulated by Arrow (1972) and Phelps (1972), grew out of this critique. The classical analysis formulated in Aigner and Cain (1977) assumes that a job applicant's group type is one variable an employer uses for inference about their unknown true productivity.[15] If groups have different aggregate characteristics, and these characteristics are controlled for by the potential employer, then individuals with the same ability from different group may have different expected productivities. In the labor market context, unequal outcomes arise from potential employers-as-statisticians using group-based information about prospective employees for inference. Statistical discrimination therefore presents a different view of inequality than the Becker model; unequal outcomes may not be the result of prejudice or distaste for certain group types, but rather the result of rational decision making.

Fundamentally, the model outlined in sections 2.2 and 2.3 is a model of statistical discrimination. To convince the reader that this inequality of outcomes should indeed be classified as statistical discrimination, consider the definition set forth in Lundberg and Startz (1983):

> Economic discrimination exists when groups with equal average initial endowments of productive ability do not receive equal average compensation in equilib-

13. For a review of the Becker model and details on testable implications, see Charles and Guryan (2008), who find empirical evidence for the existence of taste-based discrimination in the U.S. labor market.

14. A number of authors have incorporated search frictions into the Becker model to explain the long-run persistence of racial wage gaps. See Lang and Lehmann (2012) for a review of the literature. It is worth noting that racial wage gaps cannot persist in a taste-based model with firm entry; discriminatory firms will always be less profitable.

15. Note that much of the canonical discrimination literature primarily focuses on labor market discrimination, whereby employers discriminate offer lower wages to employees of a particular group. These theories of discrimination can easily be extended to alternative contexts.

rium.

This definition is employed to explicitly account for the fact that the existence of unequal outcomes may impact pre-labor market human capital investment decisions. In the model outlined above, agents with equal levels of initial human capital may not make the same specialization decisions because of differences in initial beliefs, as shown in section 2.4. Students-as-statisticians do not know their true productivity, and they use group-based information for inference. The resulting inequality of outcomes can thus be classified as statistical discrimination.

Recent discrimination literature presents variants of traditional statistical discrimination models, such as dynamic discrimination (Bohren, Imas, and Rosenberg 2019), or inaccurate statistical discrimination (Bohren et al. 2019). Dynamic belief formation, or the influence of inaccurate belief formation, are both extremely relevant to this model, and could influence the self-fulfilling nature. Additionally, thinking about this model as statistical discrimination lays the groundwork to connect this model to the literature on affirmative action. Traditional theoretical literature on affirmative action, such as Coate and Loury (1993), begin with a model of statistical discrimination. The belief updating mechanism in my model still allows for self-fulfilling prophecies, but using a different framework.

2.6 Dynamic extension of model

There is an intuitive connection between the model outlined in sections 2.2 and 2.3 and theories of statistical discrimination. In general, a dynamic framework is needed to formalize this connection, where group-based beliefs are being transmitted across cohorts over time. Below, I discuss some of the theoretical assumptions necessary for creating a dynamic version of this model. Specifically, I present notes on the potential cohort structure that preserves the optimal policy in section 2.6.1. The transmission of group beliefs over time is briefly discussed in section 2.6.2. Formalizing this connection requires a host of additional theoretical assumptions, and is beyond the scope of this paper.

2.6.1 Cohort structure

It's natural to view the simulation exercise outlined in section 2.4 as representing the decision making of a single cohort comprised of two genders. According to those results, a cohort with two group-types, whose agents are identical save for initial group-based beliefs, may see these divergent specialization outcomes. It's possible to extend this analysis to multiple cohorts, where each agent in a new cohort forms their initial beliefs about their ability based on past outcomes for members of their group type. Some steps must be taken to ensure the optimal policy from Alon and Fershtman (2019) outlined in section 2.4 still solves the individual agent's problem. Specifically, one cohort of agents will not begin their education until the previous cohort has chosen their labor specialty and begun working. Below I outline notation that allows for this type of cohort structure in sequential time.

Consider a new cohort of agents that begins studying at time τ. For convenience, I will often refer to this cohort as "cohort τ." In this cohort, there are $N_{g\tau}$ agents of each group type $g \in \{1,\ldots,G\}$. Thus, the total number of agents in cohort τ is given by $N_\tau = \sum_{g=1}^{G} N_{g\tau}$.

Assume that a new cohort only begins their education once the previous cohort has chosen their labor specialty and begun working. Suppose this happens at time $\tau + T(\tau)$ for a cohort beginning their education at time τ. Then this assumption implies that at time $\tau + T(\tau)$:

$$\sum_{i=1}^{N_{g\tau}} \sum_{j=1}^{J} \ell_{ij,\tau+T(\tau)}^{g\tau} = N_{g\tau}, \quad \forall g \in 1,\ldots,G.$$

At that point a new cohort of size $N_{\tau+T(\tau)+1} = \sum_{g=1}^{G} N_{g,\tau+T(\tau)+1}$ begins studying at time $\tau + T(\tau) + 1$. It will often be convenient to let $\tau(+)$ denote the cohort following cohort τ, so cohort $\tau(+)$ begins studying at $\tau(+) = \tau + T(\tau) + 1$. Similarly, let $\tau(-)$ denote the cohort preceding cohort τ. The following table summarizes the above indices, relative to a cohort that begins their education at time τ:

The motivation for this structure lies in the optimal policy outlined in section 2.2.3 and

Notation	Meaning
$\tau(+)$	The cohort that follows cohort τ. If all cohort members who beginning studying at time τ finish studying after $T(\tau)$ periods, then $\tau(+) = \tau + T(\tau) + 1$
$\tau + 1$	One time period after a cohort begins studying at time τ
τ	A cohort that begins studying at time τ will often be referred to as "cohort τ"
$\tau - 1$	Time period immediately preceding τ. Note that this is the time period when all members of the preceding cohort $\tau(-)$ have finished their education begin working.
$\tau(-)$	The cohort that precedes the cohort beginning at time τ

fully explored in Alon and Fershtman (2019). The optimal policy that solves the individual agent's problem requires the following assumptions on the individual agent's learning problem to hold: The first is the monotonicity assumption, which is discussed above. If the agent's problem is monotone, then the one-step-look ahead rule is optimal. The second is the locality of the learning process. Time homogeneity is required for the optimal policy to hold. Intuitively, this means that a skill that is not studied is not improved, and that payoffs from different skills are independent of one another. This assumption places certain constraints on the cohort structure: any cohort version of this model needs to ensure that an individual agent's learning process is not influenced by the experience of other agents. The structure above allows for the full separability of cohort decision making, which preserves the optimality of the policy outlined above, while allowing for complete flexibility in terms of how beliefs are formed and transmitted across cohorts.

2.6.2 Group-based beliefs

Recall that $\theta_{ij}^{g\tau}$ is the true ability in field j for agent i with group type g from cohort τ. Their initial beliefs about their value of $\theta_{ij}^{g\tau}$ are given by the beta distribution with parameters $\left(\alpha_{ij0}^{g\tau}, \beta_{ij0}^{g\tau}\right)$. When an agent i with group type g begins their education at time τ, they form their initial beliefs $\left(\alpha_{ij0}^{g\tau}, \beta_{ij0}^{g\tau}\right)$ from the outcomes of previous cohorts. Given the cohort structure outlined in section 2.6.1, this updating rule is not currently constrained by any particular facet of the model.

It helps to introduce some additional notation to make this more concrete. At time $\tau - 1$, all agents from the the previous cohort $\tau(-)$ have finished their education and make their labor specialization decisions. Let $L_j^{g\tau(-)}$ denote the total number of type g agents from cohort $\tau(-)$ who choose to specialize in field j:

$$L_j^{g\tau(-)} = \sum_{i=1}^{N_{g\tau(-)}} \sum_{j=1}^{J} \ell_{ij,\tau-1}^{g\tau(-)}$$

Note that the specialization decisions of all agents from cohort $\tau(-)$ can be represented by the following $(G \times J)$ matrix:

$$\mathbf{L}^{\tau(-)} = \begin{bmatrix} L_1^{1\tau(-)}, & \dots, & L_J^{1\tau(-)} \\ \vdots, & \ddots, & \vdots \\ L_1^{G\tau(-)}, & \dots, & L_J^{G\tau(-)} \end{bmatrix}$$

We can then define the between cohort initial belief updating rule as a function from initial beliefs and outcomes for cohort $\tau(-)$ to initial beliefs for cohort τ:

$$\Gamma : \left(\alpha_{ij0}^{g\tau(-)}, \beta_{ij0}^{g\tau(-)}, \mathbf{L}^{\tau(-)}, N_{1\tau(-)}, \dots, N_{G\tau(-)}\right) \rightarrow \left(\alpha_{ij0}^{g\tau}, \beta_{ij0}^{g\tau}\right)$$

Once a between-cohort updating rule is established, it is straightforward to simulate outcomes for particular cohorts over time.

It helps to illustrate this using a naive, imperfect, between-cohort belief updating rule in the context of the simulation exercise from section 2.4. Let's specifically consider the baseline simulation illustrated in figure 2.2a, where agents choose between two completely symmetric fields with equal payoffs. Assume that cohort $\tau = 0$ has initial beliefs equal to the uninformative (1, 1) prior. Further assume that, for each cohort τ, their initial group based beliefs are based on the initial beliefs and specialization decisions of the preceding cohort, cohort $\tau(-)$. As a somewhat simplistic example, suppose that the values of α and β only equal integer values between 1 and 5. Let α represent the relative popularity of the field for the group type, and β

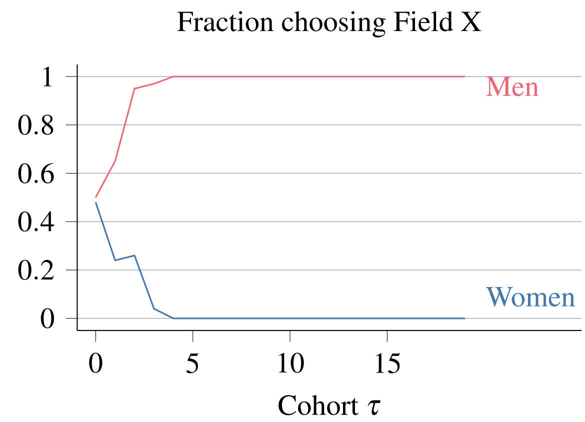

Figure 2.4. Simulation of fraction of men and women in each cohort choosing field X or Y, assuming agents' initial beliefs are based on group outcomes and updated across time.

represent the difference between the expected frequency of field for the group and the realized frequency, scaled to choose outcomes within the grid.[16]

Outcomes for cohorts are simulated in figure 2.4. For each cohort on the x-axis, the y-axis plots the fraction of men and women in that cohort who choose to specialize in field X. Beliefs for the subsequent cohort are then updated to reflect the specialization decision of the preceding cohort. Members of cohort 0 randomly choose between fields X and Y, as they are completely symmetric fields. Slightly more men choose field X than field Y than women due to randomness. However, given the nature of this particular belief updating rule, men are now more likely to specialize in field X in cohort 1. This feedback loop continues until this slight

16. Note that the expected frequency of type g agents in cohort $\tau(-)$ choosing field j is given by the fraction of total agents choosing field j, multiplied by the number in group g, or $E_j^{g\tau(-)} = \frac{L_j^{\tau(-)}}{N_{\tau(-)}} N_{g\tau(-)}$. Define the fraction a as the fraction of all agents from cohort $\tau(-)$ choosing to specialize in j, and the fraction b as a plus the the relative difference between the expected frequency and observed frequency:

$$a = \frac{L_j^{\tau(-)}}{N_{\tau(-)}}, \qquad b = a + \frac{E_j^{g\tau(-)} - L_j^{g\tau(-)}}{N_{g\tau(-)}}.$$

Then initial beliefs in field j for agents with group type g in cohort τ can be given by the above fractions, multiplied by 5 (or whatever the relevant grid size is), and rounded to the closest integer:

$$\alpha_{j0}^{g\tau} = \left\{ x : \min_{1,\dots,5} |a \cdot 5 - x| \right\}, \qquad \beta_{j0}^{g\tau} = \left\{ x : \min_{1,\dots,5} |b \cdot 5 - x| \right\}$$

difference in specialization outcomes has perpetuated into full segregation.

