

ESSAYS IN APPLIED MICROECONOMICS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Todd Randy Jones

May 2018

© 2018 Todd Randy Jones

ALL RIGHTS RESERVED

ESSAYS IN APPLIED MICROECONOMICS

Todd Randy Jones, Ph.D.

Cornell University 2018

This dissertation includes three essays. The first essay is entitled “Do Peers Influence Occupational Preferences? Evidence from Randomly-Assigned Peer Groups at West Point” and is coauthored with Michael Kofoed. We study the effect of an individual’s peers on his occupational preferences among cadets at West Point. We exploit unique institutional details at West Point, including random assignment of cadets to peer groups and the repeated elicitation of occupational preferences to overcome common biases that arise in the study of peer effects. Overall, we find little evidence of peer effects in this setting.

The second essay, coauthored with Ronald Ehrenberg, is entitled “Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?” This paper uses restricted-access *Survey of Earned Doctorates* merged with other datasets to study which attributes of a doctoral university are associated with a higher share of its undergraduate BAs who proceed to earn a PhD. We consider four fields: humanities, physical sciences, life sciences, and social sciences. We use truncation correction methodology to correct for PhDs earned after the data end. In our main specification, we find that across all fields PhD production is positively related to student test scores and the number of high-quality PhD programs an institution has. It is negatively related to the total number of students and the share of total BAs that are received in the field.

The third essay is entitled “Information and the Beauty Premium in Political Elections.” It is coauthored with Joseph Price. We study the beauty premium in U.S. political

elections. We collect subjective beauty data on candidate photos from participants on an internet survey. We find that more beautiful candidates receive a higher vote share and are more likely to win, but that this is moderated by situations in which the voter is more likely to have less information on the candidate. The beauty premium is smaller—or zero—for U.S. congressional races than for state congressional races. It is also smaller for incumbents and in elections in which there is more campaign spending.

BIOGRAPHICAL SKETCH

Todd R. Jones was born in Bedford, Texas. He grew up in Brigham City, Utah, also spending time in Stuttgart, Germany. After serving a religious mission in Sweden, he moved to Shanghai, China, to live with his family and study Chinese. Before beginning his Ph.D. studies at Cornell, Todd attended Utah State University, earning a Bachelor's in Economics and Master's in Statistics. An introductory economics course taught by Dwight Israelson at Utah State opened his eyes to the world of economics. Although difficult at times, Todd has loved his experience at Cornell and in Ithaca and will be very sad to leave. On the other hand, he is very excited about his next opportunity, a postdoc at Georgia State University working on economics of education research.

To my parents

ACKNOWLEDGEMENTS

I am very grateful for my advisor, Ron Ehrenberg, for believing in me and giving me many opportunities to grow. I have learned a great deal through our many meetings as well as being a teaching assistant multiple times for his economics of the university course. It is rare to find someone in academia who cares so incredibly much about the success of others. I would also like to thank the other two members of my committee, Mike Lovenheim and David Just. I appreciate their advice, support, kindness, and encouragement over the last several years.

It has been a pleasure working with and learning from my coauthors, Mike Kofoed, Joe Price, and Nick Sanders. I am thankful that the economics department admitted me and gave me such a great education. I will always appreciate my wonderful—although challenging—time at Cornell and how much I have grown though this experience. My life has been enriched through my association with many graduate students at Cornell, including Jason Cook, Rich Patterson, Mindi Martin, Kristy Parkinson, Sarah Prenovitz, Cassie Benson, Jeff Swigert, Vesall Nourani, Corbin Miller, Adam Shumway, and Sam Dodini. Many other students and professors have also helped me during my time at Cornell.

I have greatly enjoyed my time in my church in Ithaca, through which I have met many outstanding people and grown a great deal. I appreciate that Cornell allows religious organizations such as the LDSSA to utilize Anabel Taylor Hall. I am forever grateful to Todd Kreuger for his help and advice. I am thankful to have had an exceptional roommate and friend, Joseph Ruesch. Finally, I have been extremely blessed to be part of a truly remarkable family. Thanks for everything, Grandma Israelsen, Grandpa Jones, Katie, Nate, Ben, Lucy, Anna, Matt, Emily, Shelley, and, especially, Mom, and Dad.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vi
1 Do Peers Influence Occupational Preferences? Evidence from Randomly-Assigned Peer Groups at West Point	1
1.1 Introduction	1
1.2 Prior Literature	8
1.3 West Point Institutional Details	12
1.4 Data and Descriptive Statistics	18
1.5 Empirical Strategy	23
1.6 Check for Randomization	24
1.7 Regression Analysis	26
1.7.1 Branch-Specific Regressions	26
1.7.2 Stacked Regression	28
1.7.3 Older Peers	29
1.7.4 Roommates	30
1.7.5 Combined Peer Groups	31
1.7.6 Branch-Specific Regressions Results	31
1.7.7 Stacked Regression Results	32
1.7.8 Older Peers Results	33
1.7.9 Roommates Results	34
1.7.10 Combined Peer Groups Results	34
1.8 Rank Agreement	34
1.8.1 Rank Agreement Calculation	35
1.8.2 Using Rank Agreement to Test for Peer Effects	39
1.8.3 Rank Agreement Results	41
1.9 Conclusion	42
2 Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?	44
2.1 Introduction	44
2.2 PhD Production	47
2.3 Estimating Equation	50
2.4 Constructing the PhD Production Rate and Correcting for Truncation	52
2.4.1 Maximum TTD	52

2.4.2	Truncation Correction	54
2.4.3	PhD Production Rate	56
2.5	Descriptive Statistics	57
2.6	Results	60
2.7	Conclusion	64
3	Information and the Beauty Premium in Political Elections	65
4	Tables	73
4.1	Tables for Do Peers Influence Occupational Preferences?	73
4.2	Tables for Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?	86
5	Figures	92
5.1	Figures for Do Peers Influence Occupational Preferences?	92
5.2	Figures for Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?	98
6	Appendices for Do Peers Influence Occupational Preferences?	101
6.1	Rank Agreement Example	101
6.2	Tables	105
6.3	Figures	108
7	Appendices for Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?	109
7.1	Data	109
7.2	Classification of Subjects to Fields	115
8	Bibliography	119

CHAPTER 1

DO PEERS INFLUENCE OCCUPATIONAL PREFERENCES? EVIDENCE FROM RANDOMLY-ASSIGNED PEER GROUPS AT WEST POINT

1.1 Introduction

Occupation selection relates strongly to many life outcomes and labor market experiences. For example, large variation exists in wages across occupations; industry and occupation explained 31% of wage variations among full-time workers in the 2016 Current Population Survey.¹ Correlations have been found between occupation/industry and job satisfaction (Rose, 2003) and job tenure (Zavodny, 2003), and slightly more than half (51%) of the gender wage gap is explained by occupation and industry (Blau and Kahn, 2016). In addition to earnings, there is causal evidence that occupation influences health outcomes, such as obesity, smoking, and alcohol consumption (Kelly et al., 2014).

The authors would like to thank the Office of Economic and Manpower Analysis for providing the data and for helpful feedback. We would additionally like to thank Ronald Ehrenberg, Michael Lovenheim, David Just, Stephanie Cellini, Jeff Denning, Doug Miller, Ben Hansen, Andrew Johnston, Nicholas Jolly, Joshua Kinsler, Doug Miller, Richard Patterson, Joe Price, Evan Riehl, Barton Willage, and participants at several seminars. The views expressed herein are those of the authors and do not reflect the position of the United States Military Academy, the Department of the Army, or the Department of Defense.

¹From the authors' calculations. Full-time is defined as having more than 35 usual hours. Industry and occupation explained 34% of the variation in weekly earnings among full-time workers. Not all of this variation was caused by occupation choice, and some literature explains this variation. For example, job function (Gittleman and Pierce, 2011), individual heterogeneity (Abowd, Kramarz, and Margolis, 1999), and worker ability (Goux and Maurin, 1999) offer great explanatory power. However, Du Caju et al. (2010) do not find that worker ability explains a large share of the difference. Measuring earnings by total compensation—as opposed to wages—substantially increases the amount of interindustry difference (Gittleman and Pierce, 2012).

Despite the importance of occupations, there is little evidence on how individuals' preferences for occupations are formed, what factors influence them, and how they evolve. There are several factors through which preferences for occupations might be influenced, including parents' occupations (Lindquist, Sol, and Van Praag, 2015), mentors (Kofoed and McGovney, 2017), and the earning potential of a job.² Another factor, and the focus of this paper, is the role that peers play in an individual's occupational preferences. College is a time such preferences are particularly important, when college students need to make decisions on occupations, with long-term consequences. Exposure to certain types of peers might influence the career trajectories of college students. Colleges, for example, use peer quality as a recruitment tool to attract the brightest students. Many students and families respond to these types of arguments with the hopes that successful peers or alumni will provide positive academic benefits. Sacerdote (2001), Zimmerman (2003), and Carrell, Fullerton, and West (2009) use random assignment to peer groups, finding evidence of peer effects in academic performance. Since many other studies also find evidence of peer effects across domains, such as worker productivity (Mas and Moretti, 2009; Herbst and Mas, 2015), peers might also influence occupational preferences. In a theoretical discussion of determinants of occupation choice in the military, Lim et al. (2009) discuss social influence as a factor.

Few studies examine the causal effect peers have on occupational preferences. One reason is lack of availability of repeated occupational preference data, and another is identification challenges present in peer effects studies—selection, reflection, and common

²Additional factors such as temporal flexibility (Goldin, 2014), credit constraints (Rothstein and Rouse, 2011), and labor supply effects might influence occupational choice, without influencing occupational preferences.

shock problems (Manski, 1993). This paper uses individual-level, administrative data, including a panel of rank-ordered preference data of students (cadets) at the United States Military Academy (West Point), combined with unique institutional features, to estimate the effect of peers on occupational preferences. We overcome self-selection into peer groups by using random assignment to peers, and reflection by using preferences collected before exposure to peer groups.

West Point conditionally and randomly assigns each cadet to a set of peers called a company. The company contains freshmen through seniors, and serves as the primary social unit for a cadet during time at West Point. Most social interactions are organized around the company, including dormitories, seating at meals, and extracurricular teams (though not classes). West Point assigns entering cadets to companies after they complete basic training during the summer but before their freshman year, and again randomly assigns cadets to a company before the beginning of their sophomore year. Cadets stay in the second company for the remainder of their time at West Point. Thus, all cadets are exposed to two sets of company-level peers. Cadets are also randomly assigned to one (or occasionally two) roommate(s) during the first year.

Cadets are required to serve in the U.S. Army as officers after graduation in exchange for free tuition, room, and board while attending West Point. At the beginning of their senior years, cadets rank each of the 16 post-graduation officer occupational branches that are available to them, and the Army assigns them to be an officer in one of the branches, given their preferences, class rank, and the capacity constraints of the branches.³

³Branches are occupations in the Army, and are not to be confused with branches in the U.S. Military, such as the Navy and Air Force.

West Point polls cadets at least once per school year on how they rank each occupational branch if they had to make a decision at that moment, allowing us to assess how their preferences evolved. The ability to observe the evolution of preferences is a novel feature of this study. The closest analogues are found in college major choice literature, which uses elicitation of preferences concerning college majors (Stinebrickner and Stinebrickner, 2013; Wiswall and Zafar, 2014). We present new descriptive evidence on how occupational preferences change over time, including flows from each branch to each other branch. When we examine the top-ranked branch—a large percentage of cadets are ultimately matched to their top-ranked, final preference—there is a large amount of switching across rounds. The initially most popular branches become less popular, and the initially least popular branches become more popular. Cadets gain information about the various branches and about themselves, allowing preferences to evolve over the branches as they gain collegiate experience. However, the prevalence of switching decreases with time in college, indicating that preferences stabilize after the first several rounds. Some branches, such as Aviation and Infantry, are vastly more popular than other branches, such as Chemical Corps.

We estimate peer effects models that examine the influence that peers in the same company and cohort have on occupational preferences of a cadet's first and second years at West Point. The identifying variation we use is differences in the percentage of cadets across companies who prefer different branches. We use the timing of when cadets are randomly assigned into companies relative to when they produce their branch rankings to address identification concerns. Random assignment to companies and roommates allows us to overcome the selection problem endemic to peer-effects analyses, which arises

when unobservable characteristics drive selection into both peer groups and outcomes. The reflection problem occurs when the individual and group co-determine each others' outcomes. This is commonly addressed by using pre- or exogenously determined attributes of a group as the explanatory variable, in comparison to using the final outcome as the explanatory variable. In our main specification, we use preferences over branches that were given before exposure to the company and roommate.⁴ We do not focus on common shocks in this paper.

In addition to own-cohort peer effects, we explore mentor-peer effects, where mentors are defined as older cadets in the same company. Finally, we introduce a novel statistic called rank agreement, which is a summary measure of the proximity of a cadet's preferences to other cadets in the company. This enables us to reduce the multi-dimensional occupation measurement into a single measure, which is useful when studying occupational selection more broadly. We use rank agreement to test whether overall preference proximity decreases after cadets are exposed to their companies over the course of an academic year by comparing observed rank agreement with a distribution of rank agreement from synthetic companies to take into account aggregate West Point preferences.

Although regression specifications without controls for initial popularity of occupational branches show evidence of a peer effect, once popularity is controlled for, the effect disappears. This pattern holds at both the company and roommate levels, and among peers of the same cohort and mentors.⁵ Results demonstrate that controlling for existing

⁴Several studies, including those that feature a military academy, also use random assignment to a peer group to deal with selection bias, and pre-determined characteristics to deal with the reflection problem. See Sacerdote (2001), Lyle (2007), Carrell, Fullerton, and West (2009), and Carrell, Hoekstra, and West (2016).

⁵The one exception is that we find small, negative effects in latter rounds.

popularity of choices in similar discrete-choice settings is important; even with random assignment, results can be misleading without controlling for baseline preferences. Our main rank agreement results indicate no peer effects. However, without accounting for the reflection problem, we do find some evidence of peer effects during the first few weeks of the semester of the first year; this result fades quickly.

West Point is an informative setting in which to study the research question for several reasons. First, it is rare to observe an individual's preferences over occupations, and to have a repeated panel over preferences. These data allow us to view the evolution of preferences, which is a novel feature of the study. Second, the data allow us to address common identification concerns. This paper uses random assignment to estimate a causal effect of peers on occupational preferences or choice, requiring less-stringent assumptions than would exist in a different context by using randomization to peer groups and preferences measured before peer group assignment to overcome selection and reflection biases. Third, the entire choice set of available occupations is known, which is unusual in a typical labor market context, since individuals generally have thousands of jobs available to them all over the country and world, and a researcher would not likely observe all of them. Fourth, the context we examine is unique in that cadets are in a much more similar situation as fellow students in their cohorts are than are typical students at other colleges. Tuition is free, and everyone gets an equal stipend. In addition, most officer fields have the same starting salaries. Thus, preferences over attributes of the occupations are likely to be more important than financial considerations. Finally, since cadets are assigned to both companies and roommates, we are able to compare peer groups that are more broadly and narrowly defined in terms of peer exposure. We also examine own-cohort peer effects and

older mentor peer effects.

West Point is an instructive institution to study in its own right.⁶ It is an important training ground for those who will become commissioned officers in the U.S. Army. The Army is a multi-billion dollar organization, and approximately 22% of all Army officers graduate from West Point. Many of these officers will lead hundreds of soldiers into combat, and are responsible for carrying out national defense policies for the United States.

Despite its strengths, the sample is also limited for several reasons. West Point is an academy with very high admissions standards, with an acceptance rate comparable to Ivy League Institutions.⁷ Thus, its students are of high ability in comparison to the average individual. Although an institution of higher education, West Point is also a military academy, leading students at West Point to have a unique experience in relation to students at non-military colleges and universities. Occupations into which cadets enter after graduation are also limited, and are specific to the U.S. Army. The sample consists of only male cadets for reasons discussed in Section 7.1.

This study explores a potential area of occupational preferences, which is an important determinant of many outcomes in an individual's life. Despite having reason to believe that peers are an important part of occupational choice decisions, overall we find little evidence of peer effects in this context. This is not due to lack of variation in branch preferences

⁶West Point has a rich history, having produced two U.S. presidents, Ulysses S. Grant and Dwight D. Eisenhower, and many other notable alumni, including Buzz Aldrin, George Custer, Michael Krzyzewski, Douglas MacArthur, George Patton, and Norman Schwarzkopf. Before the First World War and the rise of the ROTC program, it was the primary source of commissioned officers.

⁷The Fall 2015 acceptance rate at West Point was 10%. Yale, Princeton, Brown, Pennsylvania, Dartmouth, and Cornell had acceptance rates of 7%, 7%, 9%, 10%, 11%, and 15%, respectively. See U.S. News & World Report (n.d.).

across companies or in preference changes over time. Lack of evidence of peer effects accords with the other main study on the subject, Arcidiacono and Nicholson (2005).

The remainder of this paper is organized as follows. First, we review the literature in Section 1.2. In Section 1.3, we describe the institutional details of West Point. Next, we discuss the data and show descriptive statistics in Section 7.1. The empirical strategy, check for randomization, regression analysis, and rank agreement are found in Sections 1.5-1.8 . We conclude in Section 1.9.

1.2 Prior Literature

This paper contributes to two topics in labor economics literature—peer effects and occupational choice. Consensus suggests that social groups affect many aspects of life. Empirical study of peer effects, which is extensive, quantifies influences that peer groups have on members.⁸ One challenge to obtaining causal estimates of peer effects is selection bias. In many contexts, individuals choose their peer groups; college students choose their roommates and classrooms, people choose their friends, and employees choose which company to work for. This endogenous choice generally introduces selection bias, making it difficult to disentangle the role that individual attributes play in their outcomes, and the role peer groups play due to individual, unobserved factors that influence which peer groups

⁸There are many peer effects studies in higher education. For example, Carrell et al. (2008) study peer effects on academic cheating. Several studies focus on health outcomes, such as obesity (Carrell, Hoekstra, and West, 2011) and risky health behaviors, including smoking and drinking (Gaviria, and Raphael, 2001; Eisenberg, Golberstein, and Whitlock, 2014). Sacerdote (2011) reviews peer-effect research in education, including both K–12 and higher education contexts.

individuals choose to be in that are also correlated with the outcome. An increasing number of studies in higher education use random assignment of individuals to peer groups to circumvent selection issue. For example, Sacerdote (2001) examines the effect that a randomly assigned roommate at Dartmouth has on the other roommate's fraternity choice and academic outcomes. Following a similar identification strategy, Zimmerman (2003) assesses peer effects on the college grades of students using the academic ability of randomly assigned Williams College roommates, as proxied by math and verbal SAT scores.

Peer-effects literature often studies military academies due to their frequent use of randomization of cadets to social groups. Lyle (2007) uses random assignment of cadets to companies at West Point to study educational and labor market outcomes such as GPA, college major, and the eventual decision of whether to remain in the Army post service commitment. Similarly, Carrell, Fullerton, and West (2009) use random assignment of cadets to peer groups (squadrons) at the United States Air Force Academy (USAFA) to study educational outcomes, finding evidence of an academic ability peer effect on student grades, though Lyle (2007) does not. Carrell, Fullerton, and West (2009) argue that the conditional randomization employed at West Point, which includes conditioning on academic ability, leads to lower variation in the companies at West Point than in squadrons at USAFA, which does not condition on ability. This observation might reconcile disparate findings across the two contexts. Using peer effects estimates from past data, Carrell, Sacerdote, and West (2013) assign cadets to peer groups in an attempt to improve outcomes, yet they find negative peer effects. They argue that this is due to endogenous peer group formation, which is important to keep in mind when bringing observational peer effects results to policy. Carrell, Hoekstra, and West (2016) find evidence that being randomly

assigned to a black roommate as a freshman increases the probability that a student at US-AFA will choose to become roommates with a black roommate during subsequent years. We take a similar approach in our paper, and use the random assignment of West Point cadets to companies and roommates to address selection.

The reflection problem arises when the outcomes of a peer group are used to estimate the peer effect, which is problematic because those outcomes might also have been influenced by other peers in the group. The ideal way to solve this problem is to use only peer characteristics that are fixed attributes, or measured before exposure to a peer group, such as gender and prior test scores. Carrell, Fullerton, and West (2009) use the SAT scores of members of a cadet's squadron, which are determined before entry into college, to proxy a squadron's academic ability. We use preference data measured before assignment to a company to overcome the reflection problem.

This study is among the few that explore the relationship between peer effects and occupational choice. Although much of the foundational literature on occupational choice⁹ does not model social influence as an input, there is reason to believe that this might be an important factor. In a theoretical paper, Mani and Mullin (2004) include a community's perceptions of occupations and whether they relate to an individual's social approval when modeling occupational choice. Focusing on the gender composition of classmates, Zöllitz and Feld (2017) show that the majors and occupations that women choose are affected by the proportion of randomly assigned classmates who are female. Using administrative Norwegian data, Markussen and Røed (2017) find that women are more likely to

⁹See Roy (1951), Siow (1984), and Keane and Wolpin (1997).

be entrepreneurs when exposed to other female entrepreneurs. In a report that discusses branch choice in the Army, with a focus on ROTC, Lim et al. (2009) discuss theoretical determinants of occupational choice, including social approval and network effects. Thus, there is an *a priori* reason to consider that peers might be a component of preferences of occupation.

Marmaros and Sacerdote (2002) use student surveys at Dartmouth to find that social groups, including fraternities, play a role in student networking. They also find suggestive evidence that fraternities and roommates influence job choice, but due to selection bias, do not claim causal interpretation. The authors use random assignment of roommates and hallmates to address the selection issue, finding that hallmate choice of high-paying occupations correlates with own choice. However, they are cautious in their interpretation due to the reflection problem present when regressing outcomes on outcomes. We address this problem by using branch choice preferences made before the formation of the peer group.

Kofoed and McGovney (2017) also study the branch choices of West Point cadets. They use an empirical strategy similar to Carrell, Page, and West (2010) to consider the effect that the gender of a cadet's randomly assigned TAC officer—a mentor who serves in an advisory role—plays in which branch a cadet chooses. They find that a female cadet is more likely to choose the branch of her TAC officer if the officer is female. The same result is found for black (but not Hispanic) cadets and TAC officers. Our paper differs in several respects, including the type of social group considered and a focus on the evolution of preferences.

The paper most similar to ours is Arcidiacono and Nicholson (2005), which empirically combines peer effects and occupational choice. The authors examine medical schools using the universe of medical school graduates and focusing on the effect of peers on the choice of medical school specialty and academic performance. The medical specialties are assigned to low-paying and high-paying groups. Since medical students are not assigned randomly to medical schools, they address the selection problem by using medical school fixed effects, which means that they control for the baseline propensity of students at a medical school to prefer one type of medical specialty over another. They find no evidence of peer effects when the peer group is defined as the entire graduating cohort, or regarding race or ability. However, they find evidence of a small peer effect among women in terms of academic performance, but not specialty choice. We do not collapse occupational choices into two categories, but use each branch individually. We are also able to make less restrictive assumptions due to cadets being assigned randomly to peer groups at West Point, and we are able to examine the evolution of preferences as opposed to only the final choice.