The example illustrated in figure 2.4 is extreme, as the belief updating rule is designed to perpetuate beliefs across cohorts. If beliefs were not transmitted across cohorts, and each cohort simply began with the uninformative prior, we may see the perpetuation of even splits between fields X and Y for men and women across time. Further, agents within a cohort-group are identical; there are no individual idiosyncrasies in preferences or abilities, just randomness in terms of outcomes while in school. Nevertheless, this presents a starting point for how group-based beliefs can act as a transmission mechanism across cohorts.

2.7 Conclusion

In the above chapter, I have outlined a model of group-based beliefs and human capital specialization. This model can be used to show how differential beliefs may perpetuate differences in outcomes between groups. This theoretical framework is meant to be considered in conjunction with the empirical results outlined in Chapter 1. According to the model, if men and women interpret signals about performance in different ways, then we may see the perpetuation of gender gaps in fields with lower female representation. This is consistent with women switching out of STEM fields at higher GPAs than men.

Appendix

A.2.1 Computing agent behavior

This section summarizes how to compute an agent's behavior given the state variables $(\tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}) = (\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0})$ and assuming the initial monotonicity assumption (2.5). Recall that the graduation region under the initial monotonicity assumption is given by (2.11):

$$\mathscr{G}_j = \left\{ \tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0} \middle| \tilde{m}_{jt} \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{jt}}{h_{j0} + \nu_j \tilde{s}_{jt}} \right) - \alpha_{j0} - \beta_{j0} \right\}$$

The index (2.12) under the initial monotonicity assumption (2.5) is given by:

$$\mathscr{I}_j(\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}) = \begin{cases} \frac{w_j}{1-\delta} h_{jt} & (\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}) \in \mathscr{G}_j, \\ \frac{w_j}{1-\delta} \mathbb{E}_t \left[\delta^{m_j^* - \tilde{m}_{jt}} h_{j,t+(m_j^* - \tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0} \right] & \text{otherwise,} \end{cases}$$

1. To find the index for all fields j, first determine whether an agent is in their graduation region (2.11).

   a. If the agent is in their graduation region for field j, set the j-index equal to $\frac{w_{jt} h_{jt}}{1-\delta}$. **Skip to step 6.**

   b. If the agent is not in their graduation region, the index must be calculated using their expected accumulation of human capital, which is a function of expected time remaining in j:

$$\mathbb{E}_t \left[\delta^{m_j^* - \tilde{m}_{jt}} h_{j,t+(m_j^* - \tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0} \right] = \mathbb{E}_t \left[g\left(m_j^* - \tilde{m}_{jt}\right) \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0} \right]$$
$$= \mathbb{E}_t \left[\delta^{m_j^* - \tilde{m}_{jt}} \left(h_{j0} + \nu_j \tilde{s}_{jt} + \left(m_j^* - \tilde{m}_{jt}\right) \frac{\alpha_{j0} + \tilde{s}_{jt}}{\alpha_{j0} + \beta_{j0} + \tilde{m}_{jt}} \right) \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0} \right]$$

   Proceed to the next step.

2. Re-index the problem to simplify analysis:

$$\hat{\alpha}_{j0} = \alpha_{j0} + \tilde{s}_{jt}, \qquad \hat{\alpha}_{j0} + \hat{\beta}_{j0} = \alpha_{j0} + \beta_{j0} + \tilde{m}_{jt},$$
$$\hat{h}_{j0} = h_{j0} + \nu_j \tilde{s}_{jt}, \qquad \hat{\psi}_{j0} = \left(\hat{\alpha}_{j0}, \hat{\beta}_{j0}, \hat{h}_{j0} \right),$$

   This way, we are considering how many courses an agent is expecting to study from time $t = 0$, instead of how many remaining courses an agent expects to study at an arbitrary t.

3. Let $N = m_{jt}^* - \tilde{m}_{jt}$ denote the number of periods an agent expects to study, and let $\underline{n}$ and

$\overline{n}$ denote the lower and upper bound of N, respectively. As described in section A.2.2, these bounds are given by:

$$\underline{n} = \underline{m_{jt}^{*} - \tilde{m}_{jt}} = \left\lceil \frac{\delta}{1-\delta} \right\rceil - \hat{\alpha}_{j0} - \hat{\beta}_{j0}$$

$$\overline{n} = \overline{m_{jt}^{*} - \tilde{m}_{jt}} = \begin{cases} \underline{n} & \text{if } 1 \leq \frac{\hat{\alpha}_{j0}\nu_j}{\hat{h}_{j0}} \leq \frac{\left\lceil \frac{\delta}{1-\delta} \right\rceil}{\frac{\delta}{1-\delta}} \\ \left\lceil \frac{\delta}{1-\delta} \frac{\hat{\alpha}_{j0}\nu_j}{\hat{h}_{j0}} \right\rceil - \hat{\alpha}_{j0} - \hat{\beta}_{j0}, & \text{otherwise.} \end{cases}$$

4. Find the probability distribution of stopping between $\underline{n}$ and $\overline{n}$ conditional on the state variables $\left(\hat{\alpha}_{j0}, \hat{\beta}_{j0}, \hat{h}_{j0}\right)$. First evaluate the probability that $N = \underline{n}$: .

$$\mathbb{P}_0 = \mathbb{P}\left(N = \underline{n} \middle| \hat{\alpha}_{j0}, \hat{\beta}_{j0}, \hat{h}_{j0}\right) = \begin{cases} 1 & \text{if } 1 \leq \frac{\alpha_{j0}\nu_j}{h_{j0}} \leq \frac{\left\lceil \frac{\delta}{1-\delta} \right\rceil}{\frac{\delta}{1-\delta}} \\ 0 & \text{otherwise.} \end{cases}$$

Next evaluate the probability of stopping at $N = \underline{n} + 1$.

$$\begin{aligned} \mathbb{P}_1 =& \mathbb{P}\left(N = \underline{n} + 1 | \, \hat{\psi}_{j0}\right) \\ =& \mathbb{P}\left(N = \underline{n} + 1 | N \neq \underline{n}, \psi_{j0}\right)\left(1 - \mathbb{P}_0\right), \end{aligned}$$

where $\mathbb{P}\left(N = \underline{n} + 1 | N \neq \underline{n}, \psi_{j0}\right)$ is the probability of a binomial random variable. For any integer $x > 1$ such that $\underline{n} + x < \overline{n}$, the probability of stopping at $\underline{n} + x$ is given by:

$$\begin{aligned} \mathbb{P}_x =& \mathbb{P}\left(N = \underline{n} + x | \, \hat{\psi}_{j0}\right) \\ =& \mathbb{P}\left(N = \underline{n} + x | N \neq \underline{n} + x - 1, \hat{\psi}_{j0}\right) \left(1 - \sum_{n=0}^{x-1} \mathbb{P}_n\right), \end{aligned}$$

where $\mathbb{P}\left(N = \underline{n} + x | N \neq \underline{n} + x - 1, \hat{\psi}_{j0}\right)$ is a conditional sum of one-to-one functions of a Bernoulli and a binomial random variable. Finally, the probability of stopping at the upper bound $\overline{n}$ is given by:

$$\mathbb{P}_{\overline{n}-\underline{n}} = \mathbb{P}\left(N = \overline{n} \middle| \hat{\alpha}_{j0}, \hat{\beta}_{j0}, \hat{h}_{j0}\right) = 1 - \sum_{n=0}^{\overline{n}-\underline{n}-1} \mathbb{P}_n$$

See section A.2.2 for details and proofs.

5. Compute expected discounted accumulation of human capital

$$\mathbb{E}_t\left[g(N) | \hat{\psi}_{j0}\right] = \sum_{n=\underline{n}}^{\overline{n}} g(N) \mathbb{P}\left(N = n \middle| \hat{\psi}_{j0}\right)$$

And set the index equal to $\frac{w_j}{1-\delta}\mathbb{E}_t\left[g(N) | \hat{\psi}_{j0}\right]$.

6. Follow the optimal policy outlined in section 2.2.3. If the agent chooses to study this period, return to step one in the next period.

A.2.2 Solving the index

This section discusses the analytical solution to the index (2.12). Specifically, I describe how to evaluate the expected value of discounted human capital accumulation, conditional on initial states:

$$\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} h_{j,t+(m_j^*-\tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \alpha_{j0}, \beta_{j0}, h_{j0}\right]. \tag{2.15}$$

To solve this, I first show how this expectation can be re-written as a function of expected time remaining in school. Computing the index therefore requires finding the conditional probability distribution of stopping times. The remainder of section is devoted to finding this distribution. This is done by first bounding the stopping times, and then recursively defining the probability distribution.

Index in terms of expected time in school

To simplify notation, let ψ_{j0} denote the initial belief parameters and human capital levels:

$$\psi_{j0} = \left(\alpha_{j0}, \beta_{j0}, h_{j0}\right).$$

the agent's state when evaluating field j at time t is now determined by $\left(\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right)$.

Ignoring other fields, an agents expects to study j for $m_j^* - \tilde{m}_{jt}$ additional periods before beginning work as a field-j specialist. Substituting in the human capital accumulation function (2.2) into (2.15):

$$\begin{aligned}
&\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} h_{j,t+(m_j^*-\tilde{m}_{jt})} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\quad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \left(h_{j0} + \nu_j \tilde{s}_{jt} + \sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x}\right) \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\quad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\left(h_{j0} + \nu_j \tilde{s}_{jt}\right) + \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right].
\end{aligned}$$

Two expectations are key. The first is the expected value of discounting the next $m_j^* - \tilde{m}_{jt}$ additional periods. The second is the expected value of the discounted term times the number of of times an agent successfully passes their field-j courses during those $m_j^* - \tilde{m}_{jt}$ periods. To simplify the second probability, use the law of iterated expectations:

$$\begin{aligned}
&\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x} \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\quad = \mathbb{E}_t\left[\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x} \middle| m_j^*, \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\quad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}} \mathbb{E}_t\left[\sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x} \middle| m_j^*, \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \middle| \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]
\end{aligned}$$

The number of times the agent successfully passes their field-j courses over the next $m_j^* - \tilde{m}_{jt}$ periods is a series of $m_j^* - \tilde{m}_{jt}$ Bernoulli trials with probability θ_j. The agent's expected value of this random variable is given by the sample size multiplied by their ability parameter θ_j. Therefore the previous equation can be written as:

$$
\begin{aligned}
&\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\mathbb{E}_t\left[\sum_{x=0}^{m_j^*-\tilde{m}_{jt}} s_{j,t+x}\middle| m_j^*, \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\qquad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\left(m_j^* - \tilde{m}_{jt}\right)\mathbb{E}_t\left[\theta_j\middle| m_j^*, \tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\qquad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\left(m_j^* - \tilde{m}_{jt}\right)\frac{\alpha_{j0}+\tilde{s}_{jt}}{\alpha_{j0}+\beta_{j0}+\tilde{m}_{jt}}\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right] \\
&\qquad = \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\left(m_j^* - \tilde{m}_{jt}\right)\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\frac{\alpha_{j0}+\tilde{s}_{jt}}{\alpha_{j0}+\beta_{j0}+\tilde{m}_{jt}}
\end{aligned}
$$

The third line follows from the agent's expected value of θ_j, conditional on their states and their time until completion, m_j^*, according to their belief distribution. Here, it's worth emphasizing a caveat to this proof, detailed in section A.2.3. If knowing m^* feeds into your expectation of θ_j, then additional assumptions may need to be made, or the following results may not hold. See section A.2.3 for details.