1.3 West Point Institutional Details

The United States Military Academy, located above the Hudson River in West Point, New York, is a four-year liberal arts college that combines academic curricula with military training. The cadets educated there receive free tuition, room, board, and a stipend. After graduation, each cadet receives a commission as an officer in the U.S. Army, and serves

for a minimum of five years on active duty in an occupational branch.¹⁰ Branches include specialties such as Infantry, Engineer, Military Police, Chemical Corps, and Military Intelligence. A full list of branches available to the cadets appears in Table 4.1.1.¹¹

Cadets at West Point are organized into one of 36 companies. Each company consists of freshmen, sophomores, juniors, and seniors. West Point also randomly assigns an Army officer to each company, whose primary job is to mentor cadets in the company. Cadets in the same cohort year spend much of their time together, especially freshmen. Companies eat, room, and exercise together. However, they do not attend academic classes as companies; beginning cadets are randomly assigned to courses, and largely take the same subjects. Although freshmen are allowed to visit hallways of other companies, socializing outside of the company is highly discouraged, especially with upperclassmen.

In a process called the “scramble,” West Point conditionally randomly assigns cadets to a company both before and after their freshman year. Thus, the typical cadet is in a first company for one year and in a final company for the remaining three years. A timeline illustrating the timing of the first and second company assignments is shown in Figure 5.1.1. As described in Lyle (2007), each cadet is randomly assigned to the initial company by a computer algorithm. The algorithm then shuffles cadets to different companies, attempting to equalize company means across eight dimensions, including gender, recruited athlete, attendance at the West Point preparatory school, College Entrance Exam Rank, and Whole

¹⁰ After active duty concludes, they also serve three years in reserve duty. Those in the Aviation branch are required to serve seven years on active duty.

¹¹ A 17th branch, Cyber, was introduced during the sample’s studies. For consistency, we refer to 16 branches throughout the study. Cyber was unavailable as a choice for the 2015 cohort, but the 2016 cohort was able to choose Cyber beginning in the third round.

Candidate Score (a composite of various high school academic and extracurricular measures).¹²

Preferences for branches are elicited from the cadets, who rank the branches from 1 to 16. Six rounds of preferences were collected. Rounds 1–3 were taken during the first half of September during the freshman, sophomore, and junior years. Round 4 was gathered in April of the junior year. Round 5 occurred in August, and Round 6 occurred at the start of September, both during the senior year. The timing of these preferences is shown in Figure 5.1.1. Round 6 represents the final branch preference elicitation, the set of preferences that factors into assignment of the branch in which the cadet will serve after graduation. All other preference rounds are non-binding, and thus are of much lower stakes.

Cadets are assigned to branches based on their final round of branch preferences, branch capacity constraints, and their rank at West Point. Rank is based on a score computed as a linear combination of the cadet’s academic, military, and physical performance grades at West Point, with weights of 0.55, 0.30, and 0.15, respectively (Colarusso et al., 2016).¹³ Assignment is an extended version of serial dictatorship. The top-ranked cadet is assigned to their most-preferred branch, the second-ranked cadet to their top preference if there is a remaining slot, and the second choice otherwise, and so on (Sönmez and Switzer, 2013). The way in which this process differs from the basic version of serial dictatorship is described below.

¹²College Entrance Exam Rank is a composite of SAT Math, SAT Verbal, and high school rank. Whole Candidate Score is a composite of various high school academic and extracurricular measures.

¹³Since ranking is based on course performance, there might be some strategic behavior in which the cadets take classes that they expect to result in a higher grade (Colarusso et al., 2016).

The branch to which a cadet is assigned is important both for the Army and cadet. In a report that discusses Army officer retention, Colarusso et al. (2016) argue that Army workforce productivity is influenced by the talent match of officers to branches: “It is not hard to imagine how an officer might be a better talent match for one branch than another, as each does decidedly different work.” Highlighting the importance of the branch, they also point out, “Poor initial matches also have significant implications for individual officer career satisfaction and thus retention” beyond minimum commitment. An assigned branch determines jobs within the Army that the new officer will supervise, and has implications for career success, including promotion. For example, Army generals come disproportionately from Combat Arms branches (Lim et al., 2009).¹⁴ Although starting salaries are the same across branches,¹⁵ some branches, such as Infantry, are more likely to earn hazard and combat pay bonuses, and Aviators receive flight pay. Anecdotally, there is also more transferability to private-sector jobs in some branches, such as Aviation, Finance, and Engineering, than in others.

To increase officer retention, the Army instituted an optional incentive program, The Officer Career Satisfaction Program.¹⁶ This program allows cadets to trade additional years in the Army—beyond the baseline commitment of 5 active and 3 reserve years—in exchange for extra privileges. Relevant to this paper is the option to give up 3 years for getting assigned to a branch of their choice.¹⁷ When cadets fill in their final round

¹⁴ Combat Arms branches are: Infantry, Armor, Field Artillery, Air Defense Artillery, Aviation, and Engineer.

¹⁵The exception is the Medical Corps, which receives higher pay for retention purposes.

¹⁶For more information see Wardynski, Lyle, and Colarusso (2009) and Sönmez and Switzer (2013).

¹⁷Other options involve receiving funding for a graduate degree and being assigned to the Army base of choice as the initial location assignment.

of preferences, they sign a contract regarding in which of the branches, if any, they are willing to serve extra years. For any branch, the first 75% of slots are determined by class rank. Following this, the following slots are allocated to those who have signed the extra years contract, which also occurs in order of class rank. If any slots remain, they are again assigned based on class rank to those unwilling to give up the extra years of service. Cadets assigned to one of the branches as part of the initial 75% do not need to serve the extra time, even if they signed the contract for the branch; the contract is only binding if assigned as one of the bottom 25%. Sönmez and Switzer (2013) find that this mechanism does not have the desirable properties of the simple serial dictatorship, such as truth-telling and being strategy-proof.

Cadets' preferences are not known publicly, but it is common for cadets to share them with each other. Excepting those in the Office of Economic and Manpower Analysis, where data are stored, neither the administration, which includes the commandant and tactical officers, nor general faculty have access to preference data. Cadets are allowed to leave West Point, but those who do after the beginning of their junior years must repay a large sum of money or enlist in the Army, so it is much more common for a cadet to leave before the start of the junior year.

In our analysis period, cadets were not required to have majored in a subject to be assigned to a branch, though some majors are advantageous for some branches (e.g., an engineering major for the Engineering branch) (Colarusso et al., 2016).¹⁸ If cadets perform

¹⁸Colarusso et al. (2016) note that several branches now “require some domain specific education (Adjutant General, Chemical, Cyber, and Finance).” In some cases, there are other requirements, such as an eyesight requirement for Aviators, which can be met with Lasik eye surgery.

highly as officers, they are more likely to be promoted, and having branch-specific training can help with performance. However, since cadets do not officially know until the fall of senior year which branch they will be in, there is a limited number of steps they can take during senior year to modify academic training if allocated to an unanticipated branch. Most cadets who graduate do so in four years, further limiting the scope of changing majors.

Colarusso et al. (2016) describe the “smart branching” program that began as a pilot program during the period of the study. The program was designed to create better matches between cadets and branches by providing cadets with information about branches, giving them feedback on their strengths and weaknesses, and recommending which branches might be a good fit for them based on their branch preferences and academic, military, and physical potential. However, these are only recommendations, and are in no way binding. When cadets are giving their preferences, they are told that this will help them choose the right branch for them. Cadets are contacted more frequently through this program during their junior and senior years than in their freshman and sophomore years.¹⁹ When cadets are provided recommendations regarding branches, they are not pushed to change their choices as part of the program, and face no consequences if their preference rankings do not accord with recommendations. We do not view the smart branching program as problematic to our identification strategy, since it applies universally. Although it is one way for cadets to be better informed about the choices they are making, it does not alter the randomization to companies and roommates.

¹⁹During the summer between junior and senior years, cadets participate in Cadet Troop Leadership Training, which resembles an internship during which cadets become more familiar with particular branches.

1.4 Data and Descriptive Statistics

Data came from the Office of Economic and Manpower Analysis (OEMA), housed at West Point. We use data on the 2015 and 2016 graduating cohorts for company-level analysis. These include preference rankings for Rounds 1–6, and final branch assignment. We supplement this information with data on older peers from the classes of 2012–2014 in the older peer regressions. We also use roommate data from 2016; roommate-level data were unavailable for the 2015 cohort. There was some attrition among cadets; we observe only cadets conditional on graduating.²⁰ In our sample, women are ineligible to serve in all branches—they cannot serve in the combat arms branches of Infantry and Armor. As an experiment, women were allowed to select these branches for every survey except for final branch preferences. However, since during the study period women were barred from selecting Infantry or Armor as a final branch preference and because they are missing some preference data, we use only male cadets in the study. In addition to excluding women from the sample, we exclude those who do not have branch rankings for Rounds 1, 2, 3, and 6. The sample, which begins at 2,024, reduces by 335 after excluding women, and by 234 and 24 when excluding those missing early and final-round preference data, respectively.²¹ The final sample consists of 1,431 cadets.

Table 4.2.3 shows summary statistics after sample selection. The sample is 8% black, 9% Hispanic, and 9% other races. The *Cadet Entrance Exam Rank (CEER)* is a composite

²⁰Since we have only data on those who graduated, we cannot assess how attrition might have influenced analyses.

²¹Forty observations were missing from the first round, 34 from the second round, and 160 from the third; these are not mutually exclusive. The final round exclusion includes two observations that do not agree across the two final-round variables in the data.

score of SAT Math, SAT Verbal, and high school rank, and is used to rank cadets academically for admissions purposes. *Whole Candidate Score* is a metric used during admissions, and includes components such as measurements of leadership and athletic fitness, all measured before starting at West Point. Approximately one-sixth of the sample entered as an NCAA athlete. Similar proportions of cadets were enlisted previously in the Army, are the child of a parent who served in the Army, and attended West Point preparatory school.

Our regressions are identified because of variation across companies in the percentage of cadets who prefer a specific branch as their top choice. Figure 5.1.2 displays this variation for Round 1 in Panel A and for Round 2 in Panel B. Each of the histograms corresponds to a specific branch. The observations for a histogram are at the company-level, and are the percentage of the company that prefers that branch first. There is a good amount of variation across companies in preferences, particularly for the more popular branches, such as Aviation and Infantry. There is less variation for the less popular branches as few people choose these branches as their top preference.

The distribution over all cadets of the branches ranked as the top preference during Rounds 1, 2, and 6 is shown in Table 4.1.3. There is substantial gravitation to some branches, especially Infantry, Aviation, Military Intelligence, and Engineering. During Round 1, Infantry and Aviation comprised just over half of all selections. In contrast, several branches are seldom chosen as a top choice, including Chemical Corps, Ordnance, Signal Corps, and Transportation. Such variation in the popularity of the branches evidences that cadets are not selecting branches randomly during early rounds. If they were, we would expect uniform distribution across the branches. There is a fair amount of

change from the first to second rounds, with a notable decline in the number who rank Infantry as their top choice, while Aviation experienced an increase. Many of the branches that were ranked lower initially gained popularity during Round 2. The initially bottom-ranked three branches continued to grow in popularity between Rounds 2 and 6. Aviation dropped sharply between these rounds. Aggregate ordering of branches during Round 6, though not identical, is similar to what it was during the initial round.

To offer another view of the stability of preferences, Figure 5.1.3 shows the distribution of the number of times cadets switched their top preferences between rounds. Since there were 6 rounds, a cadet was able to change the top preference up to five times. A roughly equal share of cadets never switched, switched once, and switched twice. Many fewer switched four or more times. A transition matrix of top preferences between Rounds 1 and 6 is shown in Table 4.1.4. Each entry refers to the number of cadets who chose the branch on the vertical axis as the top choice during the first round, and the branch on the horizontal axis as the top choice during the sixth round. Diagonal elements, which are larger than surrounding cells, count the number of cadets who kept the same top branch between Rounds 1 and 6. Although diagonals contain many cadets, many others switched their top preferences.

Figure 5.1.4 is a Sankey diagram that illustrates the overall flows of top-ranked preferences across all branches for Rounds 1, 2, 3, and 6. Round 1 appears in the leftmost column, Round 2 is adjacent to Round 1, etc. Each round contains 17 branches, represented by a unique color. Within a round, the width of the bars represents the number of cadets who chose the branch as the top choice during that round. Across adjacent

rounds, lines from one branch to another reflect the number of cadets who chose the left-hand-side branch during the former round, and the righthand branch during the subsequent round. Much switching occurred between Combat Arms branches, in part because these are among the most popular branches. Although there is much stability across rounds, there is also a number of switches. We highlight patterns arising in several of the branches. First, many individuals shifted from Infantry between Rounds 1 and 2. Of these, a sizable portion went to Aviation, but the number switching from Aviation to Infantry was much smaller. Military Intelligence is a particularly volatile branch in the sense that there was much movement both to and from the branch across rounds. Cyber Corps drew heavily from related branches of Military Intelligence and Signal Corps.

It is informative to view the stability of preferences through a different lens by looking at what percentage of cadets ranked the same branch first in a latter round that they also ranked first in a former round. Table 4.1.5 shows this statistic for all rounds. The vertical round is the former round, and the horizontal round is the latter. The entry for Rounds 1 and 2 indicates that 58.2% of cadets kept the same top branch between the rounds, meaning that 42% switched. Less than half (45.9%) of cadets ranked this choice by Round 3. A slightly higher number (60.9%) is observed between Rounds 2 and 3. Only 38.8% of cadets selected the same branch during the final round as they did during the first. Many fewer cadets switched between successive rounds, such as between Rounds 4 and 5, and between 5 and 6, than between 1 and 2, and between 2 and 3. It is important to recall the timeline of when preferences are given. There is approximately one year between Rounds 1 and 2, and between 2 and 3, but only about 4 months between Rounds 4 and 5, and an even shorter period of approximately one month between 5 and 6. Even with that short

period between the final two rounds, a quarter of candidates switched their top preferences.

College students' major preferences change a great deal over the course of their college careers (Stinebrickner and Stinebrickner, 2013). With preference data over time, we see how much that preferences for officer branches changed over time. Table 4.1.6 Panel A shows the average position that the top-ranked branch during a former round was selected during a latter round. Panel B repeats this for second-ranked preferences. The earlier (vertical axis) and subsequent rounds (horizontal axis) function the same as in Table 4.1.5.²² If all cadets continued to rank the same branch first between rounds, the value would be 1. If all cadets ranked this branch last during a latter round, the value would be 16. In Panel A, cadets ranked their top choices from the first round about 2.9 during the second. The number falls to 3.3 by the final round. The values also drop over time between successive rounds. Between the final two rounds (5 and 6), cadets ranked the branch they rated top during Round 5, on average, 1.4 during Round 6, indicating greater stability than between all previous neighboring rounds. As expected, in Panel B, all numbers in the table are larger than those in Panel A, but the same patterns are observed.

Table 4.1.7 shows the distribution of branch assignments by final Round 6 preference ranking—the percentage of cadets assigned to serve in the branch they ranked first, the fraction that gets the second choice, etc. A high percentage (78.6%) of cadets got their first choice. Described in Section 4.2.4, we focus on first preferences during regression analysis. That such a high percentage of cadets are assigned to the branch they ranked highest shows that this top preference is most importance. The majority (93%) got one of

²²In this table, we only consider the 2015 cohort as they always had only 16 choices because Cyber was never available.

their first three. Fewer than 2% are assigned to a branch they ranked fourth.

1.5 Empirical Strategy

Since our empirical strategy relies on random assignment of students into companies, we first test the proposal that West Point randomizes cadets into companies regarding top branch preference. In other words, that cadets with similar preferences are not systematically grouped together. After the check for randomization, we present two strategies to test for evidence of peer effects—regression analysis and rank agreement. The idea behind regression analysis is to assess how a cadet’s peers in the second company influence top branch preferences given near the beginning of the third year, after exposure to the second company. The approach addresses selection and reflection biases. We present several sets of regressions—branch-specific regressions, generalized stacked-branch regressions, regressions that examine the influence of older peers, and regressions that assess roommates. Rank agreement is a summary statistic of the proximity of a cadet’s preferences to those in the company. We create an estimator that allows us to assess how average rank agreement exposure to randomly assigned cadets in the company for a school year, taking into account broad West Point-wide convergence or divergence in preferences.

1.6 Check for Randomization

One threat to identification assumptions is whether West Point solicits branch preferences from cadets and then assigns cadets to companies according to preference. To demonstrate random assignment, we use empirical p-values. Proposed by Lehmann and Romano (2005) and Good (2006), several papers use this simulation method to demonstrate evidence of random assignment to peer groups in a variety of contexts.²³

First, for each company, we count how many cadets indicated a branch within the branch categories, which are Combat Arms, Combat Support Arms, and Combat Service Support. We measure preferences in the first round for both the first and company. In the first round, preferences are measured after a few weeks of exposure to the peer group; these tests are influenced if peer effects operate in that time. However, because we also measure preferences in round 1 for the second company, this concern is minimized because cadets were not yet randomly assigned to the company when they gave preferences. Next, we create 10,000 synthetic companies for each individual company during a year by randomly assigning cadets to companies.²⁴ For each of the synthetic companies, we compare the number of cadets choosing a branch grouping as their first choice, and indicate whether a synthetic company has more cadets choosing a branch group than the real company. Finally, we divide the number of synthetic companies having more cadets

²³Carrell and West (2010) and Carrell, Hoekstra, and West (2016) use empirical p-values to evidence random assignment of students at the United States Air Force Academy. Lim and Meer (2017) use it to show how elementary schools in South Korea assign students to classes and teachers. Kofoed and McGovney (2017) use a similar simulation method to show that West Point assigns cadet companies and tactical officers randomly irrespective of gender and race.

²⁴In contrast to West Point's conditional random assignment, we employ unconditional random assignment.

choosing a branch group by 10,000 to calculate the empirical p-value. After repeating this technique for every company–year, there should exist an empirical p-value for every observed company in the data. If companies are assigned randomly regarding branch preferences, each of the empirical p-values can take a value between zero and one, and should be distributed uniformly. We then use two goodness-of-fit tests to evidence a uniform distribution—the Kolmogorov-Smirnov and Pearson χ^2 tests proposed by Ammermueller and Pischke (2009). We test each branch category separately for the two years.

Table 4.1.8 shows results from the simulation and goodness-of-fit tests. Panel A shows results for the first cadet company. While we would expect that means would be closer to 0.5, we find them to be 0.428, 0.406, and 0.380 for Combat Arms, Combat Support, and Support Services, respectively.²⁵ We also find that for the two years of data in the sample, each year passes both the Kolmogorov-Smirnov test. One of two cohorts in Support Services fails the Pearson χ^2 test.²⁶ Panel B shows results for the second cadet company. We find similar results, including one of two cohorts failing the and χ^2 test. Overall, all Kolmogorov-Smirnov tests pass and one χ^2 test fails for each company, which we interpret as evidence in favor of West Point randomly assigning cadets to companies.

²⁵One factor that may be decreasing the means is that companies are not all equally-sized.

²⁶There are several companies that have zero cadets in Support Services. These take on a 0 p-value, which may influence this test.

1.7 Regression Analysis

We describe branch-specific regressions, and then condense these into a stacked regression. Regressions are performed with a peer group, defined as own-cohort peers in the company. Following this, we extend the analysis, and define the peer group as mentors, or older peers, and roommates.

1.7.1 Branch-Specific Regressions

For the branch-specific regression, we use only the top preference. Although we do not use all of the data, the top branch preference is extremely important; nearly 80% of cadets are assigned to the branch they ranked highest during the final round (see Table 4.1.7). We take as the outcome variable a cadet’s top preference during Round 3. We use the same-cohort peers from a cadet’s second company as the peer group, overcoming selection bias with the random assignment to company. It would be ideal to use preferences measured just before the beginning of random assignment to the second company, but second round preferences were measured a few weeks into the semester. Thus, we use preferences—for both cadet and company—measured near the beginning of the first year. This introduces noise, but overcomes the reflection problem.

We run each regression separately by branch b . For each branch, a cadet contributes one observation. Cadets belong to Company c in year 2: $c(2)$. This company–year contains $N_{c(2)}$ cadets. The dependent variable is an indicator of a cadet’s top choice branch b during

preference round $r(3)$, $\text{Branch}_{i,r(3)}^b$.²⁷ We control for a cadet's preference during Round 1: $\text{Branch}_{i,r(1)}^b$. The variable of interest is the leave-out percentage of cadet i 's Company 2 peers choosing branch b as a top choice during Round 1: $\frac{1}{N_{c(2)} - 1} \sum_{j \neq i} \text{Branch}_{j,r(1),c(2)}^b$.

We estimate the following model separately for each branch b :

$$\text{Branch}_{i,r(3)}^b = \alpha_0^b + \alpha_1^b \frac{1}{N_{c(2)} - 1} \sum_{j \neq i} \text{Branch}_{j,r(1),c(2)}^b + \alpha_2^b \text{Branch}_{i,r(1)}^b + \alpha_3^b X_i + \epsilon_i^b, \quad (1.1)$$

where X_i is a vector of individual-specific controls, and $\epsilon_{i,b}$ is the error term. X_i includes racial/ethnicity variables (e.g., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. It also includes a year fixed effect. The company percentage term includes only cadets in the cadet i 's second company who are in the sample. The α_1 coefficient measures the change in the probability that a cadet selected branch b if there is a 100 percentage point increase in the number of peers in the company that also selected the branch. This increase is an enormous shift.²⁸ The dummy variable for a cadet's preference during Round r is used to control for path dependency in preferences. We use inter-branch variation of peers' occupational preferences to identify peer effects on a cadet's occupational choice. Standard errors are clustered at the company level. Since 16 coefficients were measured, we present results in graphical form.

²⁷Here and in the following regressions, we do not consider Cyber Corps. If a cadet chose Cyber Corps during the top round, each branch has a zero value.

²⁸It would typically be out of the support. Dividing by 10 gives a 10 percentage point increase.

1.7.2 Stacked Regression

Next, we combine the 17 occupational branches into a parsimonious, stacked regression model.²⁹ The effect is summarized in one coefficient, making findings easier to interpret. The dataset is structured such that the data stack each of the branch-specific observations; each cadet contributes 17 observations, one for each branch.³⁰ We also include a set of branch fixed effects. For example, if an observation corresponds to Aviation, the Aviation fixed effect takes on a value of 1 for the observation.

We estimate the following model:

$$Branch_{i,r(3)} = \beta_1 \frac{1}{N_{c(2)} - 1} \sum_{j \neq i} Branch_{j,r(1),c(2)} + \beta_2 Branch_{i,r(1)} + \beta_3 X_i + \delta^b + \nu_i \quad (1.2)$$

Terms are defined as in the previous section. The branch fixed effect δ^b controls for the Round r West Point-wide preferences cadets of the same cohort had for occupational branches. For example, if for some idiosyncratic reason Infantry is initially popular at the Academy, both cadet i and peers might pick the same branch, not because of a peer effect, but because of converging tastes and preferences. We estimate another version that includes a Branch-Year fixed effect, which controls for differing initial popularity across years. Clustering at the higher level of company subsumes the level of individual to account for the correlation of standard errors within individual. We additionally examine the

²⁹One can also imagine combining the branches into one estimate in many other ways, such as weighting the branch specific estimates by branch popularity.

³⁰We include Cyber in these regressions. Only one cohort was able to select Cyber as top branch in Round 3, which may create a small amount of bias. It was not possible to select Cyber in rounds 1, so the right-hand side Cyber variables are always zero. We include a robustness check in Appendix Table 6.2.2; we find similar results.

longer-term impact of peers by changing the dependent variable to be measured in each of Rounds 2 through 6.³¹ In these regressions, we include a Branch fixed effect. For comparison, we repeat this exercise, but do so for Company 1 peers, with the knowledge that the reflection problem may exist.

1.7.3 Older Peers

So far, the methods focus on peer effects for cadets in the same cohort. We now explore the effect of older peers, or mentors, on the branch choices of younger cadets, and examine the relative importance of peers of different ages. We consider separately the effect of the branch preferences of freshman cadets on cadets who were seniors in this company. We modify the setup slightly. We now measure the dependent variable for cadet i during Round 2, and the initial preference for cadet i during Round 1. However, we now consider the peer group of Company 1. Since we observe only final, Round 6 preferences of older peers, we do not have preferences measured before cadet i was assigned to the company of peers. We address whether older peers had a lasting impact on cadet i when a freshman. Since it is less likely that seniors would be influenced by the branch preferences of freshman than sophomores and juniors would be, and since seniors' final preferences are measured toward the beginning of cadet i 's first year (while the juniors' [sophomores'] final preferences were measured when cadet i was a sophomore [junior]), we focus on

³¹We include Cyber in each of these rounds; because Cyber was not an option in Round 2, results should be interpreted accordingly

seniors. The estimation equation is:

$$Branch_{i,r(2),c(y)} = \theta_0 + \theta_1 \frac{1}{N_{c(1)}^{Senior} - 1} \sum_j Branch_{j,r(1),c(1)}^{Senior} + \theta_2 Branch_{i,r(1)} + \theta_3 X_i + \delta + \omega_i, \quad (1.3)$$

where $N_{c(1)}^{Senior}$ is the number of senior cadets in cadet i 's Company 1.

1.7.4 Roommates

In addition to being assigned randomly to a company, cadets are assigned randomly to a roommate (or occasionally more than one). This assignment applies to the first year. During the second year, some cadets are assigned roommates randomly and others are not; it depends on the tactical officer of the company. Consequently, we focus on first-year roommates. One caveat is that the reflection problem might come into play since preferences were measured a few weeks after cadets were exposed to a roommate. Instead of using the third preference round as the dependent variable, as above, we use the second round, which was measured much closer to exposure to a first-year roommate. Usable roommate data are available only for the 2016 class, so we do not include year or Branch-Year fixed effects.³²

Let the total number of individuals, including cadet i , in the room be $N_{c(1)}^{Room}$, and the top preference of roommate j be $Branch_{j,r(1),c(1)}^{Room}$, where $c(1)$ indicates that it is the first company roommate. We modify the stacked regression above to test whether there are

³²We exclude the small number of cadets with missing roommate data.

roommate effects on preferences:³³

$$Branch_{i,r(2)} = \gamma_0 + \gamma_1 \frac{1}{N_{c(1)}^{Room} - 1} \sum_{j \neq i} Branch_{j,r(1),c(1)}^{Room} + \gamma_2 Branch_{i,r(1)} + \gamma_3 X_i + \zeta_i, \quad (1.4)$$

1.7.5 Combined Peer Groups

Finally, we consider a regression of third round peer preference on several different measures of peer groups in the same regression: Company peers (excluding roommates), Company 2 peers, Company 1 seniors, and Company 1 roommate. We consider only the 2016 cohort because we do not have usable roommate data for the 2015 cohort.

1.7.6 Branch-Specific Regressions Results

First, we estimate 16 separate regressions, with sophomore branch preferences as the dependent variable with a cadet's peers' freshman occupational branch preferences. Figure 5.1.5 shows these results. Each dot, or point estimate, and line, or confidence interval, of the same color correspond to a branch regression, which is labeled using a two-letter abbreviation, with a correspondence to branch name found in Table 4.1.1. Values plotted are from the leave-out mean peer measure. Only one of the estimates is statistically significant at a 5% level, and the coefficients are not systematical, either above or below zero.

³³We do not consider Cyber Corps in this regression.

1.7.7 Stacked Regression Results

Main stacked regression model estimates are reported in Table 4.1.9. We include a cadet’s third-round preference as the outcome variable, and own and company-cohort first-round preferences as explanatory variables. The coefficient for “Freshman First Preference” in Column 1, which refers to a cadet’s own preference during the initial round, is highly predictive of the third-round preference; controlling for peer preferences, a cadet is predicted to select the same branch during the first and third rounds approximately one-third of the time. An increase of 100 percentage points in the number of peers who prefer the branch during the earlier round associates with a 38.3 percentage point increase in the probability of choosing the branch. Results are nearly identical in Column 2, after including the controls listed in Section 1.7.2. The own preference coefficient is very stable when we add branch fixed effects in Column 3. However, the statistically significant and positive peer result in the previous two columns becomes very small, negative, and not statistically significant at the 10% level. The branch fixed effects constrain the model to identify the causal effect from the within-branch variation in the percentage of cadets choosing a specific branch. Therefore, we cannot claim that the coefficient estimates in Columns 1 and 2 are a peer effect, but they might identify similar converging tastes and preferences among cadets that are a result of processes such as learning about the Army and their own abilities as they remain at West Point for an additional academic year. If the branches that cadets initially prefer correlate with the branches that latter cadets (who switch their preference) prefer, we expect a positive coefficient in the absence of branch fixed effects, even in the presence of random assignment to peers. Column 4 adds Branch-Year fixed

effects, without much change from the previous column.

We repeat this analysis, except we examine the longer-run impact of first-company peers by changing the dependent variable to the final branch preferences, with each column being the result of using a different preference round. Each regression uses branch fixed effects. Results are shown in Table 4.1.10. Column 3 of Table 4.1.9 is equivalent to column 2 in this table. The coefficients are negative and statistically significant in rounds 4 though 6, providing some evidence of small, negative peer effects. Appendix Table 6.2.3 shows the results when using Company 1. Each of the company-level variables is close to zero.

1.7.8 Older Peers Results

We consider mentor effects, specifically the effect of seniors on freshmen cadets, in Tables 4.1.11. A similar pattern as those in earlier specifications was found; when branch and branch-year fixed effects are not included, there is a positive relationship between the number of peers who prefer the branch first and the top choice of cadet i . A cadet's choice from the previous round continues to be highly predictive of the current top round choice. However, when fixed effects are included, the mentor effect disappears. In each case, it is close to zero, providing little to no evidence that there are mentor peer effects. Although results are not shown, similar qualitative findings occur with sophomore and junior older peers.

1.7.9 Roommates Results

We turn to roommate results shown in Table 4.1.12 for first-year roommates. Again, a similar pattern was found, where there is a positive, significant coefficient that is attenuated to zero after controlling for branch fixed effects. The coefficient for a cadet's own first preference is much larger than in prior results because the outcome variable is now the Round 2 top preference, as opposed to the round 3 top preference, and the own first preference from the round is more predictive of the preference in the second round than in the third. We find no evidence of peer effects among roommates.

1.7.10 Combined Peer Groups Results

Our final regression, Table 4.1.13, includes many different measures of peer groups in the same regression. The outcome variable is third round preference. After including branch fixed effects, none of the measures are statistically significant.

1.8 Rank Agreement

We introduce a summary statistic called Rank Agreement, a measure of how closely a cadet's preferences align with the preferences of the remainder of peers in the company, to simplify multi-dimensional, branch-preference rankings while still using all of the information contained within. A cadet ranks branches from 1 to 16 several times during

attendance at West Point, including near the beginning of the first year in the first company, and again near the beginning of the second year in the second company. We assign each cadet a rank agreement value for Round 1, computed in relation to other cadets in the first company. We assign another value for Round 2, but in relation to cadets in the first company. This is done so we can evaluate whether each cadet's preferences are closer to other cadets' preferences in the first company after a year with them. We repeat this process for Rounds 2 and 3, computing both in relation to the second company. We first present a generalized formulation of rank agreement using matrix notation. A simplified, artificial example of how rank agreement is computed appears in the Appendix. Finally, we discuss basic properties of rank agreement, how it is used, and potential shortcomings of the approach.

1.8.1 Rank Agreement Calculation

There are i cadets, each of whom belongs to company c in year y , $c(y)$. Each company–year consists of $N_{c(y)}$ cadets, for N_y cadets across companies during year y . The round in which a cadet's preferences are given is indexed by r . Preferences in Rounds 1, 2, and 3 are given in years 1, 2, and 3, respectively. There are B branches from which a cadet can

choose: $Branch_1, \dots, Branch_b, \dots, Branch_B$, collected in $Branch$ vector:

$$Branch = \begin{bmatrix} Branch^1 \\ \vdots \\ Branch^b \\ \vdots \\ Branch^B \end{bmatrix}. \quad (1.5)$$

Each cadets' preferences for the branches are given by:

$$Pref_{i,r} = \begin{bmatrix} Pref_{i,r}^1 \\ \vdots \\ Pref_{i,r}^b \\ \vdots \\ Pref_{i,r}^B \end{bmatrix}, \quad (1.6)$$

where $Pref_{i,r}^b$ denotes cadet i 's ranking of $Branch^b$ during Round r . Each of the B $Pref_{i,r}^b$'s takes on a distinct value from 1 to B , with the most-preferred branch having a value of 1. For each cadet i , Round r , and company–year $c(y)$, the leave-out mean of the preferences

of the other cadets during Round r , company $c(y)$ is calculated as:

$$\overline{Pref}_{i,r,c(y)}^{Leave-Out} = \begin{bmatrix} \frac{1}{N_{c(y)}-1} \sum_{j \neq i} Pref_{j,r}^1 \\ \dots \\ \frac{1}{N_{c(y)}-1} \sum_{j \neq i} Pref_{j,r}^b \\ \dots \\ \frac{1}{N_{c(y)}-1} \sum_{j \neq i} Pref_{j,r}^B \end{bmatrix}, \quad (1.7)$$

Differences in cadet i 's preferences and others in the branch are:

$$DiffPref_{i,r,c(y)} = abs(Pref_{i,r} - \overline{Pref}_{i,r,c(y)}^{Leave-Out}) \quad (1.8)$$

r and y do not require the same value in the equation. For example, $DiffPref_{i,1,c(1)}$ and $DiffPref_{i,2,c(1)}$ are both computed using the same cadets from company–year $c(y)$, but preferences used during calculation for both cadet i and the other cadets in the company are from Round 1 in the first case and Round 2 in the second.

$DiffPref_{i,r,c(y)}$ is sorted by cadet i 's preferences to produce $DiffPref_{i,r,c(y)}^{Sorted}$, such that the first element corresponds to the branch that cadet i preferred most.

We shift focus to weights applied to calculate rank agreement for cadet i . Suppose cadets at West Point are assigned to branches that they ranked 1st, 2nd, etc. with the

following percentages:

$$RealizedPercentages = \begin{bmatrix} Perc_1 \\ \dots \\ Perc_b \\ \dots \\ Perc_B \end{bmatrix}, \quad (1.9)$$

where $Perc^b$ represents the percentage of cadets assigned their b th-preferred branch. We observe the percentage of cadets assigned to their 1st, 2nd, etc. branch preferences from their final preference ordering. The final preference ordering is used to allocate cadets to branches. As in Table 4.1.7, not all cadets are assigned to their most-preferred branches, but a high percentage are. Only a small percentage are assigned to a branch ranked at the bottom of their ordering. We use this to weight the elements of the sorted difference vector such that the most-preferred branch is weighted by the percentage of cadets whose realized branch outcome is the one they ranked first, etc.

This vector is used to weight $DiffPref_{i,r,c(y)}^{Sorted}$ to produce rank agreement for cadet i during Round r in relation to company–year $c(y)$:

$$RankAgreement_{i,r,c(y)} = RealizedPercentages' \cdot DiffPref_{i,r,c(y)}^{Sorted} \quad (1.10)$$

After $RankAgreement_{i,r,c(y)}$ is computed for all cadets for Round r and year y , the average is calculated for all cadets in the cohort, N_y , to give average rank agreement:

$$\overline{RankAgreement}_{r,c(y)} = \frac{1}{N_y} \sum_i RankAgreement_{i,r,c(y)} \quad (1.11)$$

Lower rank agreement means greater agreement on preferences. If all cadets in a company share the same preference ranking, the rank agreement will be zero for each. The maximum value that rank agreement can be depends on several factors, including the number of branches and the weighting vector. Rank agreement has several desirable properties. It allows the dimensionality of the proximity of preferences to be reduced to a single number. Consequently, it uses all information of the preferences while placing more weight on rankings that more often determine to which branch a cadet is ultimately assigned.

1.8.2 Using Rank Agreement to Test for Peer Effects

To test for peer effects, we employ an empirical p-value-like strategy in which we compare the observed rank agreement with the distribution of rank agreement values calculated from assigning cadets to 1,000 synthetic companies.

At first glance, one might simply compare the rank agreement after exposure to the company with the rank agreement before exposure to the company. If it has decreased, then this is evidence of peer effects because a smaller rank agreement indicates more similar preferences. However, this is not correct because there are more forces at play that influence how cadets rank branches. One possibility is that there is common, West Point-wide convergence to some branches and away from others. This could be orthogonal to company-peer effects, and might result from cadets acquiring greater knowledge of which branches are desirable. An extreme example illustrates this point. Suppose all cadets are assigned to their top-ranked branch, so that the branch gets the entirety of the weight. Sup-

pose also that the top preference of cadets within companies and across West Point follow a uniform distribution, meaning cadets are equally likely to rank any branch as their top choice. Average rank agreement is computed to be some positive amount for the initial preference round. During the school year, all cadets learn, independent of social influence, that Branch A is more desirable than all other branches. This leads all cadets to rank Branch A in first place in the latter round, causing average rank agreement to be zero, even though no peer effects were in operation.

Because of this convergence, we account for aggregate trends by constructing synthetic companies. Keeping cadets' actual preferences, we randomly assign them to a synthetic company.³⁴ We then compute rank agreement in relation to the synthetic companies. We repeat this process 1,000 times in order to construct a distribution of rank agreement in which to compare the value completed from the true companies. We present results graphically and compute a p-value based on where the true value lies in the distribution.³⁵ Very small p-values are evidence of positive peer effects (preferences converging) and very large p-values are evidence of negative peer effects.

Similar to the regression analysis, we focus on second company peers in our main rank agreement results,³⁶ allowing us to observe rank agreement both before and after exposure to the company. Specifically, we consider the $\overline{RankAgreement}_{1,2}$ (preference round 1, company 2) and $\overline{RankAgreement}_{r3,2}$ (preference round 3, company 2), both in

³⁴We keep the sizes of the companies the observed in the data, but do not conditionally-randomly assign as is done in practice. Results should be interpreted accordingly.

³⁵To compute the p-value, we divide the rank in the distribution of the actual rank agreement by 1,001.

³⁶When computing rank agreement, we removed Cyber from the preferences and re-ranked the preferences such that each cadet had 16 preferences.

relation to their corresponding synthetic distributions. In Appendix results, we consider the first company, with preference rounds 1 and 2. Both of these are after at least some exposure to peers, as preferences are elicited several weeks into the semester.

The weighting vector, *RealizedPercentages*, that we use is constructed using the final preferences and branch assignments from the cadets in our sample.³⁷ We note that this is endogenous. We assume that cadets have a rough sense of these percentages, and that they are roughly the same over time, which appears to be the case based on past data.

1.8.3 Rank Agreement Results

Rank agreement results appear in Table 5.1.6. Panel A displays rank agreement for company 2, preference round 1. Preferences are measured before assignment to the second company. The distribution of synthetic rank agreement is plotted, with the actual rank agreement indicated by the red line. The actual rank agreement falls squarely in the distribution and has a p-value of 0.55, indicating that preferences in this round in relation to this company are similar to what happens when randomly assigning cadets to companies.

In Panel B, the synthetic rank agreement distribution has a higher than in Panel A. In other words, preferences overall have gotten further apart by round 3 in relation to round 1. While the actual rank agreement falls more to the left of the distribution as compared to its relative rank in the distribution in Panel A, it has a p-value of 0.16; we therefore reject

³⁷The weights in the weighting vector include assignments to Cyber, and the 2016 cohort had 17 choices. Nobody got assigned to their 9th or 17th choice. In calculating weights, we exclude 13 observations that were assigned Cyber despite not ranking it and 1 observation that had a missing assignment variable.

the hypothesis that the peers in the second company led to rank agreement being closer.

Appendix Table 6.3.1 repeats the rank agreement analysis, but for company 1. Panel A shows first round preferences. The p-value is very low and statistically significant, indicating that preferences of cadets are closer together than would be expected. Because preferences are measured a few weeks into the semester, this is some evidence for peer effects.³⁸ By round 2, after one year of exposure to the cadets in the company, the p-value is again in the middle of the distribution, meaning that actual preferences look like synthetic preferences. It could be that there are initially peer effects for the first preference round, but these rapidly decline after cadets spend more time at West Point.

1.9 Conclusion

This paper contributes to our understanding of how preferences generally and over occupations are shaped. West Point is a context conducive to the study of peer effects and occupation, because it assigns cadet companies and roommates randomly. The data are advantageous because West Point polls cadets each year over a finite set of occupations six times during their time at the academy. We exploit the repeated panel nature of a dataset on cadet occupational preferences to illustrate the evolution of these preferences. We find that preferences for top branches shift and stabilize over time. Using random assignment of cadets to peer groups to overcome selection bias and pre-measured preferences to

³⁸Our regression analysis focuses on the second company with preferences measured in the first round to address the reflection problem; we did not run a regression of own first-round preferences on only first-company first-round preferences.

overcome reflection bias, we assess the role of peers on occupational preferences, finding overall little evidence of peer effects in this context once overall West Point preferences for branches are controlled for in the regression analysis. However, without accounting for popularity of branches and the natural sorting of cadets across occupations given talent and comparative advantage, we find an illusion of peer effect. This occurs even with random assignment. Findings should assist researchers with ensuring that aggregate trends in preferences are not interpreted as peer effects in similar empirical context. We additionally introduce the rank agreement, which may provide some evidence of peer effects. Further research on this question in other contexts, including at earlier points in life, are warranted, as is further exploration of other determinants of preferences for occupations.

CHAPTER 2

ARE HIGH-QUALITY PHD PROGRAMS AT UNIVERSITIES ASSOCIATED WITH MORE UNDERGRADUATE STUDENTS PURSUING PHD STUDY?

2.1 Introduction

This study examines the characteristics of an institution that correlate with the share of its BAs that ultimately receive a PhD, or its PhD production rate. We examine four fields separately, the humanities, life sciences, physical sciences, and social sciences. Despite long reported poor job market prospects for new humanities PhDs, graduate programs in the humanities have continued, sometimes with smaller cohort sizes, to churn out PhD students in the humanities. On the other hand, STEM fields, the life sciences and physical sciences, strongly contribute to the economy, foster innovation, and lead PhDs in these fields to receive high salaries. We study a number of factors of institutions, such as expenditure measures, the number of total students, and other student body characteristics that might be associated with predicting the PhD production rate. In addition to producing PhDs, departments also educate undergraduate students, and one focus of this paper is to see whether the scope and quality of an institution's PhD programs are related the likelihood that its undergraduate students go on to receive PhDs. Put another way, we are interested in the dual role that departments at universities play in producing new PhDs and in generating undergraduate student interest in going on for doctoral degrees, typically at other institutions.

We focus our attention on doctoral institutions.¹ Using data from the National Science Foundation's *Survey of Earned Doctorates* (SED) linked to the IPEDS *Completions Survey*, we calculate that for 1980-2001 BA recipients, the fraction of BAs in a given field that receive PhDs in the same field is higher at *baccalaureate institutions*² than at *doctoral institutions*.³ However, because of the differences in the sizes of these two types of institutions, doctoral institutions produce far more total PhDs than baccalaureate colleges. Doctoral institutions produce 5.8 times as many PhDs than do baccalaureate colleges for the physical sciences; they also produce 2.7 times as many for the humanities, 4.7 times as many for the life sciences, and 3.7 times as many for the social sciences.⁴ With the exception of life sciences, the ratios are even higher for Master's colleges.⁵ Therefore, knowledge of the characteristics of the Carnegie doctoral institutions that lead to more undergraduates going on to receive PhDs is of importance in understanding the determinants of PhD production.

¹We define doctoral institutions as 2015 Carnegie Categories 15: Doctoral Universities: Highest Research Activity; 16: Doctoral Universities: Higher Research Activity; and 17: Doctoral Universities: Moderate Research Activity.

²We define baccalaureate as 2015 Carnegie Categories 21: Baccalaureate Colleges: Arts and Science Focus; 22: Baccalaureate Colleges: Diverse Fields; and 23: Baccalaureate/Associate's Colleges: Mixed Baccalaureate/Associate's. We do not consider category 14: Baccalaureate/Associate's Colleges: Associate's Dominant.

³This calculation is for BA-years 1980-2001 and consider only those with a time-to-degree of 15 years or less, which contrasts from the approach later in the paper, where the maximum time-to-degree varies by field. The total number of PhDs in a given field is divided by the total number of BA recipients in the same field. Some of the PhD recipients received BAs in other fields. We include only observations for which there is a match with the Completions Survey data. The institution must have produced at least 5 PhDs over the SED data, without regard to TTD. Other sample restrictions are as in Section 7.1.

⁴This calculation makes use of sample restrictions as in Section 7.1, including that the institution must have produced at least 5 PhDs, but here consider all TTDs, not just up to 15 years. The BA years are 1980-2005.