Proceeding, we can use this result to further simplify (2.15) as:

$$
\begin{aligned}
&\mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\left(h_{j0}+\nu_j\tilde{s}_{jt}\right) \\
&\qquad\qquad + \mathbb{E}_t\left[\delta^{m_j^*-\tilde{m}_{jt}}\left(m_j^* - \tilde{m}_{jt}\right)\middle|\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right]\frac{\alpha_{j0}+\tilde{s}_{jt}}{\alpha_{j0}+\beta_{j0}+\tilde{m}_{jt}}.
\end{aligned}
$$

Thus, the key expected value (2.15) is really the expected value of a function of time remaining in school.

Before proceeding, it's helpful to introduce some simplifying notation, and to re-scale the problem to start at time $t = 0$. To simplify notation, let $N = m_j^* - \tilde{m}_{jt}$ denote the time remaining in school after $\tilde{m}_{jt}$. The variable N is capitalized to emphasize the fact that N is a random quantity. Next, note that for any agent evaluating field j at time t with states $\left(\tilde{m}_{jt}, \tilde{s}_{jt}, \psi_{j0}\right)$, we can always define:

$$\begin{aligned} \hat{\alpha}_{j0} &= \alpha_{j0} + \tilde{s}_{jt}, & \hat{\alpha}_{j0} + \hat{\beta}_{j0} &= \alpha_{j0} + \beta_{j0} + \tilde{m}_{jt}, \\ \hat{h}_{j0} &= h_{j0} + \nu_j \tilde{s}_{jt}, & \hat{\psi}_{j0} &= \left(\hat{\alpha}_{j0}, \hat{\beta}_{j0}, \hat{h}_{j0}\right). \end{aligned}$$

Therefore, instead of evaluating how many courses an agent has remaining after completing $\tilde{m}_{jt}$ courses, we can re-define the agent's states and evaluate the agent's total expected time in school from $t = 0$, before the agent has taken any courses in j. In that vein, I will only condition on the initial states $\psi_{j0} = \left(\alpha_{j0}, \beta_{j0}, h_{j0}\right)$, and assume $\tilde{s}_{jt} = \tilde{m}_{jt} = 0$. Recall that $\tilde{s}_{j,N}$ is the number of times the student successfully passes their courses after matriculating N times. Using the above notation, the monotonic stopping condition (2.10) is given by:

$$N \geq \frac{\delta}{1-\delta}\left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{j,N}}{h_{j0} + \nu_j \tilde{s}_{j,N}}\right) - \alpha_{j0} - \beta_{j0}, \tag{2.16}$$

The goal in subsequent sections is to evaluate the following conditional expectations:

$$\begin{aligned} \mathbb{E}_0\left[\delta^N \middle| \alpha_{j0}, \beta_{j0}, h_{j0}\right] &= \sum_{z=0}^{\infty} \delta^z \mathbb{P}\left(N = z \middle| \alpha_{j0}, \beta_{j0}, h_{j0}\right) \\ \mathbb{E}_0\left[\delta^N N \middle| \alpha_{j0}, \beta_{j0}, h_{j0}\right] &= \sum_{z=0}^{\infty} \delta^z z \mathbb{P}\left(N = z \middle| \alpha_{j0}, \beta_{j0}, h_{j0}\right) \end{aligned} \tag{2.17}$$

The following section discusses the bounds on stopping times, so the above summation is finite.

Bounds on stopping times

This section starts by defining a lower bound larger than zero for all stopping times. I then discuss stopping times with positive probability, and use this discussion to motivate the upper bound for all stopping times.

Lemma 1. Define the positive integer $\underline{n}$ as:

$$\underline{n} = \min_N \left\{ N \geq \frac{\delta}{1-\delta} - \alpha_{j0} - \beta_{j0} \right\} = \left\lceil \frac{\delta}{1-\delta} \right\rceil - \alpha_{j0} - \beta_{j0}. \tag{2.18}$$

Then $\underline{n}$ is a lower bound for stopping times that satisfy the monotonic stopping condition (2.16).

Proof. For all possible stopping times N and all possible stochastic outcomes $\tilde{s}_{j,N}$, $\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{jN}}{h_{j0} + \nu_j \tilde{s}_{jN}} \geq 1$ under the initial monotonicity condition (2.5).[17] Therefore, $\underline{n}$ is a lower bound. □

Before defining the upper bound of N, it is helpful to discuss the stopping condition for different potential values of N. Given the stopping condition (2.16), an agent will decide to stop studying in period $\underline{n} + x$ if:

$$\begin{aligned}
& \left\lceil \frac{\delta}{1-\delta} \right\rceil - \alpha_{j0} - \beta_{j0} + x \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}}{h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}} \right) - \alpha_{j0} - \beta_{j0} \\
\Longrightarrow \quad & \frac{\delta}{1-\delta} + \varepsilon_\delta + x \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}}{h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}} \right) \\
\Longrightarrow \quad & \varepsilon_\delta + x \geq \frac{\delta}{1-\delta} \left(\frac{\nu_j \alpha_{j0} - h_{j0}}{h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}} \right) \\
\Longrightarrow \quad & (\varepsilon_\delta + x)\left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x} \right) \geq \frac{\delta}{1-\delta} \left(\alpha_{j0} \nu_j - h_{j0} \right)
\end{aligned} \tag{2.19}$$

where the rounding error $\varepsilon_\delta \in [0,1)$ equals the difference between the ceiling of the discount factor $\frac{\delta}{1-\delta}$ and its true value.[18]

Now let's consider cases where agents would only take $N = \underline{n}$ courses, meaning that

17. Define $f(s) = \frac{\nu_j \alpha_{j0} + \nu_j s}{h_{j0} + \nu_j s} = \frac{\nu_j \alpha_{j0} - h_{j0}}{h_{j0} + \nu_j s} + 1$. Because $\nu_j \alpha_{j0} \geq h_{j0}$ and $s \geq 0$, $f(s) \geq 1$.

18. Specifically, if $\frac{\delta}{1-\delta}$ is not an integer, then $\varepsilon_\delta = \left\lceil \frac{\delta}{1-\delta} \right\rceil - \left\lfloor \frac{\delta}{1-\delta} \right\rfloor - \left\{ \frac{\delta}{1-\delta} \right\} = 1 - \left\{ \frac{\delta}{1-\delta} \right\}$ is the difference between the ceiling of $\frac{\delta}{1-\delta}$ and $\frac{\delta}{1-\delta}$. If $\frac{\delta}{1-\delta}$ is an integer, then $\varepsilon = 0$. Recall that the ceiling of any real number x, denoted $\lceil x \rceil$, is the smallest integer greater than or equal to x. The floor of a x, $\lfloor x \rfloor$, is the largest integer less than or equal to x. The fractional part of x, denoted $\{x\}$, is defined by $\{x\} = x - \lfloor x \rfloor$.

$x = 0$. Using the stopping condition (2.19) with $x = 0$, this only happens if:

$$\begin{aligned} \varepsilon_\delta \left(h_{j0} + \nu_j \tilde{s}_{j\underline{n}}\right) &\geq \frac{\delta}{1-\delta}(\alpha_{j0}\nu_j - h_{j0}) \\ \implies \qquad \varepsilon_\delta \nu_j \tilde{s}_{j\underline{n}} &\geq \frac{\delta}{1-\delta}\nu_j\alpha_{j0} - \left\lceil \frac{\delta}{1-\delta} \right\rceil h_{j0}. \end{aligned}$$

The agent will only study for $\underline{n}$ periods if this inequality is satisfied for all possible stochastic outcomes. The only stochastic part of this inequality is $\tilde{s}_{j\underline{n}}$, which make take on values between 0 and $\underline{n}$. Therefore, agents will only study for $\underline{n}$ periods if:

$$\begin{aligned} 0 &\geq \frac{\delta}{1-\delta}\nu_j\alpha_{j0} - \left\lceil \frac{\delta}{1-\delta} \right\rceil h_{j0}. \\ \implies \left\lceil \frac{\delta}{1-\delta} \right\rceil h_{j0} &\geq \frac{\delta}{1-\delta}\alpha_{j0}\nu_j. \end{aligned}$$

Combining the above with the initial monotonicity condition (2.5) implies that an agent will only study for exactly $N = \underline{n}$ periods if:

$$1 \leq \frac{\nu_j\alpha_{j0}}{h_{j0}} \leq \frac{\left\lceil \frac{\delta}{1-\delta} \right\rceil}{\frac{\delta}{1-\delta}}.$$

Because the ratio of the ceiling of the discount factor to its true value will be close to 1, this inequality effectively states that the agent will only study for $N = \underline{n}$ periods if $h_{j0} = \nu_j\alpha_{j0}$ (with some adjustment for rounding error). This is the tractable case evaluated in Alon and Fershtman (2019). Specifically, assuming the slightly stronger initial condition, $h_{j0} = \nu_j\alpha_{j0}$, implies time spent in school N equals $\underline{n}$; an agent who specializes in field j will take exactly N courses in field j. Therefore, the the optimal number of field-j courses is a deterministic function of the agent's initial beliefs. However, for reasons discussed in the overview of section 2.3.4, this assumption is not necessarily appropriate for this evaluation. Thus, we now turn to evaluating the upper bound of possible stopping times.

Lemma 2. Define the positive integer $\overline{n}$ as:

$$\overline{n} = \left\lceil \frac{\delta}{1-\delta} \frac{\alpha_{j0}\nu_j}{h_{j0}} \right\rceil - \alpha_{j0} - \beta_{j0} \tag{2.20}$$

Then $\overline{n}$ is an upper bound for stopping times.

Proof. $\overline{n}$ is an upper bound for stopping times if, for all stochastic outcomes $\tilde{s}_{j\overline{n}}$:

$$\overline{n} \geq \frac{\delta}{1-\delta} \left(\frac{\alpha_{j0}\nu_j + \nu_j \tilde{s}_{j\overline{n}}}{h_{j0} + \nu_j \tilde{s}_{j\overline{n}}} \right) - \alpha_{j0} - \beta_{j0}$$

Because $\frac{\alpha_{j0}\nu_j + \nu_j \tilde{s}_{j\overline{n}}}{h_{j0} + \nu_j \tilde{s}_{j\overline{n}}}$ is decreasing in $\tilde{s}_{j\overline{n}}$,[19] $\overline{n}$ is an upper bound independent of stochastic outcomes only if the above inequality holds for $\tilde{s}_{j\overline{n}} = 0$ (i.e. when the agent has failed all of their field j courses). Therefore, $\overline{n}$ is an upper bound if:

$$\begin{aligned} \overline{n} \geq & \frac{\delta}{1-\delta} \left(\frac{\alpha_{j0}\nu_j}{h_{j0}} \right) - \alpha_{j0} - \beta_{j0} \\ \implies \left\lceil \frac{\delta}{1-\delta} \frac{\alpha_{j0}\nu_j}{h_{j0}} \right\rceil \geq & \frac{\delta}{1-\delta} \left(\frac{\alpha_{j0}\nu_j}{h_{j0}} \right). \end{aligned}$$

Therefore, $\overline{n}$ is an upper bound for stopping times. □

The details above can be used to bound the summations in (2.17):

$$\begin{aligned} \mathbb{E}\left[\delta^N \middle| \alpha_{j0}, \beta_{j0}, h_{j0} \right] &= \sum_{z=\underline{n}}^{\overline{n}} \delta^z \mathbb{P}\left(N = z \middle| \alpha_{j0}, \beta_{j0}, h_{j0} \right) \\ \mathbb{E}\left[\delta^N N \middle| \alpha_{j0}, \beta_{j0}, h_{j0} \right] &= \sum_{z=\underline{n}}^{\overline{n}} \delta^z z \mathbb{P}\left(N = z \middle| \alpha_{j0}, \beta_{j0}, h_{j0} \right) \end{aligned}$$

The next section evaluates the above conditional probabilities.