⁵We define Master's institutions as 2015 Carnegie Categories 18: Master's Colleges and Universities: Larger Programs; 19: Master's Colleges and Universities: Medium Programs; and 20: Master's Colleges and Universities: Small Programs.

Lemke (2009) is among the latest studies to examine the determinants of undergraduate institutions in “generating” (producing) PhD recipients.⁶ This study computed a PhD production rate by dividing publicly-available data on the number of PhDs in all fields earned by alumni of selective liberal arts colleges by the number of BAs earned by the students at these institutions. Because of this study’s use of public-use data, while it was known *where* a PhD recipient received her BA, it was not known *when* he or she received it. Lemke (2009) assumed a 6-year time-to-degree (TTD) window from BA to PhD, which often is too short of a window. We use restricted-access, individual-level SED data to observe the actual matches of where and when each individual PhD graduated from his or her BA institution. For each field, we compute one observation per institution-BA year on the percentage of BAs who received PhDs. This is in contrast to the one observation per institution in the Lemke (2009) study, where both the numerator and denominator of the dependent variable are averages over different 10-year periods.

When using individual-level data to compute the PhD production rate, there is an inherent truncation problem. Not all PhDs from a given BA cohort are observed because some individuals complete their PhD after the conclusion of the data. The truncation problem increases in severity for later cohorts. We use truncation correction methodology, where early cohorts’ patterns of time-to-degree are used to predict the number of PhDs that later cohorts will eventually produce. After correcting for truncation, we examine which features of the institution predict the share of graduates that ultimately receive a PhD.

⁶Many other studies have addressed the determinants of PhD students’ times to degree and completion rates. See, for example, Knapp and Goodrich (1951), Tidball (1986), Schapiro et al. (1991), Ehrenberg et al. (2007), Groen et al. (2008), Ehrenberg et al. (2009), Bowen and Bok (2016), and Groen (2016).

In our main specification using the full sample, we find the number of an institution's departments that are in the top 10th percentile of the National Research Council (NRC) 1995 doctoral rankings in each field is one of the three variables that is correlated with the PhD production rate across each field. Put simply, the number of high-quality doctoral programs is an important predictor of the rate at which undergraduates go on for PhD study. The other two variables that are statistically significant across all fields are the incoming student 75th percentile test scores, which is positively related with the PhD production rate, and the total number of students, which is negatively related. Other variables, such as institutional expenditures per student and the percentage of total BAs in the field, are significant for some, but not all, fields. We also restrict the sample to Carnegie category 15 institutions (Highest Research Activity), which have the highest level of research activity. Finally, we conduct an analysis along the lines of that done in Lemke (2009). For both, we find generally similar results to the baseline specification, although the NRC top 10th percentile variable is significant in only two of the four fields, life sciences and social sciences.

2.2 PhD Production

Following Lemke (2009), we assume that characteristics of the undergraduate institution generate, or produce, PhDs. We first list the explanatory variables considered in the analysis. We break down these characteristics into four categories: institutional, expenditure, student body, and graduate programs. Appendix 7.1 provides details on the data.

Institutional Characteristics

First, we include the student-faculty ratio. The smaller the class size, the greater the potential for students to interact with the professor, which may lead to the students being more interested in pursuing graduate studies. Next, we include a public institution indicator. If the quality or amount of interaction between faculty and students differs between public and private institutions, this may lead to differing propensities to pursue a PhD. The mission of the institution may be generally different between the two types and may also influence the propensity.

Expenditures

We include instructional expenditures per student. Greater instructional expenditures—of which faculty salaries are a large part—may attract faculty who are more research productive. Students may then want to emulate the career paths of these faculty by pursuing PhDs themselves. We also include research expenditures per full-time faculty. A higher level of research spending likely leads to higher quality research. This may influence students, making a career in research via a PhD more attractive.

Student Body Characteristics

As a proxy for the institution’s academic ability, we consider first-year incoming student test scores, specifically a combination of 75th percentile SAT and ACT scores, weighted by the percentage of scores from each. We emphasize that the scores are for

first-year, not seniors. We additionally include several measures of the size and composition of the student body: the total number of students (including graduate and professional students), the percentage of the students that are undergraduates, the percentage of undergraduates that are female, and the percentage of undergraduates that come from underrepresented minority groups. Because the importance on campus of, say, the humanities could play a role in generating interest in humanities PhD study, in the humanities regressions, we include the percentage of BAs that were received in the humanities in our humanities equations. We similarly include the percent of degrees received in the other three fields in their relevant equations.

Graduate Programs

Finally, we are interested if having highly-ranked doctoral programs is associated with having a higher PhD production rate. More highly-ranked doctoral programs is presumably correlated with having more faculty at the cutting edge of their field, and, to the extent that these faculty interact with undergraduate students, this may influence students' decisions to obtain a PhD in the humanities. Thus, we include a set of variables indicating the number of humanities' doctoral programs an institution has in several percentile bins: 0-10, 11-25, 25-50, and 51-100.

2.3 Estimating Equation

In order to obtain estimates, we employ weighted least squares regression with the PhD production rate, defined in Section 2.4.3, as the dependent variable. We run the regression separately by PhD field. Specifically:

$$PhdProdRate_{i,f,y} = \alpha_0 + \alpha_1 IndepVars + Carnegie_i + Year_y + \epsilon_{i,f,y},$$

where, for institution i , field f in BA-year y , $PhdProdRate_{i,f,y}$ is the PhD production rate defined in Section 2.4.3, $Year_y$ are year fixed effects, and $\epsilon_{i,f,y}$ is the error term. $Carnegie_i$ is the Carnegie category of the institution, with category 15 (Highest Research Activity) being the omitted category. We include these because institutions vary substantially across Carnegie categories (see Section 2.5). $IndepVars_{i,f,y}$ is a vector containing the explanatory variable (and missing variable indicators) as described in Section 2.2: *Stud – FT Fac Ratio, Public, Instr. Exp./Student, Rsch. Exp./FT Fclty, 75 Percentile Score, Total Students, % UG, % Female, % Minority, %(Field) BAs of Total, #NRC 0 – 10 in (Field), #NRC 11–25 in (Field), #NRC 26–50 in (Field), #NRC 51–100 in (Field)*. $(Field)$ refers to field of study (Humanities, Life Science, Physical Science, and Social Science), with the field matching the field currently in consideration.

If an observation is missing or has 0 BAs—the denominator of the PhD production rate—then it is not included in the regression. We weight the institution-field-BA year observation by the number of field-specific BAs.⁷ Robust standard errors are clustered at

⁷It can happen that an institution is predicted to have a very high (even greater than 1) PhD share. In such cases, it is likely that the institution awards a very small number of BAs in this field, but produced a

the institution-level.

In addition to the main regression above, we also present results where we restrict the sample to Carnegie category 15 institutions (Research Activity), which have a higher PhD production rate as well as a smaller amount of missing data as compared to the other two categories.

Finally, we perform an analysis in the spirit of Lemke (2009), although there are some differences. As opposed to constructing the PhD production rate as described in Section 2.4, we instead divide the total number of PhDs produced over the 1994-2003 time period by the total number of BAs produced over the 1989-1998 time period. Truncation correction is explicitly not used in this analysis because TTD is assumed and a PhDs are not matched to their actual BA-year. We use the same set of institutions as appear in the main analysis, and there is now only one observation per institution. If an observation produced PhDs but had no BAs, it is excluded from the regression. For the explanatory variables, we take the average value over the same 1989-1998 period among non-missing observations. If all observations for a variable are missing, we assign that variable a value of 0 and a corresponding missing variable a value of 1. We use robust standard errors, and weight by total BAs over the 1989-1998 period.⁸

disproportionately high number of PhDs. This can be because a relatively large share of the BAs in the field went on to earn a PhD or BAs in other fields went on to earn a PhD. Because there are very few BAs, these observations will receive a small weight, and not be very influential in the estimates.

⁸Results differ somewhat if we do not weight by BAs. An advantage of using weighting is that there are some observations that have very few BAs in a given field, but produce many PhDs, including those with BAs from a different field. This causes the PhD production rate to be very large, even over 100. Weighting greatly reduces the influence that such observations have in the regression.

2.4 Constructing the PhD Production Rate and Correcting for Truncation

We define the PhD production rate as the fraction of BAs for a given institution-BA year that obtain a PhD. We return to the details of its construction in Section 2.4.3 after first introducing the concept of maximum TTD in Section 2.4.1 and, in Section 2.4.2, how we employ truncation correction to address not being able to observe PhDs granted after the SED sample ends.

2.4.1 Maximum TTD

We measure the PhD production rate for students who earn a PhD within a given time, which we denote as the the maximum TTD. Earning a PhD often takes many years. To illustrate, Figure 5.2.1 plots the distribution of TTD by field for BA years 1980 and 1981, the two earliest years in the sample. Because the final year of the SED sample is 2016, the longest observable TTD for BA-year 1980 is 36 years and for BA-year 1981 is 35 years, and all PhDs earned after this are truncated. This figure shows that nearly all PhDs are earned by 30 years for the life and physical sciences, with a larger tail for the humanities and social sciences. The peaks of these latter two fields are also longer, suggesting a longer average TTD.

This truncation problem is the most pronounced for BA-year 2005, the final BA-year in the sample. For this year, the longest observable TTDs is only 11 years. We use the

truncation-correction technique presented below in Section 2.4.2 to estimate how many PhDs will be produced by the maximum TTD.

The higher the maximum TTD, the higher the number of PhDs are observed. But there is a tradeoff: The higher the maximum TTD, the fewer the number of BA-years there are for which we observe all PhDs for the maximum TTD, and the more truncation correction is necessary. In this paper, we use field-specific maximum TTDs based on Table 4.2.1. This table displays separately by field the cumulative percentage of PhDs that were received by each TTD for BA-years 1980 and 1981. For each field, we select the smallest TTD for which at least 90% of PhDs had been earned.⁹ Thus maximum TTDs of 23, 20, 16, and 22 years are used for humanities, life sciences, physical sciences, and social sciences, respectively. Choosing the maximum TTD this way puts all of the fields on an equal playing field, taking into account the different TTD distributions. This means, however, that the extent to which we must use truncation correction and thus artificially generate the outcome variable differs by field, being the most in the humanities and the least in the physical sciences. For the humanities, we observe the maximum TTD for BA-years 1980-1993 (because we observe PhDs until 2016), and use truncation correction for remaining BA-years, 1994-2005.

⁹In this table, TTDs of 36 years are combined with TTDs of 35 years; only the 1980 cohort has an observed TTD of 35 years.

2.4.2 Truncation Correction

We employ the simple truncation correction technique used in Prenovitz et al. (2016) to estimate how many PhDs would be produced if we had observed PhDs up to the maximum TTD. Using the number of PhDs produced by 2016 as a baseline, we predict how many more would be produced had we observed the entire maximum TTD. The training sample that goes into the prediction uses the BA-years for which we observe the maximum TTD; as an example, for the humanities this is 1980-1993. Using these years only, we compute the percentage of PhDs were received by each TTD, up to to maximum TTD. These cumulative percentages are then applied to the remaining years (1994-2005 for humanities) at the institution-BA year level to estimate the number of PhDs.

An example will illustrate. Suppose that in the humanities, a given institution-field with BA-year 2003 produced 6 PhDs by 2016. Because the maximum observed TTD is 13, the cumulative percentage for a TTD of 13 will be used. Suppose that this is calculated to be x . The predicted number of PhDs that will be produced by 2026 (23 years after BA-year) is computed to be $(\frac{1}{x})6$.

Here, we formalize the above idea. The prediction is based on the years for which the full TTDs exist (1980-1993 for the humanities). Restricting the sample to these years only, we compute the cumulative percentage of degrees that occur after TTDs of 4,...,MaxTTD, where MaxTTD is the maximum TTD (4,...,23 years for the humanities).

Specifically, we first compute the cumulative number of degrees by each TTD for each institution, separately by field. Define the number of degrees granted by institution i , field

f , BA-year y , at TTD t , to be $PhD_{i,f,y}^t$, where $t \in \{4, MaxTTD\}$.¹⁰ The cumulative number of degrees for institution i by TTD t is then:

$$CumPhD_{i,f,y}^t = \sum_{t=4}^t PhD_{i,f,y}^t.$$

These are aggregated across institutions to the TTD-level to produce $CumPhD_f^t$:

$$CumPhD^t = \sum_i CumPhD_{i,f,y}^t.$$

Sum across TTDs to produce the total number of PhDs at the maximum TTD:

$$TotalPhD_f = \sum_{t=4}^{MaxTTD} CumPhD_f^t.$$

Now, obtain the cumulative percentage of degrees granted for each TTD:¹¹

$$CumPct_f^t = \frac{CumPhD_f^t}{TotalPhD_f}.$$

We next apply the cumulative percentages to each institution-BA year for years for which we do not observe the max TTD (1994-2005 for humanities) to obtain the estimated total number of degrees for each institution-BA year. This is done by multiplying the number of degrees produced by the maximum observed TTD by the inverse of the corresponding cumulative percentage:

$$PredictedPhD_{i,f,y} = \left(\frac{1}{CumPct_f^{MaxObsTTD}} \right) CumPhD_{i,f,y}^{MaxObsTTD},$$

¹⁰We reassign TTDs of less than 4 years to be 4 years.

¹¹The cumulative percentage is actually a fraction, not a percentage.

where $MaxObsTTD$ is the maximum TTD that it was possible for the BA-year to have; this is equivalent to $2016 - y$. For example, $MaxObsTTD$ for BA-year 2003 is 13. Note that the $CumPct$ value takes on the value corresponding to $MaxObsTTD$. If an observation produced 0 PhDs by 2016, we assign it a predicted value of 0. This method uses information from past years to predict PhD rates for future years. It assumes that the production patterns of PhDs does not vary by year, or that the pattern from the earlier estimating sample will occur in the later prediction sample. For each field, we use the entire sample of institutions in the training sample, regardless of which Carnegie category the institution is in.

2.4.3 PhD Production Rate

After implementing the truncation correction, we compute the PhD production rate for each institution BA year, and do so separately by field. For BA-years for which we observe the max TTD (1980-1993 for humanities), this is:

$$PhdProdRate_{i,f,y} = 100 \cdot \left(\frac{CumPhD_{i,f,y}^{MaxTTD}}{TotalBA_{i,f,y}} \right),$$

and for BA-years for which we do not observe the max TTD (1994-2005 for humanities):

$$PhdProdRate_{i,f,y} = 100 \cdot \left(\frac{PredictedPhD_{i,f,y}^{MaxTTD}}{TotalBA_{i,f,y}} \right),$$

where $CumPhD_{i,f,y}^{MaxTTD}$ and $PredictedPhD_{i,f,y}^{MaxTTD}$ are defined the same as previously, and $TotalBA_{i,f,y}$ is the total number of BA degrees for institution i in field f in year y .¹² Because we multiply by 100, the PhD production rate is a percentage.

Note that the numerator includes all individuals who earned a PhD in the field, whether or not they also earned a BA in the same field. To illustrate, Table 4.2.2 shows the BA fields of study for PhDs in a given field. Over 93% of Physical Science PhDs also received a BA in Physical Sciences, while only 64% of Social Science PhDs received a BA in the Social Sciences. The denominator of the PhD production rate includes only the major that is listed as the first major. If the individual dual-majored in a field other than their first-listed major, then this will not be included in the other field.

2.5 Descriptive Statistics

We first describe several plots relating to the PhD production rate, and then proceed to present summary statistics for the independent variables considered in the analysis.

PhD Production Rate

Figure 5.2.2 plots the total number of PhDs produced for each BA-year from 1980 to 2005. This is broken up by PhD field. There were fewer PhDs granted in the humanities and social sciences than in the life and physical sciences. These latter two fields both experienced a large amount of growth beginning in the mid-1990s; the former fields also

¹²Observations with 0 BAs produced in the field in the BA-year are not included in the regressions.

experienced growth, but to a lesser degree.

Panel B of Figure 5.2.2 also plots the total number of PhDs produced over time, but breaks it up by the Carnegie category of the university. Relatively few PhDs came from institutions in category 16, and even fewer from category 17. Both categories remained fairly steady over the time horizon, with a small increase towards the end. Institutions in category 15 (Highest Research Activity), meanwhile, produced many more PhDs, and steadily increased this number beginning in the late-1990s.

The plot in Figure 5.2.3 Panel A displays the PhD production rate, in contrast to the total number of PhDs. Here, at the aggregate level, this is calculated by taking the overall number of PhDs produced divided by the overall number of BAs.¹³ The PhD production rates in the life and physical sciences largely track each other and are very similar in magnitude for much of the time period. These rates are also generally at least two times as large as the rates for the humanities and social sciences, which do not experience the same growth at the end of the sample.

Finally, Figure 5.2.3 Panel B shows that the trends in the PhD production rate among the various Carnegie categories are similar, but the levels differ substantially. The rate is highest for category 15 (Highest Research Activity), while lower rates for category 16 (Higher Research Activity), and again for category 17 (Moderate Research Activity).

Independent Variables

Table 4.2.3 displays summary statistics for the independent variables, specifically

¹³This is as opposed to taking the average of individually-computed PhD production rates.

means, standard deviations, and the percent of observations that are still missing after filling in and interpolating as described in Section 7.1. Means and standard deviations are conditional on the observation not being missing. The unit of observation is an institution-BA-year; thus, each institution contributes 26 observations.¹⁴ Observations are not weighted. Recalling that the set of institutions included in the samples of the different field varies somewhat due to the requirement that institutions produced at least five PhDs over the time frame, the sample used in this table is the life sciences sample, which contains the greatest number of institutions. For brevity, we do not present the summary statistics for the three other fields.

The 1st–3rd columns display statistics for the full sample, while the 4th–6th, 7th–9th, and 10th–12th columns display statistics for institutions with Carnegie category 15 (Highest Research Activity), 16 (Higher Research Activity), and 17 (Moderate Research Activity), respectively. After filling in and interpolating, most variables have a very small percentage of observations missing, exceptions being the 75th percentile score, and research expenditures per full-time faculty, both of which have a larger percentage missing across the Carnegie category gradient. Within field, institutions that do not have any department ranked in the NRC data are considered missing. There is a greater percentage of institutions without ranked NRC departments in the humanities and social sciences than in the life and physical sciences. Across all four fields, the percentage of category 15 (Highest Research Activity) institutions that do not have any departments ranked is small, while it is very large for category 17 institutions (Moderate Research Activity).

¹⁴Observations that do not contribute to regression results because e.g., they did not have any humanities BAs, are still included in the summary statistics.

As a whole, the mean values across Carnegie category show a clear pattern. Compared to category 16 (Higher Research Activity) and even more so to category 17 (Highest Research Activity) institutions, category 15 (Highest Research Activity) institutions have smaller student-faculty ratios, are more or equally often public institutions, and have higher institutional expenditures per student and research expenditures per full time faculty. They also have a smaller percentage of undergraduate students—meaning a higher percentage of graduate and professional students—and a smaller percentage of female and minority students. Carnegie category 15 (Highest Research Activity) institutions also have a higher average share of undergraduate majors in the humanities, life sciences, and physical sciences. They have a much smaller share in other majors, which include subjects such as business. Finally, they have many more departments ranked in the 1995 NRC doctoral rankings.

2.6 Results

Baseline Regression

Table 4.2.4 presents the baseline weighted least squares results, where each column is from a separate regression and corresponds to one of the PhD fields: humanities, life sciences, physical sciences, and social sciences. The mean value of the dependent variable, the PhD production rate, is displayed directly below the regression coefficients. Recall that this variable is a percentage, not a share. In line with Figure 5.2.3 Panel A, this value is much larger for the life and physical sciences. The explanatory power for these models is

fairly high—around 0.5—for all but the social sciences, where it is lower.

There are three variables that are consistently statistically significant across the fields: the 75th percentile test score, the number of total students at the institution, the percentage of total BAs that are in the given field, and the number of doctoral programs ranked in the top 10th percentile. The direction of each is also consistent across fields. The higher the test scores, the higher the PhD production rate. The coefficient is largest for the life sciences. Students with higher college entrance exams are more likely to excel in school and have lower costs of continuing to acquire education. That the total students variable has a negative coefficient means that smaller institutions produce more PhDs, everything else equal: having 1,000 fewer students is associated with a drop in the PhD production rate of between .04 for social sciences and .11 for physical sciences. The percentage of the total BAs that are in a given field is negatively related to the PhD production rate. Interestingly, this negative relationship goes against the hypothesis that a field having a larger presence is positively related to producing PhDs. One possible explanation is that those in smaller programs are less likely to get lost in the shuffle and are more likely to gain more individual attention from mentors such as faculty members. Having more highly-ranked doctoral programs, especially in the top 10th percentile, is strongly related to producing PhDs. For the humanities, an increase of one program in the top 10th percentile is associated with a higher PhD production rate of 0.44, which is 17% of the (weighted) mean. The coefficient is even higher for the life sciences (0.65; 13%) and the physical sciences (0.54; 12%). For the social sciences, it is 0.29 (13%). Top-ranked departments tend to have stellar faculty, which may influence students in several ways, including attracting them to the institution in the first place.

Several more variables are statistically significant for at least two of the four fields: instructional expenditures per student, the percentage of undergraduates who are female, the percentage of total BAs that majored in a particular field, and the institution being in Carnegie category 17 (Moderate Research Activity) as opposed to category 15 (Highest Research Activity). Instructional expenditure per student, measured in \$1,000s, is positively related for the humanities and physical sciences, but the magnitude is rather small. It may be the case that having higher quality faculty inspires students. Research expenditures per full time faculty is not related to the PhD production rate. For life and social sciences, the more men there are, the more likely students are to get a PhD. One factor that may be involved is the lower rate in general that women pursued PhDs, especially at the beginning of our sample timeframe. The percentage of total students who are undergraduates is never significant. The percentage of undergraduates who are a minority is positively related to the PhD production rate, but only for social sciences, and the coefficient is small.

Compared to Carnegie category 15 institutions, those in category 17 have a lower PhD production rate for both the physical and social sciences. This is after controlling for all other variables and is thus not a raw difference as in Figure 5.2.3 Panel B. Neither the student-faculty ratio nor being a public (versus private) institution relate to the PhD production rate.

In tables not shown, results are quantitatively unchanged if we use a maximum TTD based on the first year in which at least 80% or 85% of individuals receive their PhD.

Restricting to Carnegie 15 Institutions

In Table 4.2.5 everything is the same as above except that we now restrict the sample to Carnegie category 15 institutions (Highest Research Activity). The PhD production rates and explanatory of these regressions are both higher. The 75th percentile test score and total students variables are again statistically significant across all fields, with coefficients larger in magnitude compared to the baseline. The coefficients across fields for the number of doctoral programs ranked in the top 10th percentile are roughly of the same magnitude as those found in Table 4.2.4, although only humanities and physical sciences continue to be statistically significant.

With some exceptions, the same general findings of significance and sign hold for the other variables even though the standard errors have increases as the sample size has decreased.

Lemke-Style Regression

Table 4.2.6 shows the results from performing the analysis along the lines of Lemke (2009). Compared to the baseline regression, the findings are again largely the same. The test score and total student variables exhibit very similar patterns. The doctoral 0-10 percentile rankings variables are statistically significant for the same fields as when we restrict the sample to Carnegie category 15 institutions. Compared to the baseline regression, most of the same coefficients are significant statistically and have similar magnitudes.

2.7 Conclusion

By using the restricted access *Survey of Earned Doctorate* individual-level data, we are able to determine for each field each PhD recipient's time-to-degree and, given the heterogeneity in time-to-degree, to compute better estimates of the share of bachelor's graduates at an institution in a year that subsequently received PhDs in the field than did the early method used by Lemke (2009). Our method required us to address the truncation problem that occurs when using the individual level SED data, that we will not observe PhDs earned after the data end.

We conduct an weighted least squares analysis to examine which factors are associated with the PhD production rate. We point out that because we are not exploiting exogenous variation, our results should not be interpreted causally, as there are likely variables that we did not control for. With that said, our major empirical finding is the strong positive association between the number of highly-ranked PhD programs in a field at an institution in a year and the number of its undergraduate students that go on for PhDs in that field. Strong doctoral programs in a field thus seem to contribute to the supply of PhDs both through the number of PhDs they directly generate and through their impact on the number of their undergraduate students that go on for PhDs. We also find that higher entering test scores for undergraduate students and a lower percentage of BAs that are received in the field are both associated with more undergraduates at the institution going on for PhDs.

CHAPTER 3

INFORMATION AND THE BEAUTY PREMIUM IN POLITICAL ELECTIONS



INFORMATION AND THE BEAUTY PREMIUM IN POLITICAL ELECTIONS

TODD R. JONES and JOSEPH PRICE*

We use data on 800 candidates from the 2012 U.S. election cycle in U.S. and state congressional races to examine the degree to which beauty affects electoral outcomes. We find that a candidate that is one standard deviation more beautiful receives a 1.1 percentage point higher vote share and is 6.0 percentage points more likely to win the election. This beauty premium is larger in situations where voters are less likely to have more information about the candidate. The beauty premium is much smaller for U.S. congressional races than for state congressional races, and is also much smaller for incumbent candidates. In addition, we find a correlation that the beauty premium is lower when a candidate spends more money on the election. (JEL D72, J70)

I. INTRODUCTION

Individuals are influenced every day by a variety of biases. These biases affect a range of decisions and may be explicit, where the person chooses to discriminate, or implicit, where there is no conscious intent (Bertrand, Chugh, and Mullainathan 2005). Biases may lead to unfair treatment of minority groups and can result in suboptimal outcomes. Political elections are an important setting where biases based on personal characteristics including gender, ethnicity, and beauty may influence the outcome. The effect of small biases on election outcomes is driven largely by the fact that many voters are uninformed about the candidates or issues (Bartels 1996) forcing them to vote based on external cues and biases (McDermott 1997).

One particular bias in elections is a beauty bias where attractive candidates receive more votes. Past studies have documented a beauty premium with magnitudes differing based on the type of election and the degree to which voters are informed. Berggren, Jordahl, and

Poutvaara (2015) find that the beauty premium is significant for both national and low-profile elections in Finland, but that it matters more for particular candidates (right-leaning candidates) in particular elections (low-profile elections). Berggren, Jordahl, and Poutvaara (2010) finds that the beauty premium is larger for nonincumbent candidates.

Extending this analysis to the United States, we examine the degree to which the beauty premium is smaller when the voters are likely to be better informed about the candidates. First, we test whether the beauty premium differs between elections for U.S. congress, where there is considerably more media attention, and elections for the House and Senate of individual states.¹ Second, we test whether the beauty premium is larger for incumbent or nonincumbent candidates, because there is likely to be much less information about nonincumbent candidates. Third, we test whether the beauty premium is diminished in elections with increased campaign spending, since higher spending is likely to be associated with more information available to voters. We recognize that more electable candidates may be able to raise more funds, which thereby induces endogeneity.

*We would like to thank Carl Klarner and Gary Jacobson for providing data. We also thank Daniel Hamermesh, Levon Barseghyan, Daniel Benjamin, Cassandra Benson, Francine Blau, Lawrence Blume, Christine Coyer, Ronald Ehrenberg, and Lawrence Khan for providing helpful comments.

Jones: Ph.D. Candidate, Field of Economics, Cornell University, Ithaca, NY 14853. Phone (607) 255-4254, Fax (607) 255-2818. E-mail tj25@cornell.edu

Price: Associate Professor, Department of Economics, Brigham Young University, Provo, UT 84602. Phone 801-422-5296, Fax 801-422-0194, E-mail joseph_price@byu.edu

ABBREVIATIONS

IQ: Intelligence Quotient

RA: Research Assistant

Thus, the estimates regarding this issue should be interpreted as correlational rather than causal.

In all three cases, we find that the beauty gap is smallest in those settings when voters are most likely to be informed, suggesting that biases such as beauty bias are malleable to the environment in which the evaluation occurs. These results add to a growing literature about ways in which biases that affect how individuals evaluate others can be reduced (e.g., Marmaros and Sacerdote 2006; Parsons et al. 2011).

The paper is organized as follows. Section II discusses previous literature on the beauty premium in general and in politics specifically. Section III details the data and methodology. We present the results in Section IV and conclude in Section V.

II. BACKGROUND

Discrimination remains a large concern in labor economics. Beauty precipitates discrimination against individuals perceived as less attractive. The seminal paper of Hamermesh and Biddle (1994) was among the first to document a wage premium for individuals who are more attractive. Subsequently, there have been many studies which explore the beauty premium in a variety of contexts, including education and professional sports (Berri et al. 2011; Hamermesh and Parker 2005). Deryugina and Shurchkov (2015) study female college students and find no relationship between college grades and attractiveness, but do find substantial beauty-based sorting into areas of study. Ravina (2012) finds that beautiful applicants have a 1.59% higher probability of getting loans and pay 60 basis points less, despite having no difference in default rates than average looking people. The income premium of attractiveness in high school is found to persist through a person's early-50s, being robust to intelligence quotient (IQ), proxies for confidence and personality, family background, and high school experiences (Scholz and Sicinski 2015). In a laboratory setting, Wilson and Eckel (2006) find that attractiveness initially makes individuals more trusted by strangers, however they later face a beauty penalty by failing to live up to expectations.

Beauty and appearance have widely been examined in the context of elections, with significant consequences on both the election outcome and vote share. Hamermesh (2006) finds that beauty predicts a significant amount of variation in American Economic Association

elections. King and Leigh (2009) and Berggren, Jordahl, and Poutvaara (2010) find evidence of the beauty premium in Australia and Finland, with an increase in beauty associated with a significant increase in vote share. Berggren, Jordahl, and Poutvaara (2015) observe that in Finland, right-wing candidates look better than left-wing candidates, and candidates on the right experience a larger beauty premium relative to those on the left, but only in low-profile elections. Todorov et al. (2005) find that judgments of competency based on 1-second exposures to photos of two U.S. Congressional candidates predict nearly 70% of elections. Similar to beauty, Benjamin and Shapiro (2009) determine that voters make election forecasts based on candidates' personal characteristics, such as likeability, rather than inferences concerning candidates' policy positions.

We expect the beauty premium to differ between high- and low-profile elections due to varying levels of voter knowledge. Specifically, we expect the beauty premium to be lower for elections where voters are more highly informed about candidates. Recent literature corroborates these expectations. Riggle et al. (1992) show that absent all other information about candidates, attractiveness significantly influences voters' choices; however, when given other information, beauty becomes less important and possibly insignificant. Johns and Shephard (2011) additionally find that photographs of politicians have higher effects on uninformed voters than on informed voters. Todorov et al. (2005) posit that first impressions, including beauty impressions, may have a significant and lasting effect on voter choices.

Voters may choose to obtain more or less information based on their judgment of the costs and benefits of obtaining the information. For example, the costs of obtaining information are often lower in national elections, with a larger amount of advertisement in the media and a higher level of news coverage. Benefits may also differ; for example, a voter may derive more utility by becoming informed about national elections and being able to express political viewpoints to peer voters who may place a higher value on national elections.

III. DATA AND METHODOLOGY

We collected subjective data on perceived beauty of 800 different candidates from 400 randomly selected elections that took place in

2012. We randomly sampled 200 elections for the U.S. Senate and U.S. House as well as 200 elections for the State Senate and State House and focus on just the Republican and Democrat candidate for each election. The photos for each candidate come from votesmart.org, a nonpartisan, nonprofit political website. For the 2012 election year, this site contains a photo for over 1,200 candidates from 470 high-profile elections and a photo of over 8,000 candidates from over 5,600 low-profile elections. Votesmart uses the official picture from candidate campaign websites, when available. Other photographs are found via candidates' social media pages, such as Facebook or Twitter. Pictures are cropped in a fixed 1:1 ratio and generally include each candidate's head and shoulders only. The size of the photos used in the study is small (110×135 pixels), but uniform across photos. Our sample uses only color photographs of candidates and excludes elections where either the Democrat or Republican candidates has only a black and white photo on votesmart.org.

There are two potential problems with estimating a beauty premium using candidate photos. A reverse causality problem arises if more successful candidates take better, more professional photos. In a robustness check, we attempt to control for this effect by using two subjective measures of photo professionalism. These are binary indicators for whether the photo appears to be professionally produced and an indicator for the quality of the photo. We coded the picture as being low quality if the picture appeared more pixelated and we coded whether the picture was professionally done based on the dress, style, and background. We find that 741 of the 800 photos were high quality and 619 professionally done. High-profile elections were more likely to have high-quality photos but less likely to have photos that appeared to be professionally done. We include these additional controls as part of a robustness check in our analysis. In addition, we hired three research assistants (RAs) to code the professionalism and quality of each photo on a 1–3 scale in order to give us an alternate measure of photo professionalism and quality not created by the authors. We took the average (mean) and mode of the three ratings (using 2 for the mode when all three ratings differed).

In the literature regarding the beauty premium, age is often controlled for as it is a confounding factor which is likely correlated with beauty as well as vote share and winning an election. In another robustness check, we estimate the

main regressions using controls for age and an interaction term of age with beauty. We were able to identify the ages for 387 high-profile and 213 low-profile candidates and include these 600 observations in a robustness check with an age control.

Second, there may be omitted variable bias if beauty is correlated with other things that could lead to votes, such as oral skills, competence, or the number of people that campaigning candidates meet. We do not have a way to measure and/or control for these other factors; therefore, results should be interpreted accordingly. Finally, Atkinson, Enos, and Hill (2009) suggest that a bias exists when a political party chooses a more beautiful candidate for an election it expects to be close; we do not address this particular issue in this paper.

We use an online survey with 991 participants from Amazon Mechanical Turk to obtain subjective ratings of beauty for the 800 candidate photos. Workers were required to be at least 18 years old, in the United States, and to have a project acceptance percentage of 90% or greater.² Each participant rated the attractiveness of 20 photographs that were selected and ordered randomly from the sample of 800 candidates.³ Answers to the beauty question are coded on the same 1–5 scale used by Berggren, Jordahl, and Poutvaara (2015).⁴ Participants also indicated whether they recognized the persons they rated; we exclude recognized photo observations from our analysis. Participants were not informed that the photos were of political candidates in order to prevent qualities that may be associated with political ability, but not with beauty, from interfering with participants' evaluations. On average, each photo received about 25 ratings; the photo with the least number of ratings received 11 ratings. We take out rater fixed effects by regressing all scores on photo fixed effects and

2. Ipeirotis (2010) provides some demographic information about Mechanical Turk workers and finds that 65% are female, 75% were born prior to 1985, almost 50% have at least a Bachelor's degree, and the median income is approximately \$50,000.

3. The wording of the question and possible answers to the beauty question are similar to that used by Berggren, Jordahl, and Poutvaara (2015). Our survey also asked the respondent to rate the competence of the person in the photo; we do not consider competence in this analysis.

4. To provide a measure of inter-rater agreement among the survey-takers, we group ratings of 1 and 2 into a low category and 4 and 5 into a high category. We calculate a Kappa coefficient between these two groups of 0.523, a highly significant result that parallels that found by Berggren, Jordahl, and Poutvaara (2010).

TABLE 1
Mean Beauty by Candidate Characteristics

	Low-Profile Election	High-Profile Election
<i>Gender</i>		
Male	2.287 (0.503)	2.396 (0.481)
Female	2.556 (0.600)	2.619 (0.556)
<i>Inc incumbency Status</i>		
Not incumbent	2.377 (0.539)	2.452 (0.515)
Incumbent	2.301 (0.541)	2.428 (0.492)
<i>Candidate Result</i>		
Lost	2.310 (0.522)	2.443 (0.530)
Won	2.396 (0.556)	2.440 (0.480)
<i>Party</i>		
Democrat	2.321 (0.553)	2.375 (0.490)
Republican	2.385 (0.527)	2.508 (0.512)
<i>Race</i>		
White	2.343 (0.540)	2.431 (0.511)
Not White	2.465 (0.534)	2.514 (0.457)
<i>N</i>	400	400

Notes: All of the statistics are calculated with Democrats and Republicans together with the exception of results by party, which are broken up by Democrat and Republican.

rater fixed effects and use the coefficients on the photo fixed effects as the beauty score, which are then normalize to have the same mean and standard deviation as the original data. Table 1 provides the average beauty rating across a set of different candidate characteristics. Female politicians were rated considerably higher than men in terms of beauty. Low-profile candidates who won had higher beauty scores on average, but there was no difference in high-profile elections. Surprisingly, incumbents had a slightly lower beauty score, perhaps due to age differences.

The data on the elections come from several sources. We use Carl Klarner's "State Legislative Election Returns Data, 2011–2012" for the state-level elections. These data are similar to that found in Klarner et al. (2013), but has been updated to include more recent years. U.S. House election data are collected from Gary Jacobson; U.S. Senate elections data were collated using data from Carl Klarner and NBC.com. OpenSecrets.org, run by the Center for Responsive Politics, supplies data on campaign expenditures.

IV. RESULTS

For our analysis, the unit of observation is a candidate. Standard errors are clustered at the

election level, with two observations per election. Our two primary dependent variables are the fraction of the vote that the candidate received (among those votes cast for one of the two major parties) and an indicator for whether or not the candidate won. Logistic regression reporting marginal effects at the mean is used when the outcome variable is winning and ordinary least squares is used for regressions with vote share as the outcome variable. We include a basic set of controls in each regression that include the candidate's incumbency status, gender, and race/ethnicity. Beauty has been standardized over all 800 candidates. As a set of robustness checks we include additional controls for the photo quality and photo professionalism.

In Table 2, we provide the results of our analysis when we pool the data from the state and national elections together. The results in this table indicate that a one standard deviation increase in a candidate's beauty is associated with a 1.1 percentage point increase in the fraction of votes received and a 6.0 percentage point increase in the probability of winning the election.⁵ In columns 2 and 4, we include an interaction term between beauty and high-profile election and another interaction between beauty and incumbency status. The coefficient for the measure of beauty now represents the premium for nonincumbents in low-profile elections. For these candidates, we find that a one standard deviation increase in beauty is associated with a 2.5 percentage point increase in the vote share and a 10.2 percentage point increase in the probability of winning (both statistically significant at the 1% level). In contrast, the interaction term between beauty and high-profile election is -1.4 percentage points for vote share and -6.3 percentage points for winning (not significant), indicating that the beauty premium is much smaller for high-profile elections. The interaction term between beauty and incumbency status is also negative, with a coefficient of -2.2 percentage points for vote share and -9.0 percentage points in terms of winning the election, although this latter number is not significant. This result matches those of Berggren, Jordahl, and Poutvaara (2010), who also found that the effects of beauty were concentrated among nonincumbent candidates.

We further test the relationship between the beauty premium and voter information by

5. We note that in regressions excluding incumbency status, the high-profile coefficient is close to 0; including incumbent causes the coefficient to increase in magnitude.

TABLE 2
Correlation between Beauty and Election
Outcomes

	Won		Vote Share	
	(1)	(2)	(3)	(4)
Beauty	0.060*** (0.023)	0.102*** (0.033)	0.011*** (0.004)	0.025** (0.006)
Incumbent*Beauty		-0.090 (0.065)		-0.022** (0.007)
High-profile*Beauty		-0.063 (0.045)		-0.014* (0.008)
Incumbent	0.692*** (0.028)	0.691*** (0.028)	0.195*** (0.009)	0.195*** (0.009)
High profile	-0.144*** (0.031)	-0.133*** (0.031)	-0.024*** (0.005)	-0.024** (0.005)
Female	-0.023 (0.061)	-0.022 (0.060)	0.010 (0.009)	0.010 (0.009)
Black	-0.035 (0.090)	-0.038 (0.090)	0.027 (0.022)	0.026 (0.022)
Hispanic	0.079 (0.111)	0.085 (0.109)	0.014 (0.023)	0.017 (0.023)
Other race/ Ethnicity	-0.292*** (0.065)	-0.263*** (0.083)	-0.017 (0.021)	-0.010 (0.018)
N	800	800	800	800
R-squared	.358	.362	.437	.445

Notes: Vote share is based on the two major party candidates. Mean beauty has been normalized across all observations by subtracting the mean and dividing by the standard deviation. We use logistic regression when *winner* is the outcome and report the marginal effects at the mean. Robust standard errors in parentheses.

*** $p < .01$, ** $p < .05$, * $p < .1$.

considering the expenditures of a candidate during an election. Higher spending is correlated with higher voter information (Potters, Sloof, and Van Winden 1997). We thus expect elections with higher spending to have a smaller beauty premium, as voters will be better informed. In Table 3, we provide regression results, which include an interaction term between beauty and expenditures by the candidate in that election. We divide expenditures by the total number of votes cast in the election and then standardize the variable separately for low-profile and high-profile elections.⁶ We consider national and state elections separately because spending and high-profile status are so strongly correlated.

In the first column of Table 3, we find that there is no beauty premium, on average, in elections for the U.S. House and Senate elections. However, in the fourth column when we include the interaction with spending, we do find a negative and statistically significant interaction term. The coefficients on the main effect for beauty and the interaction term suggest that for each standard deviation a candidate is above the beauty mean the candidate loses a beauty premium of

6. When we do not divide by the votes, the spending-beauty interaction terms at the national level are not significant.

1.6 percentage points in vote share for every standard deviation they spend above the sample mean. Generally, the positive 3.3% direct effect of spending outweighs the added beauty premium but it does leave the possibility that spending more could outweigh the beauty premium for candidates more than two standard deviations below the sample mean. For the state-level elections, we also find a negative interaction term between beauty and campaign spending that is slightly smaller and not statistically significant. We recognize that the funding a candidate receives may be correlated with traits that attract votes. Owing to this endogeneity the relationship cannot be interpreted as causal. Campaign spending is often viewed as being an expense that is socially wasteful. Our results suggest that increased campaign spending may be socially beneficial by reducing biases that affect how individuals vote.

As robustness checks, in Table 4, we re-estimate the specification of Table 2 and include subjective indicator variables for photo quality and for photo professionalism. We also estimate specifications with a control for candidate's age and an interaction of age with beauty. In the first two columns of Table 4 we find that the estimated beauty coefficient is about 25% smaller when we include controls for photo quality and professionalism. This suggests that a part of the beauty premium may operate through a candidate's presentation of themselves to the public. The final two columns of Table 4 include age controls; the marginal effect for beauty on winning is slightly smaller as compared to that in Table 2, but is slightly larger when vote share is the outcome variable. We also find that including these quality and professionalism controls reduces the magnitude of the interaction terms in Table 2. These changes in results are observed when professionalism and quality are included (shown using the mean of RA coding), but are not sensitive to how these variables were coded: the author rating, mean of RA coding, or mode of RA coding.

In terms of the direct effects of these other two measures, we find that photo quality is positively related to election outcomes, but not significant in any specification, while photo professionalism is positive, significant, and quite large. These results indicate that candidates who take more professionally-appearing photos achieve a higher vote share. This may mean that the measured beauty premium may, to some degree, be a side effect of better candidates—perhaps better candidates take more professional photos

TABLE 3
Difference in Beauty Premium Based on Election Expenditures

	National				State			
	Winner		Vote Share		Winner		Vote Share	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Beauty	-0.007 (0.037)	-0.011 (0.042)	-0.001 (0.005)	-0.004 (0.005)	0.059 * (0.031)	0.052 (0.034)	0.012 ** (0.006)	0.012 ** (0.006)
Expenditures	0.279*** (0.059)	0.291*** (0.069)	0.029*** (0.009)	0.033*** (0.008)	0.132 ** (0.060)	0.161 ** (0.071)	0.018*** (0.007)	0.022*** (0.008)
Expend*Beauty		-0.061 (0.077)		-0.016 * (0.008)		-0.063 (0.052)		-0.009 (0.006)
Observations	399	399	400	400	400	400	400	400
R-squared	.571	.573	0.578	0.584	.259	.264	0.363	0.368

Notes: Vote share is based solely on the two major party candidates. *Expenditures* was divided by total votes for the two major parties, then normalized separately for national and state elections. Observations are 399 in National with *winner* as the outcome because *other race/ethnicity* predicts the outcome perfectly. Mean beauty has been normalized across all 800 observations by subtracting the mean and dividing by the standard deviation. Each regression includes controls for incumbency, race/ethnicity, and gender. We use logistic regression when "winner" is the outcome and report the marginal effects at the mean. Robust standard errors in parentheses.

*** $p < .01$, ** $p < .05$, * $p < .1$.

TABLE 4
Correlation between Beauty and Election
Outcomes Accounting for Additional Controls

	Winner	Vote Share	Winner	Vote Share
Beauty	0.078 ** (0.034)	0.020*** (0.006)	0.090 ** (0.044)	0.036*** (0.009)
Incumbent*Beauty	-0.084 (0.066)	-0.019*** (0.007)	-0.066 (0.063)	-0.026*** (0.008)
High profile*Beauty	-0.050 (0.047)	-0.012 (0.008)	-0.061 (0.052)	-0.019 ** (0.009)
Incumbent	0.658*** (0.032)	0.172*** (0.010)	0.660*** (0.032)	0.186*** (0.010)
High profile	-0.124*** (0.033)	-0.022*** (0.005)	-0.205*** (0.033)	-0.029*** (0.007)
Photo professionalism	0.289*** (0.052)	0.049*** (0.008)		
Photo quality	0.010 (0.052)	0.005 (0.007)		
Age			-0.024 (0.028)	0.003 (0.005)
Age*Beauty			0.026 (0.025)	0.000 (0.005)
Observations	800	800	600	600
R-squared	.395	0.474	.373	0.444

Notes: Vote share is based on the two major party candidates. Mean beauty has been normalized across all observations by subtracting the mean and dividing by the standard deviation. We use logistic regression when *winner* is the outcome and report the marginal effects at the mean. *Age* is in standard deviation units. *Photo quality* and *professionalism* use the average of RA coding. Robust standard errors in parentheses.