19. Define $f(s) = \frac{\nu_j \alpha_{j0} + \nu_j s}{h_{j0} + \nu_j s} = \frac{\nu_j \alpha_{j0} - h_{j0}}{h_{j0} + \nu_j s} + 1$. Note that $f'(s) = -\frac{\nu_j(\nu_j \alpha_{j0} - h_{j0})}{(h_{j0} + \nu_j s)^2}$. This is nonpositive when $\nu_j \alpha_{j0} \geq h_{j0}$.

Conditional probabilities

Recall that the lower and upper bounds of N are given by $\underline{n}$ and $\overline{n}$, respectively. The conditional probability that N equals some integer z can be evaluated as:

$$\begin{aligned}
\mathbb{P}\left(N=z\,\middle|\,\psi_{j0}\right) =& \mathbb{P}\left(z \geq \frac{\delta}{1-\delta}\frac{\alpha_{j0}\nu_j+\nu_j\tilde{s}_{jz}}{h_{j0}+\nu_j\tilde{s}_{jz}}-\alpha_{j0}-\beta_{j0}\,\middle|\,\psi_{j0}\right)\\
=& \mathbb{P}\left(z-(\underline{n}-\varepsilon_\delta) \geq \frac{\delta}{1-\delta}\frac{\alpha_{j0}\nu_j-h_{j0}}{h_{j0}+\nu_j\tilde{s}_{jz}}\,\middle|\,\psi_{j0}\right)\\
=& \mathbb{P}\left((z-\underline{n}+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{jz}\right) \geq \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right)\,\middle|\,\psi_{j0}\right) \qquad (2.21)\\
=& \mathbb{P}\left((z-\underline{n}+\varepsilon_\delta)\,\tilde{s}_{jz} \geq \frac{\delta}{1-\delta}\alpha_{j0}-\left(\left\lceil\frac{\delta}{1-\delta}\right\rceil+z-\underline{n}\right)\frac{h_{j0}}{\nu_j}\,\middle|\,\psi_{j0}\right) \qquad (2.22)
\end{aligned}$$

First consider the case where $N=\underline{n}$. Let $\mathbb{P}_0$ denote this probability. Then, as discussed above:

$$\mathbb{P}_0=\mathbb{P}\left(N=\underline{n}\,\middle|\,\psi_{j0}\right)=\begin{cases}1 & \text{if } 1\leq\frac{\alpha_{j0}\nu_j}{h_{j0}}\leq\frac{\left\lceil\frac{\delta}{1-\delta}\right\rceil}{\frac{\delta}{1-\delta}}\\ 0 & \text{otherwise.}\end{cases}$$

Before evaluating cases where $N>\underline{n}$, note the probability of stopping at some positive integer z is always given by:

$$\begin{aligned}
\mathbb{P}\left(N=z|\,\psi_{j0}\right) &= \mathbb{P}\left(N=z,N\neq z-1|\,\psi_{j0}\right)+\mathbb{P}\left(N=z,N=z-1|\,\psi_{j0}\right)\\
&= \mathbb{P}\left(N=z|N\neq z-1,\psi_{j0}\right)\mathbb{P}\left(N\neq z-1|\,\psi_{j0}\right)
\end{aligned}$$

The second line follows from Bayes' Rule, and the fact that an agent would never stop at time z if they already stopped at time $z-1$. In words, this states that the probability of stopping at time z is given by the product of (1) the probability of stopping at time z conditional on having not stopped at time $z-1$; and (2) the probability of having not stopped by $z-1$.

To evaluate the probability that $N=\underline{n}+1$, first evaluate the conditional probability that

$N = \underline{n} + 1$, conditional on the stopping time not equaling $\underline{n}$:

$$\begin{aligned}
\mathbb{P}\left(N = \underline{n}+1 | N \neq \underline{n}, \psi_{j0}\right) =& \mathbb{P}\left((1+\varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j\underline{n}+1}\right) \geq \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j - h_{j0}\right) \middle| \psi_{j0}\right) \\
=& \mathbb{P}\left(\tilde{s}_{j\underline{n}+1} \geq \frac{1}{\nu_j}\left(\frac{\delta}{1-\delta}\frac{1}{1+\varepsilon_\delta}\left(\alpha_{j0}\nu_j - h_{j0}\right) - h_{j0}\right) \middle| \psi_{j0}\right) \\
=& \mathbb{P}\left(\tilde{s}_{j\underline{n}+1} \geq \hat{k}_1 \middle| \psi_{j0}\right)
\end{aligned}$$

The RHS of the above inequality is a function of initial parameters and can be treated as a constant. The variable $\tilde{s}_{j\underline{n}+1}$ is a random variable:

$$\tilde{s}_{j,\underline{n}+1} = \sum_{t=0}^{\underline{n}} s_{jt} \sim \text{Binomial}(\underline{n}+1, \theta_j)$$

Define k_1 as:

$$k_1 = \begin{cases} \hat{k}_1 - 1 & \text{if } \hat{k}_1 \text{ an integer,} \\ \lfloor \hat{k}_1 \rfloor & \text{otherwise.} \end{cases}$$

The conditional probability can be written as:

$$\begin{aligned}
\mathbb{P}\left(N = \underline{n}+1 | N \neq \underline{n}, \psi_{j0}\right) =& 1 - \mathbb{P}\left(\tilde{s}_{j\underline{n}+1} < k_1 \middle| \psi_{j0}\right) \\
=& 1 - \sum_{i=0}^{k_1} \binom{\underline{n}+1}{i} \theta^i (1-\theta)^{\underline{n}+1-i}
\end{aligned}$$

Now we can fully evaluate the probability that $N = \underline{n} + 1$:

$$\begin{aligned}
\mathbb{P}_1 = \mathbb{P}\left(N = \underline{n}+1 \middle| \psi_{j0}\right) =& \mathbb{P}\left(N = \underline{n}+1 | N \neq \underline{n}, \psi_{j0}\right) \mathbb{P}\left(N \neq \underline{n} \middle| \psi_{j0}\right) \\
=& \mathbb{P}\left(N = \underline{n}+1 | N \neq \underline{n}, \psi_{j0}\right)(1 - \mathbb{P}_0) \\
=& \left(1 - \sum_{i=0}^{k_1} \binom{\underline{n}+1}{i} \theta^i (1-\theta)^{\underline{n}+1-i}\right)(1 - \mathbb{P}_0).
\end{aligned}$$

To evaluate the probability that $N = \underline{n} + x$ for some integer x, we have to evaluate:

$$\mathbb{P}_x = \mathbb{P}\left(N = \underline{n} + x \middle| \psi_{j0}\right) = \mathbb{P}\left(N = \underline{n} + x \middle| \bigcap_{i=0}^{x-1} (N \neq \underline{n} + i), \psi_{j0}\right) \mathbb{P}\left(\bigcap_{i=0}^{x-1} (N \neq \underline{n} + i) \middle| \psi_{j0}\right)$$

First consider the probability that N is not equal any value between the lower bound $\underline{n}$ and $\underline{n} + x - 1$. By De Morgan's law and countability additivity:

$$\begin{aligned}\mathbb{P}\left(\bigcap_{i=0}^{x-1} (N \neq \underline{n} + i) \middle| \psi_{j0}\right) &= 1 - \mathbb{P}\left(\bigcup_{i=0}^{x-1} (N = \underline{n} + i) \middle| \psi_{j0}\right) \\ &= 1 - \sum_{i=0}^{x-1} \mathbb{P}\left(N = \underline{n} + i \middle| \psi_{j0}\right) = 1 - \sum_{i=0}^{x-1} \mathbb{P}_i.\end{aligned}$$

Now we turn to the conditional probability of interest, which can be written as:

$$\begin{aligned}&\mathbb{P}\left(N = \underline{n} + x \middle| \bigcap_{i=0}^{x-1} (N \neq \underline{n} + i), \psi_{j0}\right) \\ &= \mathbb{P}\Bigg((x + \varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+x}\right) \geq \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j - h_{j0}\right) \\ &\qquad \left| \bigcap_{k=0}^{x-1} (k + \varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j\underline{n}+k-1}\right) < \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j - h_{j0}\right)\right.\Bigg).\end{aligned} \tag{2.23}$$

To simplify this conditional, recall that an agent will decide to keep studying at time $\underline{n} + y$ if:

$$\begin{aligned}\frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j - h_{j0}\right) &> (y + \varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+y}\right) \\ &= (y - 1 + \varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+y-1}\right) \\ &\quad + \left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+y-1}\right) + (y + \varepsilon_\delta)\nu_j s_{j,\underline{n}+y-1} \\ &> (y - 1 + \varepsilon_\delta)\left(h_{j0} + \nu_j \tilde{s}_{j,\underline{n}+y-1}\right)\end{aligned}$$

This is simply the monotonicity of the stopping problem in reverse; if an agent would decide to continue on at time $\underline{n} + y$, then they also would have wanted to continue on at time $\underline{n} + y - 1$.

This simplifies the conditional expression in equation (2.23):

$$\begin{aligned}
&\mathbb{P}\left(N=\underline{n}+x \middle| \bigcap_{i=0}^{x-1}(N\neq \underline{n}+i), \psi_{j0}\right)\\
&=\mathbb{P}\Bigg((x+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{j,\underline{n}+x}\right)\geq \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right)\\
&\qquad\qquad \Bigg| (x-1+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{j\underline{n}+x-1}\right)<\frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right)\Bigg)
\end{aligned}$$

This probability can be re-written to reflect the fact that $\tilde{s}_{j,\underline{n}+x}=\tilde{s}_{j,\underline{n}+x-1}+s_{j,\underline{n}+x-1}$. In words, the total number of successes seen by time $\underline{n}+x$ equals the total number of successes seen by time $\underline{n}+x-1$ plus the course outcome during period $\underline{n}+x-1$:

$$\begin{aligned}
&\mathbb{P}\Bigg((x+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{j,\underline{n}+x-1}\right)+(x+\varepsilon_\delta)\,\nu_j s_{j,\underline{n}+x-1}\geq \frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right)\\
&\qquad\qquad\qquad \Bigg| (x-1+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{j\underline{n}+x-1}\right)<\frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right)\Bigg).
\end{aligned}$$

Define the random variables Y and Z and the constant c as:

$$\begin{aligned}
&Y=g(\tilde{s}_{j,\underline{n}+x-1})=(x+\varepsilon_\delta)\left(h_{j0}+\nu_j\tilde{s}_{j,\underline{n}+x-1}\right), \qquad c=\frac{\delta}{1-\delta}\left(\alpha_{j0}\nu_j-h_{j0}\right),\\
&Z=h(s_{j,\underline{n}+x-1})=(x+\varepsilon_\delta)\,\nu_j s_{j,\underline{n}+x-1}
\end{aligned}$$

The conditional probability that $N=\underline{n}+x$ for $x>1$ can now be written as:

$$\mathbb{P}\left(Y+Z\geq c|Y<c\frac{x+\varepsilon_\delta}{x-1+\varepsilon_\delta}\right),$$

where Y and Z are independent random variables whose distributions are one-to-one functions of binomial distributions:

$$\begin{aligned}\mathbb{P}(Y=y) =&\mathbb{P}\left(g\left(\tilde{s}_{j,\underline{n}+x-1}\right)=y\right)=\mathbb{P}(\tilde{s}_{j,\underline{n}+x-1}=g^{-1}(y)))\\ =&\binom{\underline{n}+x-2}{g^{-1}(y)}\theta_j^{g^{-1}(y)}(1-\theta_j)^{\underline{n}+x-2-g^{-1}(y)},\\ \mathbb{P}(Z=z) =&\mathbb{P}\left(h\left(s_{j,\underline{n}+x-1}\right)=z\right)=\mathbb{P}\left(s_{j,\underline{n}+x-1}=h^{-1}(z)\right)\\ =&\theta_j^{h^{-1}(z)}\left(1-\theta_j\right)^{1-h^{-1}(z)}.\end{aligned}$$

The joint conditional distribution can be solved using Theorem 20.3 in Billingsley (2012, pg. 280).