*** $p < .01$, ** $p < .05$, * $p < .1$.

or respondents to the survey rated professional-looking photos higher because they perceived the candidate to be more beautiful due to the professionalism of the photo.⁷

7. Across all observations, the correlation of standardized beauty and photo professionalism is 0.11. Incumbent status is also strongly predictive of photo professionalism.

V. CONCLUSION

Corroborating current literature, we find that beauty matters in elections. These results are robust to inclusion of controls for gender, race/ethnicity, age, and incumbency status. We find that candidates with higher levels of beauty experience electoral success. This effect is particularly large for nonincumbents in low-profile elections, a situation in which the voters are likely to have the least amount of information in the candidate. In contrast, we find no beauty premium for candidates in high-profile elections. This implies that voters use beauty as a substitute for information about candidates' ability. This finding may provide some amount of justification for the large amounts of resources invested in election campaigns as they may help eliminate the beauty bias of voters, again keeping in mind the endogeneity of the spending variable. In general, increased access to information can play a role in reducing various biases such as those based on beauty, gender, or race.

If beauty causes less-qualified or less-capable candidates to be elected due to attractiveness, policymakers who desire to prevent such outcomes will be motivated to explore ways of making beauty less of a factor in elections. It may be that beauty accurately signals to voters about the competency of a candidate and that the beauty premium may not be detrimental. Mobiüs and Rosenblat (2006) find that beauty is correlated with confidence and attractive workers have greater oral and social skills; however, they also find that for a given confidence level, physically

attractive workers are (wrongly) considered more able by employers. We recognize our results depend on beauty not signaling competence or ability, and acknowledge this is an assumption that is difficult to empirically test. As Todorov et al. (2005) suggest, the beauty premium may be a symptom of a deeper inclination voters possess to vote according to their initial impression of a candidate, regardless of how this first impression may form. This analysis suggests that potential political candidates considering an investment of time and money to run for office will also be better able to estimate their chances for success by accounting for beauty effects. Finally, voters equipped with knowledge of the beauty premium may guard against its detrimental effects by increasing their knowledge about political candidates in order to cast less biased votes.

REFERENCES

- Atkinson, M. D., R. D. Enos, and S. J. Hill. "Candidate Faces and Election Outcomes: Is the Face-Vote Correlation Caused by Candidate Selection?" *International Quarterly Journal of Political Science*, 4(3), 2009, 229–49.
- Bartels, L. M. "Uninformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science*, 40(1), 1996, 194–230.
- Benjamin, D. J., and J. M. Shapiro. "Thin-Slice Forecasts of Gubernatorial Elections." *Review of Economics and Statistics*, 91(3), 2009, 523–36.
- Berggren, N., H. Jordahl, and P. Poutvaara. "The Looks of a Winner: Beauty and Electoral Success." *Journal of Public Economics*, 94(1–2), 2010, 8–15.
- , "The Right Look: Conservative Politicians Look Better and Voters Reward It." Munich Discussion Papers in Economics, 2015–10, 2015.
- Berri, D. J., R. Simmons, J. Van Gilder, and L. O'Neill. "What Does It Mean to Find the Face of the Franchise? Physical Attractiveness and the Evaluation of Athletic Performance." *Economics Letters*, 111(3), 2011, 200–2.
- Bertrand, M., D. Chugh, and S. Mullainathan. "Implicit Discrimination." *American Economic Review*, 95(2), 2005, 94–98.
- Deryugina, T., and O. Shurchkov. "Does Beauty Matter in Undergraduate Education?" *Economic Inquiry*, 53(2), 2015, 940–61.
- Hamermesh, D. S. "Changing Looks and Changing 'Discrimination': The Beauty of Economists." *Economics Letters*, 93(3), 2006, 405–12.
- Hamermesh, D. S., and J. E. Biddle. "Beauty and the Labor Market." *American Economic Review*, 84(5), 1994, 1174–94.
- Hamermesh, D. S., and A. Parker. "Beauty in the Classroom: Instructors' Pulchritude and Putative Pedagogical Productivity." *Economics of Education Review*, 24(4), 2005, 369–76.
- Ipeirotis, P. G. "Demographics of Mechanical Turk." Working Paper CeDER-10-01, NYU Center for Digital Economy Research, 2010.
- Johns, R., and M. Shephard. "Facing the Voters: The Potential Impact of Ballot Paper Photographs in British Elections." *Political Studies*, 59(3), 2011, 636–58.
- King, A., and A. Leigh. "Beautiful Politicians." *Kyklos*, 62(4), 2009, 579–93.
- Klarner, C., W. Berry, T. Carsey, M. Jewell, R. Niemi, L. Powell, and J. Snyder. "State Legislative Election Returns Data, 1967–2010," hdl:1902.1/20401, Harvard Data-verse, V1, UNF:5:q/n5C9RQGfjy6AjbLG6JWQ==;
- Marmaros, D., and B. Sacerdote. "How Do Friendships Form?" *Quarterly Journal of Economics*, 121(1), 2006, 79–119.
- McDermott, M. L. "Voting Cues in Low-Information Elections: Candidate Gender as a Social Information Variable in Contemporary United States Elections." *American Journal of Political Science*, 41(1), 1997, 270–83.
- Mobius, M., and T. Rosenblat. "Why Beauty Matters." *American Economic Review*, 96(1), 2006, 222–35.
- Parsons, C. A., J. Sulacman, M. C. Yates, and D. S. Hamermesh. "Strike Three: Discrimination, Incentives, and Evaluation." *American Economic Review*, 101(4), 2011, 1410–35.
- Potters, J., R. Sloof, and F. Van Winden. "Campaign Expenditures, Contributions and Direct Endorsements: The Strategic Use of Information and Money to Influence Voter Behavior." *European Journal of Political Economy*, 13(1), 1997, 1–31.
- Ravina, E. "Love & Loans: The Effect of Beauty and Personal Characteristics in Credit Markets." Working Paper, 2012, http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1107307
- Riggle, E., V. Ottati, R. Wyer, J. Kuklinski, and N. Schwarz. "Bases of Political Judgments: The Role of Stereotypic and Nonstereotypic Information." *Political Behavior*, 14(1), 1992, 67–87.
- Scholz, J. K., and K. Sicinski. "Facial Attractiveness and Lifetime Earnings: Evidence from a Cohort Study." *Review of Economics and Statistics*, 97(1), 2015, 14–28.
- Todorov, A., A. N. Mandisodza, A. Goren, and C. C. Hall. "Inferences of Competence from Faces Predict Election Outcomes." *Science*, 308(5728), 2005, 1623–6.
- Wilson, R. K., and C. C. Eckel. "Judging a Book by Its Cover: Beauty and Expectations in a Trust Game." *Political Research Quarterly*, 59(2), 2006, 189–202.

CHAPTER 4

TABLES

4.1 Tables for Do Peers Influence Occupational Preferences?

Table 4.1.1: Branches Available to Cadets

Combat Arms	Combat Support	Combat Service Support
Air Defense Artillery (AD)	Chemical Corps (CM)	Adjutant General (AG)
Armor (AR)*	Military Intelligence (MI)	Finance (FI)
Aviation (AV)	Military Police (MP)	Medical Services (MS)
Engineer (EN)	Signal (SC)	Ordnance (OD)
Field Artillery (FA)		Quartermaster (QM)
Infantry (IN)*		Transportation (TC)

Notes: This table shows all branches available to cadets after graduation, broken into three branch categories: Combat Arms, Combat Support, and Combat Service Support. Initials, which are used elsewhere in the paper, follow the branch name. * indicates a branch that the Department of Defense did not allow women to join at the beginning of the period. Cyber Corps (CY) was added as a branch option during the sample timeframe.

Table 4.1.2: Summary Statistics

	Mean	SD	Min	Max
Black	0.08	0.27	0	1
Hispanic	0.09	0.29	0	1
Other Race	0.09	0.29	0	1
CEER Score	619.45	70.17	417	790
Whole Candidate Score	6208.19	491.46	4714	7335
NCAA Athlete	0.15	0.36	0	1
Prior Enlisted	0.13	0.34	0	1
Army Parent	0.15	0.35	0	1
Prep School	0.15	0.35	0	1
Observations	1431			

Notes: This table shows summary statistics. *Black*, *Hispanic*, and *Other Race* are indicator variables for race/ethnicity, and are mutually exclusive. *CEER Score* is the Cadet Entrance Exam Rank, a composite score used to rank cadets academically for admissions, and is comprised of SAT Math, SAT Verbal, and high school rank. *Whole Candidate Score* is also used during admissions, and includes assessments of leadership potential and physical fitness. *NCAA Athlete* indicates that the cadet was recruited as an NCAA athlete. *Prior Enlisted* denotes prior service as a soldier. *Army Parent* and *Prep School* are indicator variables of whether at least one of the cadet's parents was also in the Army, and whether the cadet attended West Point's preparatory academy, respectively.

Table 4.1.3: Top Branch Selections, Rounds 1, 2, and 6, with Percentage Changes

Branch	Round 1		Round 2			Round 6			
	Freq.	Pct.	Freq.	Pct.	%Δ(1-2)	Freq.	Pct.	%Δ(1-6)	%Δ(2-6)
Infantry	440	30.75	322	22.50	-8.25	356	24.88	-5.87	2.38
Aviation	285	19.92	343	23.97	4.05	204	14.26	-5.66	-9.71
Military Intelligence	181	12.65	132	9.22	-3.42	146	10.20	-2.45	0.98
Engineer	143	9.99	162	11.32	1.33	192	13.42	3.42	2.10
Armor	113	7.90	117	8.18	0.28	121	8.46	0.56	0.28
Field Artillery	63	4.40	82	5.73	1.33	70	4.89	0.49	-0.84
Medical Services	49	3.42	40	2.80	-0.63	43	3.00	-0.42	0.21
Air Defense Artillery	45	3.14	52	3.63	0.49	66	4.61	1.47	0.98
Finance	30	2.10	27	1.89	-0.21	25	1.75	-0.35	-0.14
Quartermaster	21	1.47	45	3.14	1.68	40	2.80	1.33	-0.35
Signal Corps	19	1.33	41	2.87	1.54	42	2.94	1.61	0.07
Military Police	16	1.12	24	1.68	0.56	20	1.40	0.28	-0.28
Adjutant General	11	0.77	20	1.40	0.63	14	0.98	0.21	-0.42
Ordnance	6	0.42	14	0.98	0.56	36	2.52	2.10	1.54
Transportation	6	0.42	8	0.56	0.14	22	1.54	1.12	0.98
Chemical Corps	3	0.21	2	0.14	-0.07	4	0.28	0.07	0.14
Cyber	0	0.00	0	0.00	0.00	30	2.10	2.10	2.10
Total	1431	.	1431	.	.	1431	.	.	.

Notes: This table shows the distribution of the top-ranked branch preferences for Rounds 1, 2, and 6. The *Freq.* and *Pct.* columns refer to the frequency and percentage of cadets who ranked the branch as their top choice during the corresponding round. The $\%Δ(1-2)$ column is the percentage change in the branch from Round 1 to Round 2, with analogous $\%Δ(1-6)$ and $\%Δ(2-6)$ columns.

Table 4.1.4: Top Branch Transitions between Rounds 1 and 6

	Round 6												Round 1					
	IN	AV	MI	EN	AR	FA	MS	AD	FI	QM	SC	MP	AG	OD	TC	CM	CY	Total
IN	239	40	35	29	36	13	6	6	2	4	5	4	4	6	6	1	4	440
AV	33	110	21	33	23	13	9	16	4	1	5	6	1	7	1	0	2	285
MI	22	12	52	12	9	9	4	9	4	10	8	2	1	7	4	1	15	181
EN	18	19	8	58	7	4	1	10	2	4	6	0	0	3	1	0	2	143
AR	18	6	13	20	31	8	0	5	1	2	4	2	0	2	1	0	0	113
FA	15	2	1	8	7	13	2	2	1	3	1	2	1	4	1	0	0	63
MS	4	4	5	11	3	2	15	1	0	2	0	0	0	1	0	0	1	49
AD	1	2	4	7	2	3	0	11	2	4	3	0	1	1	3	1	0	45
FI	1	1	2	3	1	2	2	1	6	4	1	0	3	2	1	0	0	30
QM	2	1	2	4	0	1	0	2	2	5	0	0	0	2	0	0	0	21
SC	1	1	0	2	0	0	0	1	0	1	7	0	0	0	0	0	6	19
MP	0	2	1	1	1	0	2	0	0	0	2	4	1	1	1	0	0	16
AG	2	2	1	0	0	1	1	2	0	0	0	0	1	0	1	0	0	11
OD	0	1	0	3	0	1	0	0	0	0	0	0	1	0	0	0	0	6
TC	0	1	1	0	1	0	0	0	1	0	0	0	0	0	2	0	0	6
CM	0	0	0	1	0	0	1	0	0	0	0	0	0	0	0	1	0	3
Total	356	204	146	192	121	70	43	66	25	40	42	20	14	36	22	4	30	1431

Notes: This table shows the frequencies of transitions of top preferences between Rounds 1 and 6 for all branches. The vertical axis corresponds to the branch selected first in Round 1, and the horizontal axis is the same for Round 6. The top-left cell, with value 239, indicates that 239 cadets ranked Infantry as their top preference during Round 1, and did so during Round 6. The cell directly to the right, 40, indicates that there were 40 cadets who ranked Infantry first during Round 1, and Aviation during Round 6.

Table 4.1.5: Percent Choosing Same Top Ranked Branch in Later Rounds

Former Round	Latter Round				
	2	3	4	5	6
1	56.18	45.91	41.09	38.99	38.78
2		60.87	52.13	48.71	47.59
3			71.98	64.29	61.50
4				69.88	66.74
5					77.99

Notes: This table shows the percentage of cadets who ranked the same branch as their most-preferred branch across rounds. The rounds on the vertical (horizontal) axis are the former (latter) round of the calculation. For example, 38.78% of cadets had the same top-ranked branch during Round 6 as they did during Round 1.

Table 4.1.6: Average Ranking of First Branch in Later Rounds

Panel A: First Preference					
	Latter Round				
Former Round	2	3	4	5	6
1	2.90	3.31	3.59	3.79	3.81
2		2.21	2.50	2.82	3.05
3			1.71	2.05	2.26
4				1.64	1.78
5					1.40

Panel B: Second Preference					
	Latter Round				
Former Round	2	3	4	5	6
1	4.00	4.80	5.06	5.28	5.34
2		3.42	3.92	4.45	4.70
3			2.97	3.68	3.81
4				3.27	3.44
5					2.56

Notes: Panel A (B) shows the average slot the first-ranked (second-ranked) branch in the former round was ranked during the latter round. The initial (final) round used in the calculation is the round on the vertical (horizontal) axis. A value of 1 means that all cadets also ranked the branch they ranked first during the former round first the latter round. Sample only includes the 2015 cohort as they had only 16 choices every round because Cyber was never available. Calculations involving rounds 3, 4, and 5 have fewer observations due to missing preference data.

Table 4.1.7: Percent of Realized Branches that Were Ranked as 1st, 2nd, etc.

Preference	Count	Pct.	Cum. Pct.
1	1118	78.9	78.9
2	155	10.9	89.8
3	55	3.88	93.7
4	23	1.62	95.3
5	13	0.92	96.3
6	14	0.99	97.2
7	8	0.56	97.8
8	5	0.35	98.2
10	2	0.14	98.3
11	6	0.42	98.7
12	3	0.21	98.9
13	4	0.28	99.2
14	3	0.21	99.4
15	3	0.21	99.6
16	5	0.35	100
Total	1417	100	

Notes: This table uses branch preferences from the sixth and final preference round, and shows the percentage of cadets whose branch to which they were assigned to be an Army officer after graduation was the one they ranked 1st, 2nd, 3rd... For example, the percentage for ranking 3 means that 3.88% of cadets were assigned to the branch they ranked third. No cadets were assigned their 9th or 17 preferences. Sample is less than 1431 due to one missing assignment observation and 13 cadets who were assigned Cyber despite not ranking it.

Table 4.1.8: Empirical P-value Check for Random Assignment to Peers

	(1) Combat Arms	(2) Combat Support	(3) Support Services
<i>A. Tests for Random Assignment of Cadets to First Company</i>			
Empirical p-values (mean and standard deviation)	0.428 (0.292)	0.406 (0.296)	0.380 (0.303)
Kolmogorov-Smirnov test (no. failed/ total tests)	0/2	0/2	0/2
χ^2 goodness of fit test (no. failed/ total tests)	0/2	0/2	1/2
<i>B. Tests for Random Assignment of Cadets to Second Company</i>			
Empirical p-values (mean and standard deviation)	0.423 (0.294)	0.415 (0.307)	0.397 (0.324)
Kolmogorov-Smirnov test (no. failed/ total tests)	0/2	0/2	0/2
χ^2 goodness of fit test (no. failed/ total tests)	0/2	0/2	1/2

Notes: This table reports the results from an empirical p-value test for random assignment of branch with respect to preferences. The dependent variable is the empirical p-value. Each column is a test for a different grouping of branches: combat arms, combat support, or support services. The mean and standard deviation of the empirical p-values are reported, followed by Kolmogorov-Smirnov and χ^2 tests, which test that the empirical p-value is distributed uniformly. Panel A tests for random assignment to the first company, while Panel B tests for assignment to second company. Section 1.6 describes further details.

Table 4.1.9: Second Company Same-Cohort Peer Preferences on Third Round Ranking

	(1) First Pref, R3	(2) First Pref, R3	(3) First Pref, R3	(4) First Pref, R3
% of 2nd Company Peers, R1	0.3830*** (0.0179)	0.3830*** (0.0179)	-0.0637 (0.0524)	-0.0472 (0.0521)
Own First Pref, R1	0.3801*** (0.0128)	0.3801*** (0.0128)	0.3685*** (0.0135)	0.3686*** (0.0134)
Controls	No	Yes	Yes	Yes
Branch FE	No	No	Yes	No
Branch \times Year FE	No	No	No	Yes
Observations	24327	24327	24327	24327
R^2	0.200	0.200	0.209	0.211
Adjusted R^2	0.199	0.199	0.207	0.208

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates of a stacked regression of third round top preferences on average preferences of a cadet's second company peers and a cadet's own initial preference. These last two objects were measured during the first round, before assignment to the second company. The data were constructed such that each cadet reported 17 observations, one for each branch. Each observation contains an indicator for whether the cadet chose the branch as top choice during the third round (*First Pref, R3*), the same for the first round (*Own First Pref, R1*), and the leave-out fraction of peers in the second company that chose that branch as top choice (*% of 2nd Company Peers, R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. *Branch FE* means that an indicator variable for the branch to which the observation refers is included. *Branch x Year FE* is a fixed effect for Branch times Year. Cyber was an option during the third round for one cohort, but it was unavailable as a choice during the first two rounds; it is included in this table. Standard errors are clustered at the company-level.

Table 4.1.10: Second Company Same-Cohort Peer Preferences on 2nd–6th Round Preferences

	(1) 1st Pref, R2	(2) 1st Pref, R3	(3) 1st Pref, R4	(4) 1st Pref, R5	(5) 1st Pref, R6
% of 2nd Cpy. Peers, R1	0.0446 (0.0466)	-0.0637 (0.0524)	-0.1446** (0.0593)	-0.1218* (0.0667)	-0.1309** (0.0572)
Own First Pref, R1	0.4941*** (0.0152)	0.3685*** (0.0170)	0.3398*** (0.0176)	0.3163*** (0.0156)	0.2940*** (0.0134)
Exog. Controls	Yes	Yes	Yes	Yes	Yes
Branch FE	Yes	Yes	Yes	Yes	Yes
Observations	24327	24327	23171	23205	24327
R ²	0.306	0.209	0.182	0.163	0.149
Adjusted R ²	0.305	0.209	0.181	0.162	0.148

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimates of a stacked regression of top preference on average preferences of a cadet's second company peers, and a cadet's own initial preference. These last two objects were measured during the first round, before assignment to the second company. Each column corresponds to a different preference round, beginning in Round 2 and ending in Round 6. The data are constructed such that each cadet reported 17 observations, one for each branch. The outcome variable is an indicator for whether the cadet chose the branch as top choice in the final (sixth) round (*First Pref, R6*). the same for the first round (*Own First Pref, R1*), and the leave-out fraction of seniors in the second company that chose that branch as top choice (*% of 2nd Company Peers, R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. All columns also have Branch FEs, which means that an indicator variable for the branch to which the observation refers is included. Cyber was an option during the third round for one cohort, but it was unavailable as a choice during the first two rounds; it is included in this table for all rounds. Standard errors are clustered at the company-level.

Table 4.1.11: First Company Senior Mentor Preferences on Second Round Ranking

	(1) First Pref, R2	(2) First Pref, R2	(3) First Pref, R2	(4) First Pref, R2
% of 1st Company Seniors, R6	0.259*** (0.019)	0.259*** (0.019)	0.019 (0.040)	0.032 (0.044)
Own First Pref, R1	0.508*** (0.012)	0.508*** (0.012)	0.494*** (0.013)	0.494*** (0.013)
Controls	No	Yes	Yes	Yes
Branch FE	No	No	Yes	No
Branch \times Year FE	No	No	No	Yes
Observations	22896	22896	22896	22896
R ²	0.292	0.292	0.303	0.304
Adjusted R ²	0.291	0.291	0.302	0.303

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates of a stacked regression of second round top preference on the average preferences of the seniors in a cadet's first company, and a cadet's own initial preference. The latter was measured during the first round. The former was measured this year, but it is Round 6, the final round for seniors. The data were constructed such that each cadet reported 17 observations, one for each branch. Each observation contains an indicator for whether the cadet chose the branch as top choice during the third round (*First Pref, R2*), the same for the first round (*Own First Pref, R1*), and the fraction of seniors in the first company that chose that branch as top choice (% of 1st Company Seniors, *R6*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. *Branch FE* means that an indicator variable for the branch to which the observation refers is included. *Branch x Year FE* is a fixed effect for Branch times Year. Cyber was an option during the third round for one cohort, but it was unavailable as a choice during the first two rounds; it is included in this table. Standard errors are clustered at the company-level.