A.2.3 Caveat in proof

When detailing how to solve the Index in terms of expected time in school in section A.2.2, a caveat needs to be made. To keep the presentation compact, I include all relevant equations below for a minimum working example below; field indices j are omitted for clarity.

Suppose human capital evolves in line with equation (2.2):

$$h_{t+1}=h_t+\nu s_t m_t, \qquad s_t\sim\text{Bernoulli}(\theta), \tag{2.24}$$

Let P_t denote an agent's beliefs about θ at time t. Assume initial beliefs $P_0=\mathscr{B}(\alpha_0,\beta_0)$, where $\mathscr{B}(\alpha_0,\beta_0)$ is a beta distribution with initial parameters α_0 and β_0. Further, assume beliefs update according to Bayes' Rule, so:

$$P_{t+1}=\mathscr{B}(\alpha_{t+1},\beta_{t+1}), \qquad (\alpha_{t+1},\beta_{t+1})=\begin{cases}(\alpha_t+1,\beta_t) & \text{if } s_t=1\\ (\alpha_t,\beta_t+1) & \text{if } s_t=0\end{cases} \tag{2.25}$$

I want to evaluate, for some T:

$$\mathbb{E}\left[\delta^T h_T\,\middle|\,\alpha_0,\beta_0\right],$$

where $\delta \in (0,1)$ is the discount rate. Assume h_0 is given. By the human capital accumulation function (2.24):

$$\mathbb{E}\left[\delta^T h_T \middle| \alpha_0, \beta_0\right] = \mathbb{E}\left[\delta^T \left(h_0 + \nu \sum_{t=0}^{T} s_T\right) \middle| \alpha_0, \beta_0\right]$$

If T is known, this simplifies to:

$$\begin{aligned}
\mathbb{E}\left[\delta^T h_T \middle| \alpha_0, \beta_0\right] &= \delta^T h_0 + \delta^T \nu \mathbb{E}\left[\sum_{t=0}^{T} s_T \middle| \alpha_0, \beta_0\right] \\
&= \delta^T h_0 + T\delta^T \nu \mathbb{E}\left[\theta \middle| \alpha_0, \beta_0\right] \\
&= \delta^T h_0 + T\delta^T \nu \frac{\alpha_0}{\alpha_0 + \beta_0}
\end{aligned}$$

In an optimal stopping time problem, T is endogenously determined by stochastic outcomes. This complicates the above expectations, since T is now random. Using the law of iterated expectations:

$$\begin{aligned}
\mathbb{E}\left[\delta^T h_T \middle| \alpha_0, \beta_0\right] &= \mathbb{E}\left[\mathbb{E}\left[\delta^T \left(h_0 + \nu \sum_{t=0}^{T} s_T\right) \middle| T, \alpha_0, \beta_0\right] \middle| \alpha_0, \beta_0\right] \\
&= \mathbb{E}\left[\delta^T h_0 + \delta^T \nu \mathbb{E}\left[\sum_{t=0}^{T} s_T \middle| T, \alpha_0, \beta_0\right] \middle| \alpha_0, \beta_0\right] \\
&= \mathbb{E}\left[\delta^T h_0 + T\delta^T \nu \underbrace{\mathbb{E}\left[\theta \middle| T, \alpha_0, \beta_0\right]} \middle| \alpha_0, \beta_0\right]
\end{aligned}$$

I now turn my attention to the conditional expectation, $\mathbb{E}[\theta | T, \alpha_0, \beta_0]$. Intuitively, this is the expected probability of success, conditional on stopping after T periods. If you do not get any information about θ from knowing T, then:

$$\mathbb{E}\left[\theta \middle| T, \alpha_0, \beta_0\right] = \frac{\alpha_0}{\alpha_0 + \beta_0}.$$

This is in line with the initial beliefs assumption, and how I proceed with the proof in section

A.2.2. However, there is an intuitive case to be made that knowledge of T influences the conditional expectation. Comments on this problem would be more than welcome.

Chapter 3

Agricultural Climate Change Adaptation: A review of recent approaches

Chapter 3 is a reprint of the material published in the Environmental Defense Fund Economic Discussion Paper Series in June 2021. This work would not have been possible without financial support from the Environmental Defense Fund Pre-Doctoral Internship. I am particularly thankful for the help of my summer advisor, Jeremy Proville, along with the Chief Economist at EDF, Suzi Kerr, and the staff of the Office of the Chief Economist for enabling such a productive experience. The dissertation author was a primary author of this chapter.

1. Introduction

Estimating the impact of climate change on food production is not straightforward. Overall, a warming climate appears to have a negative impact on yields of maize, wheat, rice and soybeans (Field et al., 2014). However, increasing temperature is not the sole metric of climate change; other relevant changes in climatic systems include shifts in precipitation conditions, extreme weather and atmospheric composition (Stocker et al., 2013). The interaction of these physical processes will directly impact agricultural outcomes.

This link is further complicated by adaptation, or the response of economic agents to realized or expected climate change.[1] Consumers, producers and governments may respond to climate change by, for example, adjusting production technologies, improving institutional capacity or participating in global food systems. Accounting for these adjustments is central to accurately estimating the impact of climate change on agricultural outcomes. However, it is often difficult to measure the role of adaptation within these causal chains.

The definition of climate change adaptation is purposefully vague and encompasses a host of different activities.[2] Occasionally, these activities are explicit, discrete actions; for example, a farmer may experience several seasons of elevated temperatures and choose to plant heat-resistant seed varieties. But the identification of seemingly explicit climate change adaptation is not straightforward, as these reponses are the result of heterogeneous agents making complex optimization choice. Further, many adaptive actions may be implicit responses to changing conditions.[3] For example, a region experiencing higher temperatures may shift labor away from agriculture and begin relying on imported food to meet its needs. In either case, estimating adaptive behavior from observed data is empirically complicated and prone to endogeneity concerns, as will be discussed in the next section.

---

[1] Economists may define adaptation slightly differently than scientists or policy makers. It is common in policy realms for resilience and adaptation to be defined separately; see Walker and Salt (2006), who define resilience as encompassing adaptation (specifically, resilience comprises resistance, adaptation and transformation). Economists tend to conflate the definition of resilience and adaptation. See Fankhauser (2017) for a discussion on this.
[2] To illustrate this, see the list of adaptation options in Field et al. (2014), table 14-1.
[3] Some subsets of the literature have attempted more detailed taxonomies of adaptation. For example, Burke and Lobell (2010) distinguish between autonomous and planned adaptation. However, as discussed in Fankhauser (2017), classifying any adaptation as autonomous is a misnomer, as “autonomous” adaptation is the result of complex optimization decisions made by multiple agents. Overall, specificity is required when discussing adaptation strategies.

In this report, I review recent approaches to studying agricultural climate change adaptation. I consider two broad methodologies. The first of these is panel data techniques, which use available weather data to identify whether adaptations have occurred. I then consider spatial general equilibrium models, which can be used to predict future adaptive behaviors. My intention is to provide the reader with a clear overview of how these tools have been used to study agricultural climate change adaptation, a discussion of the benefits and limitations of these approaches, and a list of resources that can be consulted for future inquiries.

Before proceeding, I would like to make a few clarificatory points. First, this review will generally focus on the response of agents to rising temperatures. As noted above, this is far from the only climate change metric. However, recent innovations in the agricultural adaptations literature outlined in this report specifically estimate responses to increasing temperatures.

Second, it's worth mentioning that the approaches reviewed differ from each other in meaningful ways. Both panel data methods and spatial general equilibrium models estimate the adaptive behaviors arising from heterogeneous agents making complex optimization choices. But panel data methods are often retrospective, estimating past adaptive behavior from observed data. In contrast, spatial general equibirum frameworks attempt to model expected responses to predicted climate outcomes. Thus, the methods discussed below differ both in purpose and in scope.

Finally, spatial general equilibrium models may be seen as a subset of recent advances in macroeconomic climate change analysis. Structural models that account for dynamics or general equilibrium effects have become powerful tools in climate change research. This review focuses on a small subset of this literature that specficially pertains to agricultural adaptations. As such, I do not touch on some key innovations in this field.[4] Nevertheless, I hope this review provides a concrete introduction to one dimension of this evolving toolkit.

I begin by discussing the role that agricultural adaptive behaviors have played in climate impact assessments and the development of panel data methods for studying these behaviors. I then review the use of spatial general equilibrium models in studying adaptive behaviors in agriculture.

[4] For work that focuses on bridging the gap between macroeconomic and microeconomic climate analysis, see Bakkensen and Barrage (2020, 2021). For an example of quantitative macroeconomics being used to model adaptation behaviors, see Fried (2021).

2. Impact assessments and adaptation

Insights about climate change adaptation are often closely tied to studies that estimate the effect of changing temperature on agricultural outcomes (Fankhauser, 2017).[5] For example, consider early impact assessments that employ the so-called "production function" approach. These studies use carefully constructed production functions to estimate the physical response of, say, crop yields to increased temperatures. However, these estimates assume farmers do not adapt any aspect of their production process to changing temperatures, and the assessments may therefore be biased by basic adaptation strategies, such as switching crop varieties. For this reason, the production function approach is sometimes called the "dumb-farmer approach" (Schneider et al., 2000).

The Ricardian approach developed by Mendelsohn et al. (1994) addresses this criticism. Instead of measuring the impact of temperature on agricultural yields, this seminal paper uses cross-sectional data to estimate the impact of temperature on the value of farmland. The authors assume that the value of farmland reflects the best possible use of land, and that farmers have perfect knowledge of this best-use value function. This is an improvement over the production function approach, as the best-use value function implicitly accounts for adaptations such as alternative crop production or substitution of inputs. In practice, the authors use cross-sectional data on agricultural outcomes (y_i), climatic variables ($\mathbf{c}_i$) and agricultural inputs and controls ($\mathbf{x}_i$) for various regions to estimate the following:

$$y_i = \alpha + \beta \mathbf{c}_i + \gamma \mathbf{x}_i + \epsilon_i \tag{1}$$

Here, α is a constant, and ϵ_i is the unobserved prediction error. This specification is identified under the "unit homogeneity" assumption (Hsiang, 2016).[6] Unfortunately, however, this assumption is not always plausible and is prone to endogeneity concerns. For instance, temperature may be correlated with unobserved characteristics like institutional quality (Acemoglu et al., 2002). As such, the specification may be vulnerable to omitted variable bias. In addition, this method merely accounts for adaptation; it does not attempt to measure it explicitly.

---

[5] The review of the impact assessments literature in this section follows Hsiang (2016) and Fankhauser (2017).

[6] The identifying assumption is that differences in output between areas with the same characteristics/inputs ($\mathbf{x}$) are driven solely by differences in climate. Hsiang (2016) describes this assumption as "unit homogeneity."

To address endogeneity concerns, climate econometricians began incorporating panel methods into their impact assessments. In addition to providing more plausible identification assumptions, panel methods enable researchers to estimate adaptation efforts. To explicitly measure adaptation instead of just accounting for it, climate economists have recently turned to structural models and general equilibrium analysis.

3. Panel methods

Deschênes and Greenstone (2007) took the endogeneity concerns associated with the Ricardian approach seriously. Their impact assessments measured the response of land values to changes in weather over time by estimating the following specification:

$$y_{it} = \alpha_i + \theta_t + \beta_{FE}\mathbf{c}_{it} + \gamma\mathbf{x}_{it} + \epsilon_{it} \tag{2}$$

Here, α_i are region-fixed effects and θ_t are time-fixed effects.[7] Regional-fixed effects account for unobserved heterogeneity that is constant over time. Therefore, α_i corrects for some of the omitted variable bias present in the cross-sectional specification (1). Naturally, the fixed-effects approach may still be vulnerable to endogeneity concerns.[8] In addition, the approach in (2) estimates the expected response of variable y to marginal changes in *weather*, conditional on region- and time-fixed effects. But weather is not the same thing as climate; if we wish to estimate the impact of climate change on outcomes, we must assume that this is equal to the expected response of y to changes in *climate*. Thus, the fixed-effects approach relies on the assumption that response to short-run changes in weather are comparable to the response to changes in climate.[9]

Long differencing serves as a compromise between the cross-sectional and fixed-effects approaches. Variables are averaged over two distinct periods in time (t_1 and t_2). The impact is thus estimated from differences in averages over the two time periods:

---

[7] I focus on region- and time-fixed effects for notational simplicity. In practice, Deschênes and Greenstone (2007) include county and state time-fixed effects in their preferred specification. In general, two-way fixed effects may be included.
[8] For example, weather data are collected from weather stations that may have differences in coverage due to exogenous policy shocks. If the availability of weather data and economic outcomes is correlated, results may suffer from attenuation bias. See Auffhammer et al. (2013) for a discussion on this.
[9] Hsiang (2016) calls this the marginal treatment comparability assumption, and notes that this assumption is weaker than the unit homogeneity assumption necessary in (1).

$$\bar{y}_{i,t_2} - \bar{y}_{i,t_1} = \alpha_i + \beta_{LD}(\bar{\mathbf{c}}_{i,t_2} - \bar{\mathbf{c}}_{i,t_1}) + \gamma(\bar{\mathbf{x}}_{i,t_2} - \bar{\mathbf{x}}_{i,t_1}) + \epsilon_i \tag{3}$$

This method estimates the impact of long-term changes in climate using the cross-sectional correlations.

None of the methods discussed so far attempts to explicitly measure adaptation. Both of the time-series specifications above are used to estimate response functions; the coefficients of interest, β_{FE} and β_{LD}, each estimate how agricultural outcomes respond to temperature. But the estimated responses are calculated from different variations in the data. The fixed-effect coefficient, β_{FE}, is estimated from short-run variation in weather, whereas the long-differences coefficient, β_{LD}, is estimated from short- and long-run responses to weather.

Burke and Emerick (2016) use this distinction to derive insights about adaptation. If farmers do not adapt to long-run changes in weather, then the response of crop outcomes to weather in the short run would be the same, suggesting that $\beta_{LD} = \beta_{FE}$. If farmers adapt to long-run changes in weather, then one would expect that the impact of weather on outcomes could be mitigated somewhat in the long run. This would suggest that β_{FE} is larger than β_{LD}.[10] Burke and Emerick use this insight to define an adaptation measure equal to the share of short-run impacts that are offset in the long run:

$$\frac{\beta_{FE} - \beta_{LD}}{\beta_{LD}} = 1 - \frac{\beta_{LD}}{\beta_{FE}} \tag{4}$$

They then test the null hypothesis that this adaptation measure equals zero. Their simulations of soy and maize outcomes suggest limited adaptation to extreme heat in the U.S.

This insight provides a foundation for using response functions derived from time-series analyses to estimate adaptation. Response functions can be estimated in different physical contexts; for example, Burke and Emerick (2016) measure the response of yields to extreme heat according to short-run variation in the data, and according to short- and long-run variation in the data. Different responses to physically similar events in dissimilar contexts may indicate the presence or absence of adaptation (Schlenker and Roberts, 2009; Auffhammer and Schlenker,

[10] To fully motivate this intuition, see section A.2.1 of Burke and Emerick (2016).

2014). In particular, similar responses may indicate the existence of an adaptation gap, whereby certain adaptations could improve outcomes over time (Carleton and Hsiang, 2016).

This methodology could be expanded by improving the estimation of response functions. These are often calculated in the adaptation literature using single-equation models.[11] However, it may be more appropriate to consider economic outcomes as part of a larger system experiencing structural shocks. An extensive macroeconomic literature estimates how responses to some impulses propagate through an economy.[12] Using tools such as structural vector autoregression could help build more realistic response functions. Then, as noted above, these response functions could be used to identify adaptation gaps.

More generally, the climate adaptation literature has begun combining short-run responses to weather within a larger aggregate framework, in line with that of an integrated assessment model (Auffhammer, 2018; Carleton et al. 2020). This allows researchers to estimate long-run responses to climate change, while allowing for spatial heterogeneity in those responses in the cross-section. This methodology provides a promising framework for analyzing agricultural climate impacts while accounting for long-run adaptive behaviors (Moore and Lobell, 2014). Further, these types of heterogenous responses can then be embedded in aggregate structural models, including spatial general equilibrium models, to capture heterogeneity in certain types of adaptations (Nath, 2020). Thus, the panel approaches discussed above may also be used in conjunction with the structural methods discussed below.

4. Spatial general equilibrium models

The Ricardian approach relies on the assumption that local farmers employ the optimal production adaptations for their particular farm. However, many adaptations are not the result of single agent's optimization of a plot of land. Adaptations such as international trade may be the result of intersecting market forces. Policy makers may want to know if international trade can buffer against the adverse effects of climate change on agricultural outcomes. But reduced-

[11] This includes the response functions estimated in Burke and Emerick (2016) and most of the functions featured in Carleton and Hsiang (2016).
[12] See Ramey (2016) for a comprehensive review, and Stock and Watson (2016) for a precise review of identification techniques.

form assessments of trade may be biased by general equilibrium effects. Therefore, climate economists should consider employing other techniques.

In many contexts, researchers are interested in studying economic interactions across physical space; these approaches can be broadly classified as spatial general equilibrium models. For example, trade models (a type of spatial general equilibrium model) are often used to study how differences in technology drive production specialization and bilateral trade flows. This section reviews how spatial general equilibrium models can be used to study climate change adaptations. This is an extremely broad class of models; thus, to make this review concrete, I focus on how international trade in agricultural goods can be used to adapt to climate change.[13] Therefore, I will first give a brief overview of some of the methodological advantages and disadvantages of trade models that are relevant to understanding climate change adaptation. I then discuss the small but growing literature that uses trade models to estimate how international trade can be employed to adapt to climate change. I conclude this section by discussing spatial general equilibrium more broadly, and identify possible research gaps.

4.1 Overview of trade models

Before describing how trade models are used to understand climate change adaptation, it is useful to review some important elements about trade models more generally. For the purposes of this review, I use the term “trade model” to refer to Ricardian gravity trade models.[14] In a Ricardian model of trade, countries are endowed with different technologies for producing goods. A country has an absolute advantage at producing a good if it is more productive at doing so than other countries. A country has a comparative advantage over other countries in producing a good if it can produce it at a lower relative opportunity cost. In a Ricardian model, countries specialize in producing goods according to their comparative advantage. In a gravity trade model, size and distance have multiplicative impacts on bilateral flows.[15] Thus, a Ricardian

[13] It is worth highlighting terminology here. This section reviews how spatial general equilibrium models can be used to study climate change adaptation. To make this review concrete, I will focus on how a particular category of spatial general equilibrium models (trade models) can be used to understand a particular type of climate change adaptation (international trade). Thus, the methodological tool of trade models should be understood as distinct from international trade as a type of adaptation.

[14] In fact, most of the models discussed here are variants of the Eaton and Kortum (2002) framework. A significant advantage of this model over other types of quantitative general equilibrium models is parsimony; counterfactual analysis can be conducted with only one structural parameter estimate (Adao et al., 2017).

[15] Gravity trade models are meant to capture two key empirical facts often seen in international trade data: (1) trade is proportional to size, and (2) trade is inversely proportional to distance. See Head and Mayer (2014) for a review.

gravity trade model focuses on how comparative advantage and geographical barriers govern trade patterns.

Trade models are designed to capture productivity heterogeneity across space. This makes them well poised to study climate change, as this is expected to have differential impacts on productivity across different regions. For instance, some areas of the world are expected to become better at producing certain crops, whereas other regions are expected to become worse. Crucially, data are available to measure these differences in productivity. Consider the Food and Agriculture Organization's Global Agro-Ecological Zones (GAEZ) database,[16] which contains detailed data on agricultural resources. Of particular interest are the crop-specific potential yield variables.[17] These potential yields are calculated for all crops, not just crops being grown. As such, they can be used to estimate a region's agricultural comparative advantage in a trade model.[18] Given this productivity data, economists can calibrate their trade model.

Once a trade model is estimable, it can be used for counterfactual analysis. Trade models are designed to capture how heterogeneity in productivity leads to particular trade patterns in equilibrium. In general, measures of consumer welfare can be recovered from these models. Assuming a trade model is plausible, an economist can calibrate their model in a world that allows for international trade, and in a world that does not. Because most trade models approximate consumer welfare, it is possible to calculate welfare in a world with and without international trade. The difference between these two arms is called the "gains from trade." This type of counterfactual analysis is a major methodological advantage of trade approaches.

These advantages come at a cost. For parsimony, we generally need strong functional form assumptions that may not be realistic. In addition, counterfactual analysis often describes the steady states, not the transitions. Given the pace at which climate change occurs, transitions may be crucial for understanding economic development. Models are also necessarily limited in scope; they cannot account for all the real-world frictions that might prevent the steady state

---

[16] See http://www.fao.org/nr/gaez/en/ for details.
[17] These estimates of crop production potential are constructed using high-resolution measures of crop characteristics, climatic variables (i.e., temperature, precipitation), soil resources, water supply systems and estimated management intensity. Thus, the GAEZ database provides agronomically possible crop yields for 49 different crops across the globe.
[18] This methodology was first employed in Costinot and Donaldson (2012) to evaluate the predictions of the Ricardian model. It has been used to study a number of scenarios; for example, Costinot and Donaldson (2016) estimate the impact of the U.S. railroad system and subsequent market integration on agricultural markets. They do this by applying a Ricardian framework to the U.S. and using productivity estimates from GAEZ.

ever being reached. International trade economists are well acquainted with these trade-offs; some precision is sacrificed to gain an idea about the broader dynamics at play. Even though this type of analysis may never perfectly predict the future, it can, perhaps, provide bounds on impacts or guide thinking about future dynamics. Nevertheless, policy makers and researchers alike are interested in these outcomes, and there is therefore a small but growing literature using trade models to measure agricultural climate change adaptation.

4.2 Application to climate change adaptation

Trade models are adept at summarizing how spatial heterogeneity in crop productivity drives international trade in agricultural goods. However, climate change will have a differential impact on crop productivities across the world, which may shift regional comparative advantage. Given data on how climate change impacts crop productivities, one can estimate subsequent shifts in comparative advantage. This can then be used in a trade model to estimate trading patterns in a world experiencing climate change.

This type of approach was first developed by Costinot et al. (2016). In this paper, the authors use a Ricardian trade model and GAEZ data to estimate how production and trading patterns shift with climate change. Comparative advantage is evaluated using GAEZ potential yield data, which are available for the current climate and estimated under various Intergovernmental Panel on Climate Change scenarios for climate change.[19] These data on comparative advantage can be used in a model of agricultural trade to evaluate various counterfactual scenarios. Specifically, the authors are able to estimate the expected impact of climate change on consumer welfare once trade and production adjustments are taken into account.

The framework provided by Costinot et al. (2016) leaves room for refinement and alternative mechanisms. For example, Nath (2020) employs a similar framework but considers another type of adaptation: sectoral reallocation. Many low-income, agrarian economies are located in hot climates that are particularly susceptible to climate change. However, these areas could adapt to climate change by shifting out of agriculture and into manufacturing or services; food

[19] The GAEZ database is not the only set of data available to study this phenomenon. Cline (2007) estimates the effect of climate change on agricultural productivity, but accounts for certain types of common adaptations within the agricultural sector. This can be useful, as the GAEZ data require researchers to make assumptions about production decisions such as irrigation and input quality that may not be representative of local conditions. For this reason, these are the data used for counterfactual analysis in Nath (2020).

could then be imported from regions that are less impacted. However, this dynamic is complicated by the fact that the agricultural sector in low-income countries is often less productive relative to that of high-income countries, yet maintains a high share of the labor population due to subsistence needs (Lagakos and Waugh, 2013). If climate change drives more workers into agriculture in low-productivity countries due to subsistence requirements, then sectoral reallocation may not occur. Nath builds and estimates a model of international trade and sectoral reallocation to forecast which effect may dominate. His results suggest that climate change may drive more people into subsistence farming, and that current trade policies are not open enough to prevent the subsistence pull toward agriculture. Thus, trade openness may be a key policy for encouraging sectoral reallocation as an adaptation strategy.