Table 4.1.12: First Company Roommate Preferences on Second Round Ranking

	(1) First Pref, R2	(2) First Pref, R2	(3) First Pref, R2
1st Company Roommate, R1	0.030** (0.014)	0.030** (0.014)	0.004 (0.014)
Own First Pref, R1	0.641*** (0.017)	0.641*** (0.017)	0.614*** (0.018)
Controls	No	Yes	Yes
Branch FE	No	No	Yes
Observations	11776	11776	11776
R^2	0.417	0.417	0.427
Adjusted R^2	0.417	0.416	0.426

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates of a stacked regression of second round top preference on the average preferences of a cadets' roommate(s) during the first company, and a cadet's own initial preference. These latter two objects were both measured during Round 1. The data were constructed such that each cadet reported 16 observations, one for each branch. Due to data limitations, only the 2016 cohort is considered. A small number of observations are excluded due to missing roommate data. Each observation contains an indicator for whether the cadet chose the branch as top choice during the second round (*First Pref, R2*), the same for the first round (*Own First Pref, R1*), and the fraction of the cadet's first company roommates that chose that branch as top choice (% of *1st Company Roommate, R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. *Branch FE* means that an indicator variable for the branch to which the observation refers is included. Cyber was unavailable as a choice during the first two rounds. We exclude it from this analysis. Standard errors are clustered at the company-level.

Table 4.1.13: Variously-Defined Peer Preferences on Third Round Ranking

	(1) First Pref, R3	(2) First Pref, R3	(3) First Pref, R3
% of 2nd Company Peers, R1	0.141** (0.056)	0.141** (0.056)	-0.011 (0.068)
% of 1st Company Peers, R1	0.125** (0.057)	0.125** (0.057)	-0.014 (0.063)
% of 1st Company Seniors, R6	0.123*** (0.039)	0.123*** (0.039)	-0.008 (0.056)
1st Company Roommate, R1	0.015 (0.017)	0.015 (0.017)	0.006 (0.017)
Own First Pref, R1	0.442*** (0.018)	0.442*** (0.018)	0.434*** (0.018)
Controls	No	Yes	Yes
Branch FE	No	No	Yes
Observations	12512	12512	12512
R ²	0.259	0.259	0.267
Adjusted R ²	0.258	0.258	0.266

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates of a stacked regression of third round top preferences on average preferences of a cadet's second company peers, average preferences of a cadet's first company peers (excluding roommates), senior preferences from a cadet's first company, the first company roommate, and a cadet's own initial preference. These last two objects were measured during the first round, before assignment to the second company. The data were constructed such that each cadet reported 17 observations, one for each branch. Each observation contains an indicator for whether the cadet chose the branch as top choice during the third round (*First Pref, R3*), the same for the first round (*Own First Pref, R1*), the leave-out fraction of peers in the second company that chose that branch as top choice (*% of 2nd Company Peers, R1*), and the leave-out fraction of peers in the first company that chose that branch as top choice, excluding roommates (*% of 1st Company Peers, R1*), and the leave-out fraction of seniors in the first company that chose that branch as top choice (*% of 2nd Company Peers, R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. *Branch FE* means that an indicator variable for the branch to which the observation refers is included. *Branch x Year FE* is a fixed effect for Branch times Year. Cyber was an option during the third round for one cohort, but it was unavailable as a choice during the first two rounds; it is included in this table. Standard errors are clustered at the company-level.

4.2 Tables for Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?

Table 4.2.1: Cumulative Percentage of PhDs Received by Time-to-Degree, by Field, BA-Years 1980 and 1981

TTD	Humanities	Life Science	Physical Science	Social Science	Total
4	0.8	1.3	3.6	1.2	2.1
5	2.0	7.6	14.5	4.4	9.1
6	5.1	20.2	30.7	10.9	20.9
7	10.8	32.9	45.2	17.8	32.6
8	17.5	44.4	56.0	27.3	42.8
9	25.4	53.0	64.6	35.5	51.2
10	34.4	60.5	70.6	44.0	58.5
11	42.7	66.9	75.8	52.9	65.0
12	50.7	71.8	80.0	59.1	70.2
13	57.6	75.6	83.4	65.0	74.6
14	63.2	78.9	86.3	69.9	78.3
15	68.6	81.5	88.5	74.0	81.4
16	72.8	83.6	90.7	76.8	83.9
17	76.7	85.7	92.2	80.1	86.2
18	79.9	87.3	93.5	82.8	88.1
19	82.7	88.9	94.3	85.5	89.7
20	85.2	90.2	94.9	86.8	90.9
21	87.2	91.4	95.5	88.6	92.0
22	89.0	92.4	96.0	90.2	93.0
23	90.6	93.3	96.4	91.3	93.9
24	91.8	94.1	96.9	92.4	94.7
25	93.1	94.9	97.3	93.5	95.4
26	94.2	95.7	97.6	94.3	96.0
27	95.2	96.4	97.9	95.2	96.6
28	95.8	96.8	98.1	95.7	97.1
29	96.5	97.3	98.5	96.3	97.6
30	97.2	97.9	98.7	97.0	98.0
31	98.1	98.4	99.0	97.6	98.5
32	98.6	98.8	99.3	98.4	98.9
33	99.1	99.3	99.5	99.0	99.3
34	99.5	99.6	99.7	99.6	99.6
35	100.0	100.0	100.0	100.0	100.0

Notes: This table shows the cumulative percentage of PhDs received by the time-to-degree, TTD. This is displayed separately by field as well as aggregated together. BA-years of 1980 and 1981 are used. TTDs of less than 4 are combined with TTD of 4. TTD of 36 is combined with TTD of 35. The entry corresponding to the first TTD for which at least 90 percent of PhDs were earned is indicated in bold.

Table 4.2.2: BA Fields of Study of PhD Recipients, by PhD Field

BA Field	PhD Field			
	Humanities	Life Science	Physical Science	Social Science
Humanities	82.4	2.2	0.9	13.2
Life Science	1.3	78.0	3.6	4.5
Physical Science	2.3	11.7	93.6	7.0
Social Science	6.6	2.7	0.7	64.0
Other	7.5	5.4	1.2	11.3
<i>N</i>	45521	89488	104591	34509

Notes: This table displays the percentage of PhD recipients in a given field (Humanities, Life Science, Physical Science, Social Science) who obtained their BA in each of these four fields as well as in subjects outside of these; this category is referred to as “Other.” PhD fields are shown in the columns; BA fields are shown in the rows. Observations with no BA field listed are excluded. Observations are not restricted by TTD, and only PhDs actually observed are shown (meaning no truncation correction was used).

Table 4.2.3: Summary Statistics, Full Sample and by Carnegie Classification

	All			Carnegie 15			Carnegie 16			Carnegie 17		
	Mean	SD	Miss	Mean	SD	Miss	Mean	SD	Miss	Mean	SD	Miss
Stud-FT Fac Ratio	18.2	14.4	0	15.5	20	0	18.9	9.6	0	21.2	7	0
Public	.7	.5	0	.7	.5	0	.7	.4	0	.5	.5	0
Instr Exp/Student	7.5	6.8	.06	10.5	8.8	.04	6.3	4.7	.09	4.6	2.3	.03
Rsch Exp/FT Fac	48.4	75.6	.14	76.1	93.2	.04	37.3	58.6	.09	10.1	13.6	.32
75 Percentile Score	26.6	3.1	.13	28.3	3	.1	26	2.7	.11	24.6	2.3	.2
Total Students	13.1	8.9	0	19.7	9.4	0	10.7	5.3	0	6.7	4.5	0
% UG	77.3	13.1	0	72.5	12.8	0	79.8	12.6	.01	81.1	12.1	0
% Female	50.7	9.7	0	48.5	6.8	0	49.5	10.8	.01	55.5	9.9	0
% Minority	16.4	18.2	0	12.8	8.7	0	16.7	20.8	.01	21.2	23.1	0
% Hum	12	7.7	0	14.4	7.4	0	10.8	8.2	0	10	6.4	0
% LifSci	11.7	7.3	0	12.2	5.2	0	11.1	6.6	0	11.5	10.2	0
% PhysSci	15.6	16.3	0	17.5	13.1	0	18.7	21.5	0	9	9.4	0
% SocSci	10.1	8.1	0	14.1	9	0	8	6.7	0	7.1	5.5	0
% Other	50.6	18.2	0	41.7	16.8	0	51.5	18.1	0	62.4	12.3	0
Hum NRC 0-10	.5	1.6	.52	.7	1.9	.17	0	0	.61	0	0	.91
Hum NRC 11-25	.7	1.2	.52	1	1.3	.17	0	0	.61	0	0	.91
Hum NRC 26-50	1.1	1.4	.52	1.5	1.4	.17	.1	.3	.61	0	0	.91
Hum NRC 51-100	2.2	1.7	.52	2.4	1.8	.17	1.8	1.5	.61	1	0	.91
LifSci NRC 0-10	.6	1.5	.43	.9	1.7	.04	0	0	.5	0	0	.92
LifSci NRC 11-25	.8	1.3	.43	1.2	1.5	.04	0	.1	.5	0	0	.92
LifSci NRC 26-50	1.3	1.7	.43	1.8	1.8	.04	.3	.8	.5	0	0	.92
LifSci NRC 51-100	2.3	2	.43	1.9	2	.04	3.1	1.7	.5	2	1.4	.92
PhySci NRC 0-10	.8	2.4	.35	1.4	3	.03	0	0	.3	0	0	.9
PhySci NRC 11-25	1	1.9	.35	1.7	2.2	.03	.1	.4	.3	0	0	.9
PhySci NRC 26-50	1.8	2.4	.35	2.8	2.6	.03	.4	1.2	.3	.1	.3	.9
PhySci NRC 51-100	3.4	2.7	.35	3.5	3	.03	3.4	2.3	.3	1.8	1.6	.9
SocSci NRC 0-10	.4	1	.55	.5	1.2	.12	0	0	.73	0	0	.96
SocSci NRC 11-25	.5	.9	.55	.6	1	.12	0	0	.73	.3	.5	.96
SocSci NRC 26-50	.8	1.1	.55	1	1.1	.12	.1	.3	.73	0	0	.96
SocSci NRC 51-100	1.7	1.5	.55	1.7	1.6	.12	1.8	1	.73	1	0	.96
N	7592			2964			2600			2028		

Notes: This table displays means, standard deviations, and percent of observations that are missing. Observations are unweighted and are displayed separately by Carnegie category of the institution. The sample shown is the Life Science sample, as this was the field with the most institutions. Observations are at the institution-BA year, and each institution contributes 26 observations, one for each year between 1980 and 2005. Means and standard deviations do not include missing observations in the calculations. Total students, instructional expenditures per student, and research expenditures per student are in units of 1,000. The missing column displays the fraction of observations that are missing after interpolating and filling in as described in the text. The missing column for the NRC variables is the fraction that were not observed with at least one department for the respective field (Humanities, Life Science, Physical Science, Social Science).

Table 4.2.4: Baseline Regression

	Humanities	Life Sci.	Physical Sci.	Social Sci.
Stud-FT Fac Ratio	0.005 (0.008)	-0.026 (0.017)	0.006 (0.008)	-0.004 (0.007)
Public	-0.251 (0.403)	-0.317 (0.597)	-0.208 (0.570)	-0.092 (0.273)
Instr Exp/Student	0.061*** (0.023)	0.041 (0.046)	0.098* (0.051)	0.023 (0.017)
Rsch Exp/FT Fac	-0.001 (0.002)	0.008 (0.006)	-0.003 (0.002)	0.001 (0.002)
75 Percentile Score	0.218*** (0.043)	0.757*** (0.087)	0.376*** (0.072)	0.225*** (0.034)
Total Students	-0.046*** (0.014)	-0.112*** (0.023)	-0.107*** (0.034)	-0.044*** (0.012)
% UG	0.004 (0.013)	0.002 (0.025)	-0.009 (0.024)	-0.002 (0.012)
% Female	-0.013 (0.015)	-0.077** (0.035)	-0.034 (0.037)	-0.042*** (0.013)
% Minority	0.005 (0.005)	0.006 (0.007)	-0.003 (0.007)	0.008** (0.004)
% of Total BAs in (Field)	-0.041* (0.021)	-0.145*** (0.028)	-0.048*** (0.018)	-0.064*** (0.016)
#NRC 0-10 in (Field)	0.444*** (0.135)	0.648** (0.254)	0.539*** (0.090)	0.288** (0.136)
#NRC 11-25 in (Field)	0.238** (0.101)	0.168 (0.163)	-0.028 (0.123)	0.104 (0.121)
#NRC 26-50 in (Field)	-0.070 (0.110)	0.110 (0.116)	0.128* (0.075)	-0.023 (0.080)
#NRC 51-100 in (Field)	0.004 (0.050)	-0.109 (0.097)	-0.126* (0.068)	-0.003 (0.045)
Carnegie 16	-0.220 (0.217)	-0.541 (0.395)	-0.959** (0.406)	-0.034 (0.209)
Carnegie 17	-0.543 (0.337)	-0.146 (0.625)	-1.802*** (0.523)	-0.635** (0.284)
PhD Prod. Rate Mean	2.61	5.11	4.65	2.21
Missing Dummies	Yes	Yes	Yes	Yes
N	7237	7494	7368	7122
Adjusted R-squared	0.48	0.55	0.58	0.25

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results. The dependent variable is the PhD production rate, which is a percentage, not a share. Maximum TTD and truncation correction are as described in Section 2.4. Each column is a separate regression and corresponds to a field. Observations are weighted by the number of field-specific BAs from the BA-year. Observations are at the institution-BA year, and each institution in the field's sample contributes 26 observations, one for each year between 1980 and 2005. Total students, instructional expenditures per student, and research expenditures per student are in units of 1,000. Percentage undergraduate, female, minority, and total majors in (Field) are percentages, out of 100. BA-year dummies are included, as are missing variable dummies. Robust standard errors are clustered at the institution level.

Table 4.2.5: Regression, Carnegie Category 15 Institutions Only

	Humanities	Life Sci.	Physical Sci.	Social Sci.
Stud-FT Fac Ratio	0.001 (0.012)	-0.052** (0.022)	0.012 (0.013)	-0.010 (0.010)
Public	0.036 (0.618)	0.023 (1.119)	0.546 (0.714)	0.321 (0.408)
Instr Exp/Student	0.067** (0.028)	0.031 (0.054)	0.083 (0.063)	0.013 (0.019)
Rsch Exp/FT Fac	0.000 (0.003)	0.016** (0.007)	-0.004 (0.004)	0.003 (0.003)
75 Percentile Score	0.239*** (0.065)	1.003*** (0.127)	0.436*** (0.105)	0.287*** (0.051)
Total Students	-0.056*** (0.020)	-0.133*** (0.031)	-0.112*** (0.037)	-0.052*** (0.015)
% UG	0.004 (0.021)	0.017 (0.041)	-0.049 (0.036)	-0.013 (0.016)
% Female	-0.038 (0.023)	-0.131*** (0.050)	-0.075 (0.067)	-0.047*** (0.017)
% Minority	0.012 (0.015)	0.013 (0.018)	-0.035 (0.024)	0.013 (0.011)
% of Total BAs in (Field)	-0.035 (0.043)	-0.229*** (0.069)	-0.062 (0.042)	-0.066*** (0.022)
#NRC 0-10 in (Field)	0.406** (0.161)	0.427 (0.301)	0.479*** (0.117)	0.215 (0.145)
#NRC 11-25 in (Field)	0.227** (0.109)	0.074 (0.165)	-0.088 (0.152)	0.100 (0.127)
#NRC 26-50 in (Field)	-0.078 (0.149)	0.082 (0.152)	0.074 (0.108)	-0.014 (0.094)
#NRC 51-100 in (Field)	0.040 (0.059)	-0.116 (0.146)	-0.224** (0.108)	0.034 (0.052)
PhD Prod. Rate Mean	3.05	6.44	5.56	2.42
Missing Dummies	Yes	Yes	Yes	Yes
N	2923	2959	2959	2943
Adjusted R-squared	0.52	0.58	0.61	0.32

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table displays regression results. The sample is restricted to Carnegie category 15 institutions. The dependent variable is the PhD production rate, which is a percentage, not a share. Maximum TTD and truncation correction are as described in Section 2.4. Each column is a separate regression and corresponds to a field. Observations are weighted by the number of field-specific BAs from the BA-year. Observations are at the institution-BA year, and each institution in the field's sample contributes 26 observations, one for each year between 1980 and 2005. A maximum TTD of 15 years is used, and BA-years 2002 and after are truncation-corrected. Total students, instructional expenditures per student, and research expenditures per student are in units of 1,000. Percentage undergraduate, female, minority, and total majors in (Field) are percentages, out of 100. BA-year dummies are included, as are missing variable dummies. Robust standard errors are clustered at the institution level.

Table 4.2.6: Lemke-Like Regression

	Humanities	Life Sci.	Physical Sci.	Social Sci.
Stud-FT Fac Ratio	0.021 (0.016)	-0.032 (0.034)	0.011 (0.027)	0.010 (0.013)
Public	0.135 (0.427)	0.281 (0.587)	0.369 (0.596)	0.199 (0.243)
Instr Exp/Student	0.102* (0.053)	-0.006 (0.066)	0.157* (0.091)	0.030 (0.029)
Rsch Exp/FT Fac	-0.001 (0.004)	0.015* (0.008)	-0.004 (0.006)	0.003 (0.003)
75 Percentile Score	0.327*** (0.066)	0.942*** (0.127)	0.623*** (0.139)	0.315*** (0.052)
Total Students	-0.047*** (0.015)	-0.082*** (0.025)	-0.104*** (0.039)	-0.035*** (0.012)
% UG	0.010 (0.013)	-0.000 (0.027)	0.010 (0.025)	0.001 (0.011)
% Female	0.008 (0.018)	-0.056 (0.040)	0.000 (0.043)	-0.021 (0.014)
% Minority	0.005 (0.007)	0.008 (0.010)	-0.000 (0.010)	0.008* (0.005)
% of Total BAs in (Field)	-0.050** (0.021)	-0.121*** (0.026)	-0.045** (0.021)	-0.057*** (0.015)
#NRC 0-10 in (Field)	0.423*** (0.129)	0.321 (0.252)	0.529*** (0.118)	0.208 (0.132)
#NRC 11-25 in (Field)	0.197* (0.104)	-0.063 (0.166)	-0.047 (0.134)	0.007 (0.112)
#NRC 26-50 in (Field)	-0.157 (0.136)	-0.047 (0.118)	0.135 (0.094)	-0.021 (0.067)
#NRC 51-100 in (Field)	0.050 (0.054)	-0.135 (0.100)	-0.100 (0.072)	0.003 (0.043)
Carnegie 16	-0.191 (0.214)	-0.192 (0.442)	-0.514 (0.454)	0.167 (0.197)
Carnegie 17	-0.258 (0.402)	0.708 (0.648)	-1.187* (0.646)	-0.033 (0.292)
PhD Prod. Rate Mean	2.32	4.53	4.97	1.83
Missing Dummies	Yes	Yes	Yes	Yes
N	281	291	284	276
Adjusted R-squared	0.48	0.65	0.67	0.45

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

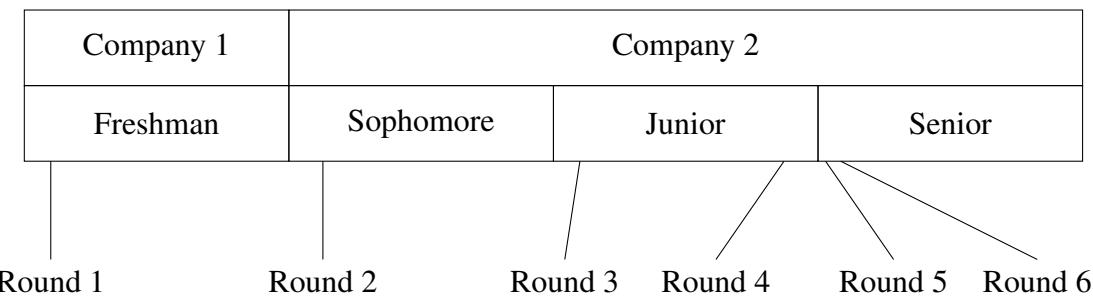
Notes: This table displays regression results. The dependent variable is the number of PhDs produced between 1994-2003 divided by the number of BAs earned from 1989-1998; this is a percentage, not a share. Each column is a separate regression and corresponds to a field. Observations are weighted by the number of field-specific BAs from 1989-1998. Each institution in the field's sample contributes one observation. The independent variables are averages over the 1989-1998 period. Averages are computed among non-missing years; if all years are missing, the missing indicator is one. Total students, instructional expenditures per student, and research expenditures per student are in units of 1,000. Percentage undergraduate, female, minority, and total majors in (Field) are percentages, out of 100. Missing dummies are included. Robust standard errors are used.

CHAPTER 5

FIGURES

5.1 Figures for Do Peers Influence Occupational Preferences?

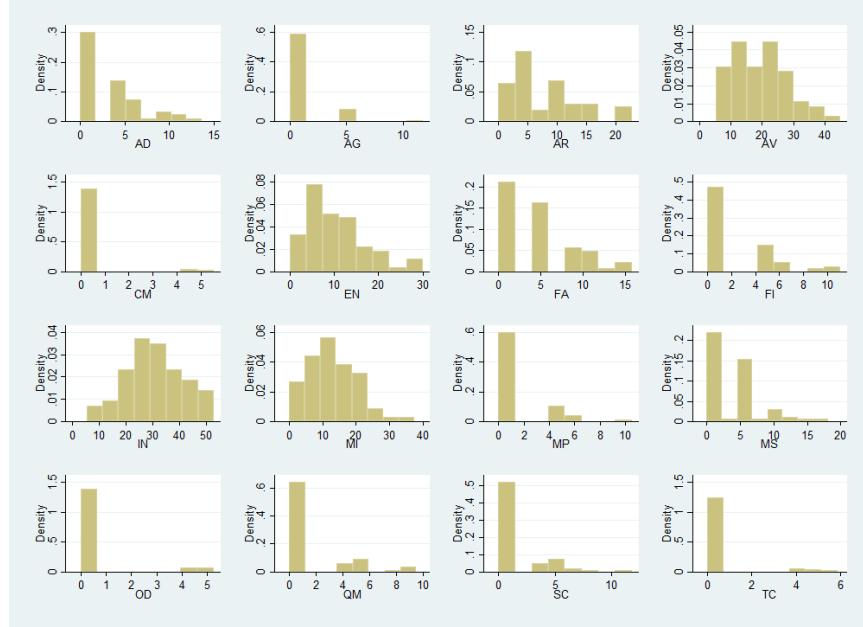
Figure 5.1.1: Timing of Company Assignment and Branch Preference Rounds



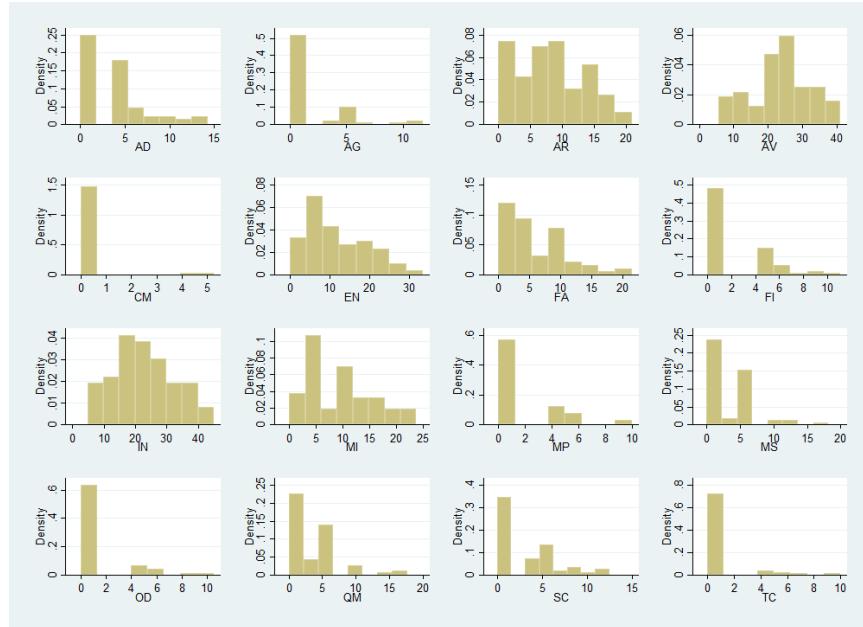
Notes: This figure shows a timeline of the first and second company assignments, and the timing of the preference rounds. A cadet is in Company 1 during freshman year, and Company 2 during sophomore through senior years. Preferences were elicited 6 times, once each during the freshman and sophomore years (during approximately the first and second weeks of September, respectively), twice during the junior year (the first during approximately the first week of September and the second during April), and twice during the senior year (the first in August, and the second during approximately the first week of September).