Recent advances in trade models can provide further insights for climate researchers. A large class of trade models characterized in Arkolakis et al. (2012) derives the welfare implications of international trade from micro data on the share of expenditure on domestic goods and the consumers' trade elasticity with respect to trade costs. This approach is extended in Dingel et al. (2020) to consider the implications of spatial correlation in agriculture for global welfare inequality. Many of the determinants of economic activity are spatially correlated; if one country is productive at growing a crop, a neighbor with a similar climate is more likely to also be productive at growing the same crop. Dingel et al. empirically validate the prediction that countries gain more from trade in cereals when they have highly productive neighbors, as they are more likely to trade with their neighbors than with distant countries.[20] Overall, the welfare gains from trade appear unequal, being larger for countries that are more productive and smaller for countries that are less productive. This result, while illustrative for understanding current global trade dynamics, is also important for understanding global trade as a type of adaptation to climate change. If changes in cereal productivity due to climate change are spatially correlated, and the predictions of the above model are true, then there may be greater climate-driven welfare losses for low-productivity regions. According to the authors' projections,

---

[20] It may be worth clarifying that spatial correlation in agricultural productivity implies shared absolute advantage. Dingel et al. (2020) focus on spatial correlation in absolute advantage, which is distinct from spatial correlation in comparative advantage. The latter produces its own interesting implications for the gains from trade: countries that are nearby, but are productive at different things, may have a lot to gain from trade; countries that are far apart, or countries that are nearby and very similar, may have relatively less to gain from trade. The idea that countries with dissimilar technologies have more to gain from trade than those with similar technologies is part of Ricardo's original theory of trade; see Lind and Ramondo (2018) for details.

forecasted welfare losses in a model without spatial correlation in productivities understates the increases in welfare inequality.

In the main, spatial general equilibrium models are useful for predicting the economic dynamics associated with climate change and various macroeconomic types of adaptations.[21] Thus, there is room for a rich crossover between the literature of economic geography and climate change adaptation. And this literature is growing. Desmet and Rossi-Hansberg (2015) build a complex model that, in particular, highlights the role international migration in adapting to climate change. Balboni (2019) builds on the economic geography literature of Redding (2016) and the model from Caliendo et al. (2019) to consider whether infrastructure investments should continue to favor coastal areas. Given the availability of detailed agricultural data, applying these models specifically to agricultural climate change adaptation is particularly promising. Further, although the above summary focuses on trade models, other types of spatial general equilibrium models seem well poised for application to climate change adaptation, including optimal transport networks (Fajgelbaum and Schaal, 2020) and knowledge diffusion models (Buera and Oberfeld, 2020). Overall, broad classes of spatial general equilibrium models are well suited to tackling climate change adaptation questions.

5. Conclusions

In this report, I briefly reviewed the role of panel data analysis and spatial general equilibrium models for studying agricultural climate change adaptation. Recent innovations in the tools used to measure and predict climate change adaptation present promising avenues for future reasarch. Each of these tools on its own cannot perfectly identify climate change adaptation or predict future responses. However, when used as part of a broader toolkit, or perhaps even in conjunction with each another, they can allow us to better measure and predict how our food systems will respond to climate change.

[21] It is worth mentioning that the spatial units studied may not be countries; the unit of analysis may equally be counties within a country or even grid cells within a satellite image. However, all spatial models (including trade models) may be applied in alternative spatial contexts to understand spatial interactions. See, for example, the use of trade models in a regional setting to study intranational pricing (Atkin and Donaldson 2015). See also Costinot and Donaldson (2016); Donaldson and Hornbeck (2016); Donaldson (2018).

References

Acemoglu, D., Johnson, S. and Robinson, J. A. (2002). Reversal of fortune: geography and institutions in the making of the modern world income distribution. *Quarterly Journal of Economics*, *117*(4), 1231–94. https://doi.org/10.1162/003355302320935025

Adao, R., Costinot, A. and Donaldson, D. (2017). Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review*, *107*(3), 633–89. https://doi.org/10.1257/aer.20150956

Arkolakis, C., Costinot, A. and Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, *102*(1), 94–130. https://doi.org/10.1257/aer.102.1.94

Atkin, D. and Donaldson, D. (2015). *Who's getting globalized? The size and implications of intra-national trade costs* (NBER Working Paper No. 21439). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/w21439

Auffhammer, M. (2018). Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data (NBER Working Paper No. 24397). Cambridge, MA: National Bureau of Economic Research, Inc." https://www.nber.org/papers/w24397

Auffhammer, M. and Schlenker, W. (2014). Empirical studies on agricultural impacts and adaptation. *Energy Economics*, *46* (November), 555–61. https://doi.org/10.1016/j.eneco.2014.09.010

Auffhammer, M., Hsiang, S. M., Schlenker, W. and Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, *7* (2), 181–98. https://doi.org/10.1093/reep/ret016

Bakkensen, L. A., and Barrage, L. (2020). *Climate shocks, cyclones, and economic growth: Bridging the micro–macro gap* (NBER Working Paper No. 24893). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/24893

Bakkensen, L. A., and Barrage, L. (2021). *Flood risk belief heterogeneity and coastal home price dynamics: Going under water?* (NBER Working Paper No. 23854). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/23854.

Balboni, C. (2019). In harm's way? Infrastructure investment and the persistence of coastal cities. Revise and Resubmit: *American Economic Review*. https://economics.mit.edu/files/21167 .

Buera, F. J. and Oberfield, E. (2020). The global diffusion of ideas. *Econometrica*, *88*(1), 83–114. https://doi.org/10.3982/ECTA14044

Burke, M. and Emerick, K. (2016). Adaption to climate change: Evidence from U.S. agriculture. *American Economic Journal: Economic Policy*, *8*(3), 106–40. https://doi.org/10.1257/pol.20130025

Burke, M. and Lobell, D. (2010). Food security and adaptation to climate change: What do we know? In D. Lobell and M. Burke (Eds.), *Climate change and food security: Adapting agriculture to a warmer world* (pp. 133–53). Springer. https://doi.org/10.1007/978-90-481-2953-9_8

Caliendo, L., Dvorkin, M. and Parro, F. (2019). Trade and labor market dynamics: general equilibrium analysis of the China trade shock. *Econometrica*, *87*(3), 741–835. https://doi.org/10.3982/ECTA13758

Carleton, T. A. and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, *353*(6304), aad9837. https://doi.org/10.1126/science.aad9837.

Carleton, T., Jina, A., Delgado, M.. Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Kopp, R. E., McCusker, K. E., Nath, I. B., Rising, J., Rode, A., Seo, H. K., Viaene, A., Yuan, J. and Zhang, A. T. (2020). *Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits* (NBER Working Paper No. 227599). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/w27599.

Cline, W. (2007). *Global warming and agriculture: Impact estimates by country*. Peterson Institute for International Economics.

Costinot, A. and Donaldson, D. (2012). Ricardo's theory of comparative advantage: Old idea, new evidence. *American Economic Review: Papers & Proceedings*, *102* (3), 453–58. https://doi.org/10.1257/aer.102.3.453

Costinot, A. and Donaldson, D. (2016). *How large are the gains from economic integration? Theory and evidence from U.S. agriculture, 1880–1997* (NBER Working Paper No. 21439). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/w22946

Costinot, A., Donaldson, D. and Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, *124*(1), 205–48. https://doi.org/10.1086/684719

Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review*, *97*(1), 354–85. https://doi.org/10.1257/aer.97.1.354

Desmet, K. and Rossi-Hansberg, E. (2015). On the spatial impact of global warming. *Journal of Urban Economics*, *88* (July), 16–37. https://doi.org/10.1016/j.jue.2015.04.004

Dingel, J. I., Meng, K. C. and Hsiang, S. M. (2020). *Spatial correlation, trade, and inequality: Evidence from the global climate* (NBER Working Paper No. 25447). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/w25447

Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review*, *108*(4–5), 899–934. https://doi.org/10.1257/aer.20101199

Donaldson, D. and Hornbeck, R. (2016). Railroads and American economic growth: A "market access" approach. *Quarterly Journal of Economics*, *131*(2), 799–858. https://doi.org/10.1093/qje/qjw002

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, *70*(5), 1741–79. https://doi.org/10.1111/1468-0262.00352

Fajgelbaum, P. D. and Schaal, E. (2020). Optimal transport networks in spatial equilibrium. *Econometrica*, *88*(4), 1411–52. https://doi.org/10.3982/ECTA15213

Fankhauser, S. (2017). Adaptation to climate change. *Annual Review of Resource Economics*, *9*, 209–30. https://doi.org/10.1146/annurev-resource-100516-033554

Field, C. B., Barros, V. R., Dokken, D. J., Mach, K. J., Mastrandrea, M. D., Bilir, T. E., Chatterjee, M., Ebi, K. L., Estrada, Y. O., Genova, R. C., Gorma, B., Kissel, E. S., Levy, A. N., MacCracken, S., Mastrandrea, P. R. and White, L. L. (Eds.). (2014. *Climate change 2014: Impacts, adaptation, and vulnerability part A: Global and sectoral aspects: Working Group II contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC. https://www.ipcc.ch/report/ar5/wg2

Fried, S. (2021). *Seawalls and stilts: A quantitative macro study of climate adaptation* (Federal Reserve Bank of San Francisco Working Paper No. 2021-07). San Francisco, CA: Federal Reserve Bank of San Francisco. https://doi.org/10.24148/wp2021-07

Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helman and K. Rogoff (Eds.), *Handbook of international economics. Vol. 4*, (pp. 131–95). https://doi.org/10.1016/B978-0-444-54314-1.00003-3

Hsiang, S. (2016). Climate econometrics. *Annual Review of Resource Economics*, *8*, 43–75. https://doi.org/10.1146/annurev-resource-100815-095343

Lagakos, D. and Waugh, M. (2013). Selection, agriculture, and cross-country productivity differences. *American Economic Review*, *103* (2), 948–80. https://doi.org/10.1257/aer.103.2.948

Lind, N. and Ramondo, N. (2018). *Trade with correlation*. (NBER Working Paper No. 24380). Cambridge, MA: National Bureau of Economic Research, Inc. https://www.nber.org/papers/w24380.

Mendelsohn, R., Nordhaus, W. D. and Shaw, D. (1994). The impact of global warming on agriculture: A Ricardian analysis. *American Economic Review*, *84*(4), 753–71. https://www.jstor.org/stable/2118029

Moore, F. C. and Lobell, D. B. (2014). Adaptation potential of European agriculture in response to climate change. *Nature Climate Change*, *4*, 610–14. https://doi.org/10.1038/nclimate2228

Nath, I. (2020). The food problem and the aggregate productivity consequences of climate change. (NBER Working Paper No. 27297). Cambridge MA: National Bureau of Economic Research Inc. https://www.nber.org/papers/w27297.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of macroeconomics. Vol. 2* (pp. 71–162). https://doi.org/10.1016/bs.hesmac.2016.03.003

Redding, S. J. (2016). Goods trade, factor mobility and welfare. *Journal of International Economics*, *101* (July), 148–67. https://doi.org/10.1016/j.jinteco.2016.04.003

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, *106*(37), 15594–98. https://doi.org/10.1073/pnas.0906865106

Schneider, S. H., Easterling, W. E. and Mearns, L. O. (2000). Adaptation: Sensitivity to natural variability, agent assumptions and dynamic climate changes. *Climatic Change*, *45*(1), 203–21. https://doi.org/10.1023/A:1005657421149

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of macroeconomics. Vol. 2* (pp. 415–525). https://doi.org/10.1016/bs.hesmac.2016.04.002

Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M. M. B., Allen, S. K., Boschung, J., Nauels, A., Xia, Y., Bex, V. and Midgley, P. M. (Eds.). (2013). *Climate change 2013: The physical science basis: Working Group 1 contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC. https://www.ipcc.ch/report/ar5/wg1

Walker, B. and Salt, D. (2006). *Resilience thinking: Sustaining ecosystems and people in a changing world*. Island Press.

Bibliography

Aigner, Dennis J., and Glen G. Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *ILR Review* 30, no. 2 (January): 175–87. https://www.jstor.org/stable/2522871.

Alon, Titan, and Daniel Fershtman. 2019. "A Theory of Human Capital Specialization." https://drive.google.com/file/d/1LXjGzQSAxl_ijlPMDE11dTPzvRLkV92q/view?usp=sharing.

Altonji, Joseph G. 1993. "The Demand for and Return to Education When Education Outcomes are Uncertain." *Journal of Labor Economics* 11, no. 1 (January): 48–83. https://doi.org/10.1086/298317.

Altonji, Joseph G., Peter Arcidiacono, and Arnaud Maurel. 2016. "The Analysis of Field Choice in College and Graduate School: Determinants and Wage Effects." Chap. 7 in *Handbook of the Economics of Education,* edited by Eric A. Hanushek, Stephen Machin, and Ludger Woessmann, 5:305–96. https://doi.org/10.1016/B978-0-444-63459-7.00007-5.

Altonji, Joseph G., Erica Blom, and Costas Meghir. 2012. "Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers." *Annual Review of Economics* 4, no. 1 (July): 185–223. https://doi.org/10.1146/annurev-economics-080511-110908.

Archer, Kellie J., and Stanley Lemeshow. 2006. "Goodness-of-fit Test for a Logistic Regression Model Fitted using Survey Sample Data." *The Stata Journal* 6 (1): 97–105. https://doi.org/10.1177/1536867X0600600106.

Arcidiacono, Peter. 2004. "Ability sorting and the returns to college major." *Journal of Econometrics* 121 (2): 343–75. https://doi.org/10.1016/j.jeconom.2003.10.010.

Arcidiacono, Peter, Esteban Aucejo, Arnuad Maurel, and Tyler Ransom. 2016. "College Attrition and the Dynamics of Information Revelation." NBER Working Paper no. 22325, https://doi.org/10.3386/w22325.

Arcidiacono, Peter, V. Joseph Hotz, and Songman Kang. 2012. "Modeling college major choices using elicited measures of expectations and counterfactuals." *Journal of Econometrics* 166, no. 1 (January): 3–16. https://doi.org/10.1016/j.jeconom.2011.06.002.

Arcidiacono, Peter, V. Joseph Hotz, Arnaud Maurel, and Teresa Romano. 2020. "Ex-ante Returns and Occupational Choice." *Journal of Political Economy* 128, no. 12 (December): 4475–4522. https://doi.org/10.1086/710559.

Arrow, Kenneth J. 1972. "Some Mathematical Models of Race in the Labor Market." In *Racial Discrimination in Economic Life,* Anthony H. Pascal, 187–203. Lexington, MA: D.C. Health.

Bayer, Amanda, Syon P. Bhanot, and Fernando Lozano. 2019. "Does Simple Information Provision Lead to More Diverse Classrooms? Evidence from a Field Experiment on Undergraduate Economics." *AEA Papers and Proceedings* 109 (May): 110–14. https://doi.org/10.1257/pandp.20191097.

Becker, Gary S. 1957. *The Economics of Discrimination.* 2nd. Chicago: University of Chicago Press.

———. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy* 7, no. 5 (October): 9–49. https://doi.org/10.1086/258724.

Becker, Gary S., William H.J. Hubbard, and Kevin M. Murphy. 2010. "Explaining the Worldwide Boom in Higher Education of Women." *Journal of Human Capital* 4 (3): 203–41. https://doi.org/10.1086/657914.

Ben-Porath, Yoram. 1967. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy* 75, no. 4 (August): 352–65. https://doi.org/10.1086/259291.

Billingsley, Patrick. 2012. *Probability and Measure.* Hoboken, NJ: A John Wiley & Sons, Inc.

Black, Dan A., Amelia M. Haviland, Seth G. Sanders, and Lowell J. Taylor. 2008. "Gender Wage Disparities among the Highly Educated." *Journal of Human Resources* 43 (3): 630–59. https://doi.org/10.3368/jhr.43.3.630.

Blau, Francine D., and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55 (3): 789–865. https://doi.org/10.1257/jel.20160995.

Bleemer, Zachary, and Aashish Mehta. 2021. "College Major Restrictions and Student Stratification." December. https://zacharybleemer.com/wp-content/uploads/Working-Papers/Restrictions_Paper.pdf.

Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope. 2019. "Inaccurate Statistical Discrimination: an identification problem." NBER Working Paper No. 25935, https://doi.org/10.3386/w25935.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "The Dynamics of Discrimination: Theory and Evidence." *American Economic Review* 109 (10): 3395–436. https://doi.org/10.1257/aer.20171829.

Bryan, Michael, Samuel Austin Caperton, Darryl Cooney, and Kristin Dudley. 2020. *2012 Beginning Postsecondary Students Longitudinal Study (BPS:12) Postsecondary Education Transcript Study (PETS): Data File Documentation.* NCES 2021176. NCES, December. https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2021176.

Bryan, Michael, Darryl Cooney, and Barbara Elliott. 2019. *Beginning Postsecondary Students Study 12/17 (BPS:12/17): Data File Documentation.* NCES 2020522. NCES, November. https://nces.ed.gov/pubsearch/pubsinfo.asp?pubid=2020522.

Casella, George, and Roger L. Berger. 2002. *Statistical Inference.* Second Edition. Belmont, CA: Brooks/Cole Cengage Learning.

Charles, Kerwin Kofi, and Jonathan Guryan. 2008. "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination." *Journal of Political Economy* 116, no. 5 (October): 773–809. https://doi.org/10.1086/593073.

Coate, Stephen, and Glenn C. Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *The American Economic Review* 83, no. 5 (December): 1220–40. https://www.jstor.org/stable/2117558.

De Giorgi, Giacomo, Michelle Pellizzari, and Silvia Redaelli. 2010. "Identification of Social Interactions through Partially Overlapping Peer Groups." *American Economic Journal: Applied economics* 2, no. 2 (April): 241–275. https://doi.org/10.1257/app.2.2.241.

Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104, no. 4 (April): 1091–119. https://doi.org/10.1257/aer.104.4.1091.

Hosmer, David W., and Stanley Lemeshow. 1980. "Goodness of fit tests for the multiple logistic regression model." *Communications in Statistics - Theory and Methods* 9 (10): 1043–1069. https://doi.org/10.1080/03610928008827941.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87, no. 5 (September): 1439–74. https://doi.org/10.3982/ECTA11427.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11, no. 4 (October): 181–209. https://doi.org/10.1257/app.20180010.

Lang, Kevin, and Jee-Yeon K. Lehmann. 2012. “Racial Discrimination in the Labor Market: Theory and Empirics.” *Journal of Economic Literature* 50 (4): 959–1006. https://doi.org/10.1257/jel.50.4.959.

Li, Hsueh-Hsiang. 2018. “Do mentoring, information, and nudge reduce the gender gap in economics majors?” *Economics of Education Review* 64 (June): 165–183. https://doi.org/10.1016/j.econedurev.2018.04.004.

Lim, Jaegeum, and Jonathan Meer. 2020. “Persistent Effects of Teacher - Student Gender Matches.” *Journal of Human Resources* 55 (3): 809–35. https://doi.org/10.3368/jhr.55.3.0218-9314R4.

Lundberg, Shelly J., and Richard Startz. 1983. “Private Discrimination and Social Intervention in Competitive Labor Markets.” *American Economic Review* 73, no. 3 (June): 340–7. https://www.jstor.org/stable/1808117.

Mincer, Jacob. 1974. *Schooling, Experience, and Earnings.* New York, NY: National Bureau of Economic Research. ISBN: 0-87014-265-8. http://www.nber.org/books/minc74-1.

Neilson, Christopher, Adam Altmejd, Andrés Barrios-Fernández, Marin Drlje, and Dejan Kovac. 2020. “Sibling’s Effects on College and Major Choices: Evidence from Chile, Croatia and Sweden.” https://dataspace.princeton.edu/jspui/bitstream/88435/dsp01z890rx142/3/633a.pdf.

Ost, Ben. 2010. “The role of peers and grades in determining major persistence in the sciences.” *Economics of Education Review* 29, no. 6 (December): 923–34. https://doi.org/10.1016/j.econedurev.2010.06.011.

Owen, Stephanie. 2020. “College Field Specialization and Beliefs about Relative Performance.” https://stephanieowen.github.io/OwenStephanie_JMP.pdf.

Patnaik, Arpita, Matthew Wiswall, and Basit Zafar. 2020. “College Majors.” NBER Working Paper No. 27645. https://doi.org/10.3386/w27645.

Phelps, Edmund S. 1972. “The Statistical Theory of Racism and Sexism.” *American Economic Review* 62, no. 4 (September): 659–61. https://www.jstor.org/stable/1806107.

Porter, Catherine, and Danila Serra. 2020. “Gender Differences and the Choice of Major: The Importance of Female Role Models.” *American Economic Journal - Applied Economics* 12 (3): 226–254. https://doi.org/https://doi.org/10.1257/app.20180426.

Rask, Kevin, and Jill Tiefenthaler. 2008. “The role of grade sensitivity in explaining the gender imbalance in undergraduate economics.” *Economics of Education Review* 27, no. 6 (December): 676–87. https://doi.org/10.1016/j.econedurev.2007.09.010.

Riley, Emma. 2019. “Role models in movies: the impact of Queen Katwe on students’ educational attainment.” https://emmaalriley.files.wordpress.com/2020/02/role_models_in_movies.pdf.

Rosen, Sherwin. 1983. “Specialization and Human Capital.” *Journal of Labor Economics* 1, no. 1 (January): 43–9. https://doi.org/10.1086/298003.

Roy, A. D. 1951. “Some Thoughts on the Distribution of Earnings.” *Oxford Economic Papers* 3 (2): 135–146. http://www.jstor.org/stable/2662082.

Sloane, Carolyn, Erik Hurst, and Dan Black. 2020. “A Cross-Cohort Analysis of Human Capital Specialization and the College Gender Wage Gap.” NBER Working Paper No. 26348, https://doi.org/10.3386/w26348.

Stinebrickner, Ralph, and Todd R. Stinebrickner. 2014a. “A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout.” *Review of Economic Studies* 81:427–472. https://doi.org/10.1093/restud/rdt025.

———. 2014b. “Academic Performance and College Dropout: Using Longitudinal Expectations Data to Estimate a Learning Model.” *Journal of Labor Economics* 32 (3). https://doi.org/10.1086/675308.

“The CSU Fact Book 2021.” 2021. https://www.calstate.edu/csu-system/about-the-csu/facts-about-the-csu/Documents/facts2021.pdf.

Wiswall, Matthew, and Basit Zafar. 2015. “Determinants of College Major Choice: Identification Using an Information Experiment.” *The Review of Economic Studies* 82, no. 2 (April): 791–824. https://doi.org/10.1093/restud/rdu044.

———. 2018. “Preference for the Workplace, Investment in Human Capital, and Gender.” *Quarterly Journal of Economics* 133, no. 1 (February): 457–507. https://doi.org/10.1093/qje/qjx035.

Wooldridge, Jeffrey M. 2010. *Econometric Analysis of Cross Section and Panel Data.* Second Edition. Cambridge, Massachusets: The MIT Press.

Zafar, Basit. 2013. “College Major Choice and the Gender Gap.” *Journal of Human Resources* 48 (3): 545–95. https://doi.org/10.3368/jhr.48.3.545.