Figure 5.1.2: Variation in Branch Preferences across Companies, Rounds 1 and 2

Panel A: Preference Round 1

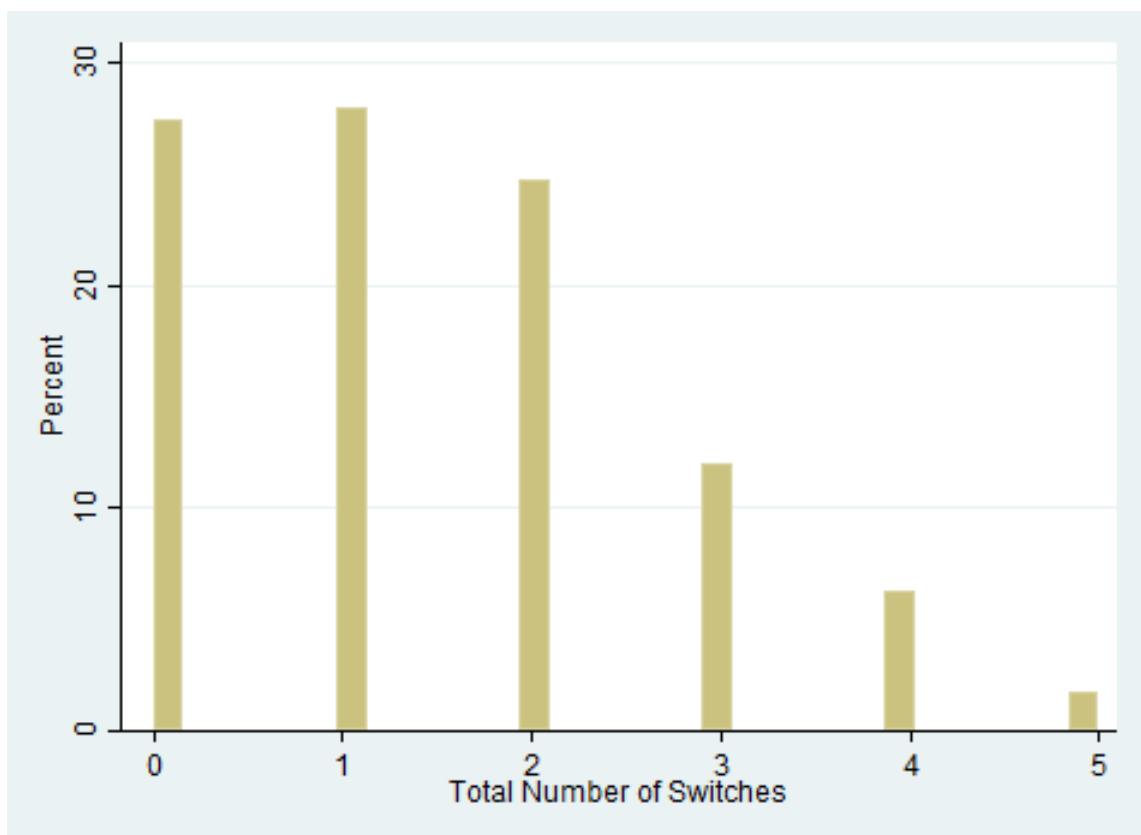


Panel B: Preference Round 2



Notes: These figures show variation in branch preferences across companies. Each histogram refers to a given branch, with labels as defined in Table 4.1.1. Each observation of a given histogram is the percentage of a company that prefers the given branch as top choice. Panel A is for preferences in round 1, while Panel B is for preferences in round 2. Cyber is not shown as it is always zero for these rounds as it was not an available choice.

Figure 5.1.3: Distribution of Number of Times a Cadet Switches Top Preference



Notes: This figure shows the percentage of cadets who switched their top branch preferences 0,...,5 times. The sample size is less than 1431 because it is limited to those with no missing preference data. Switching was calculated between rounds. Since there were 6 rounds, the maximum number of switches is five, which means that a top-rated branch changed between every round, but a cadet might have listed the same branch more than once, but not consecutively.

Figure 5.1.4: Evolution of Preferences across Rounds 1, 2, 3, and 6

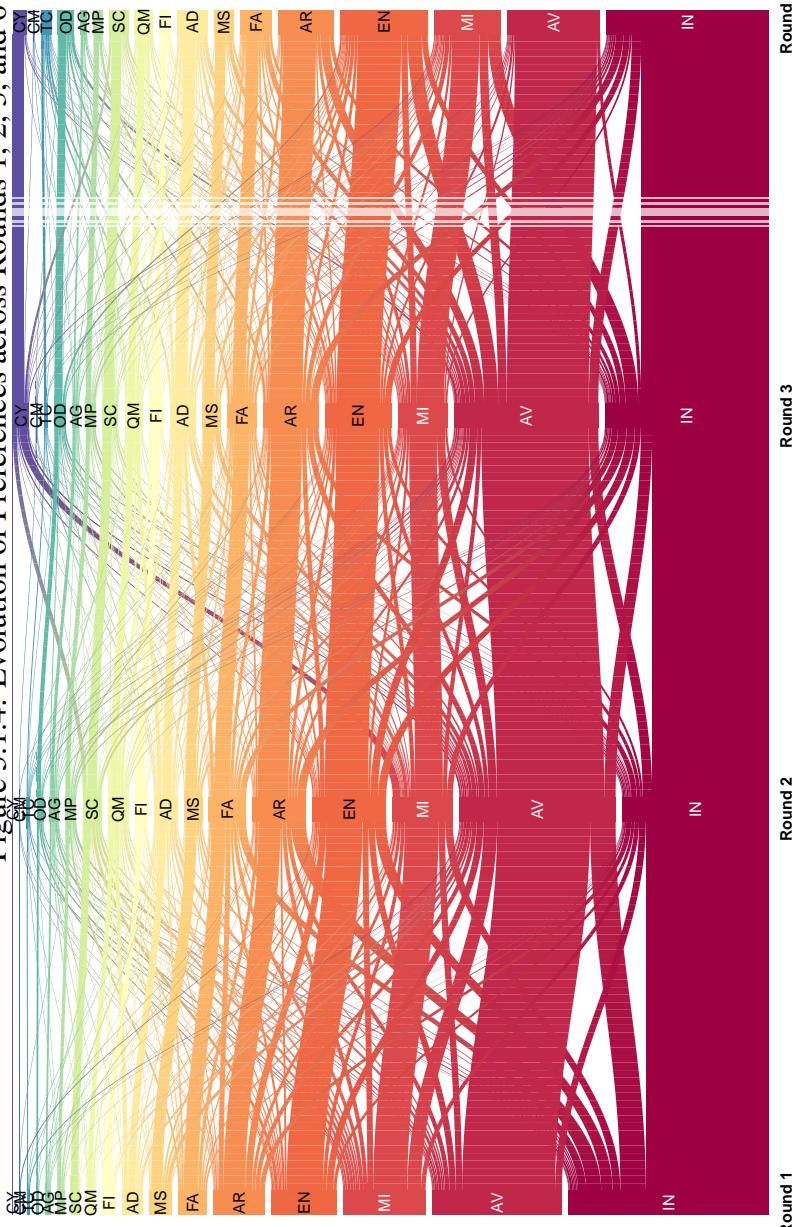
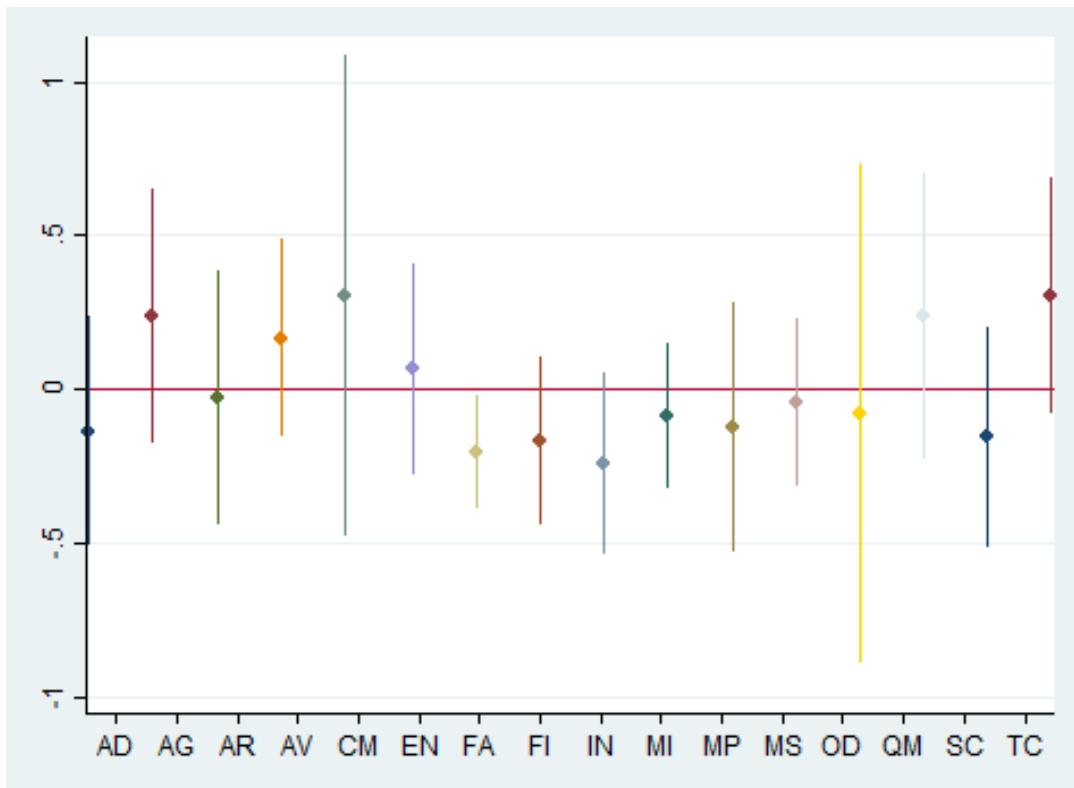
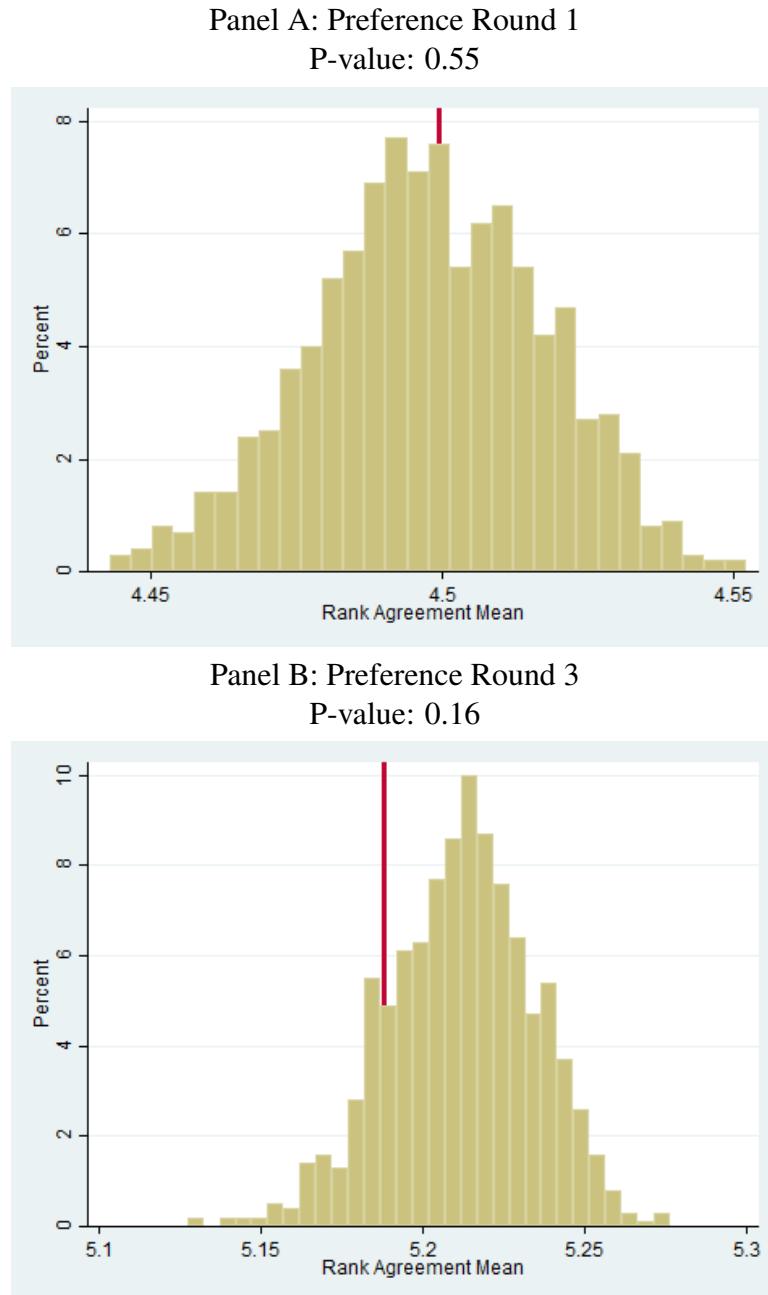


Figure 5.1.5: Branch-Specific Regressions; Second Company Same-Cohort Peer Preferences on Third Round Ranking



Notes: This figure shows estimates of 16 individual regressions, one for each branch, of third round top preference on the average preferences of a cadet's second company peers, and a cadet's own initial preference for the branch. These last two objects were measured during the first round, before assignment to the second company, to avoid the reflection problem. Each observation contains an indicator for whether a cadet chose the branch as top choice during the third round, the same for the first round, and the leave-out fraction of cadets in the second round company that chose that branch as top choice. Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. Year fixed effects are included. Cyber was an option during the third round, at least for one cohort, but it was unavailable as a choice during the first two rounds. We exclude it from this analysis. See Table 4.1.1 for the meaning of branch initials. Confidence intervals refer to an alpha of 0.05. Standard errors are clustered at the company-level.

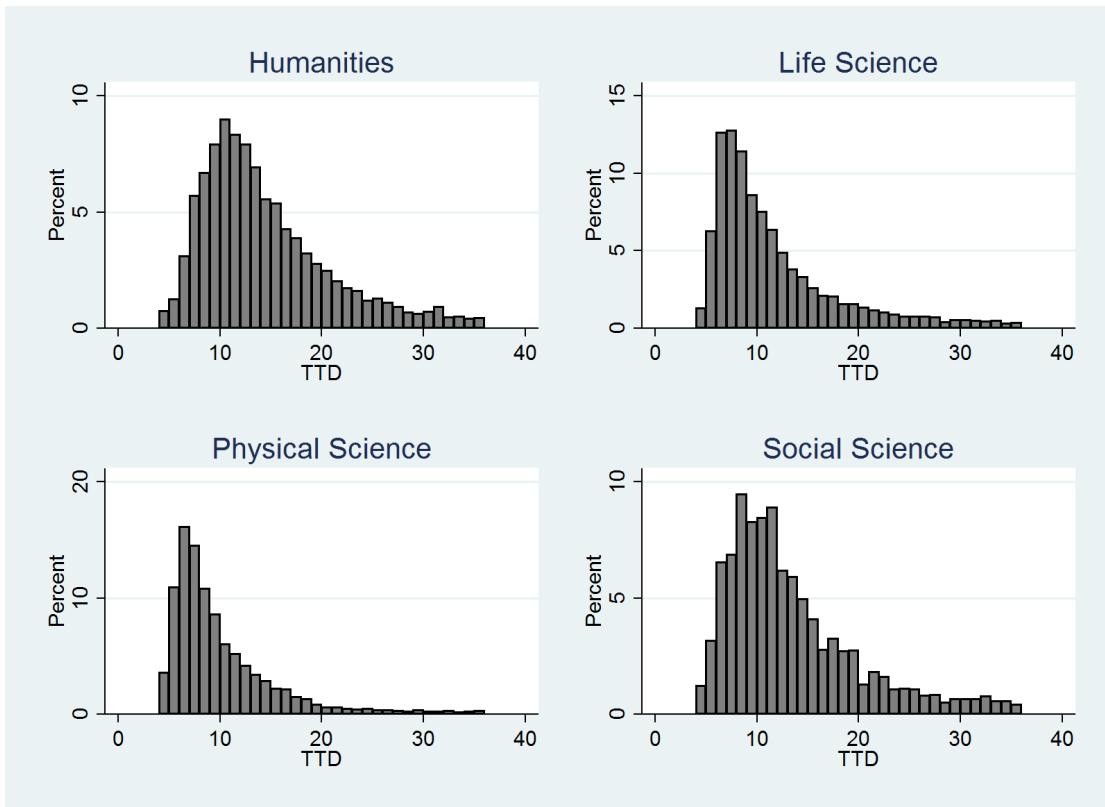
Figure 5.1.6: Rank Agreement, Company 2, Round 1 and Round 3



Notes: These figures show the average rank agreement distribution of 1,000 synthetic companies, with the red line indicating the rank agreement from the actual data. Panel A is for preference round 1, calculated in relation to company 2. Panel B is for preference round 3, also calculated in relation to company 2. Rank agreement calculation is described in Section 1.8. P-values are also indicated.

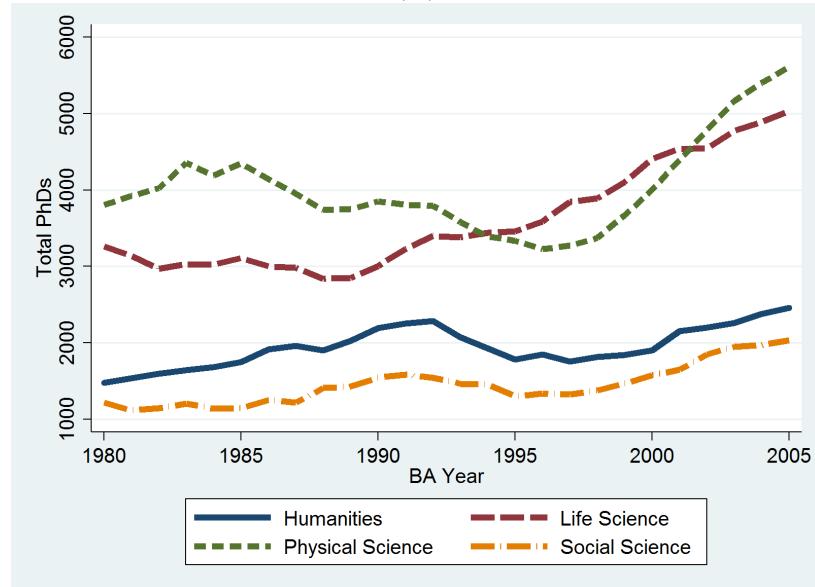
5.2 Figures for Are High-Quality PhD Programs at Universities Associated with More Undergraduate Students Pursuing PhD Study?

Figure 5.2.1: Distribution of Time-to-Degree, BA-Years 1980 and 1981

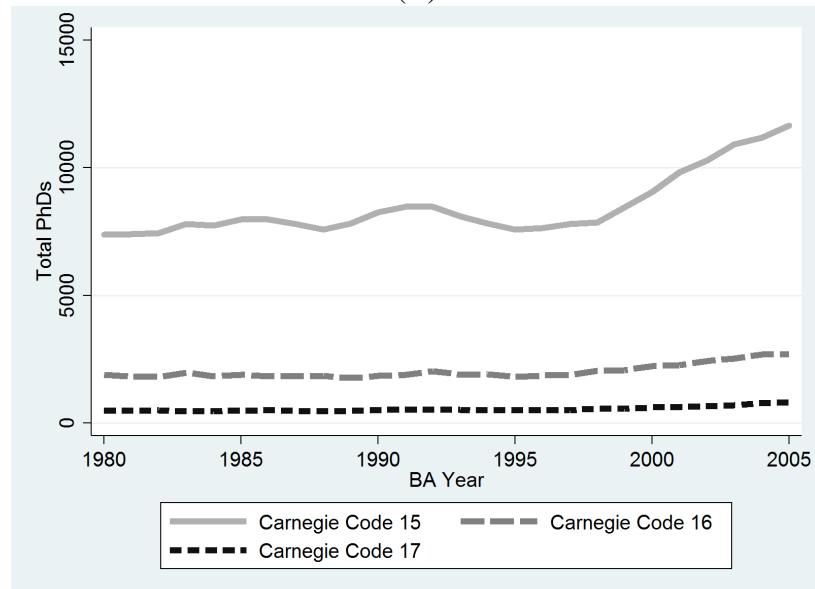


Notes: This figure plots the number of PhDs received by TTD, the number of years after the BA-year that the PhD was received. Only BA-years of 1980 and 1981 are considered. TTDs of less than 4 are grouped with TTD of 4; TTD of 36, which is only observed for 1980, is grouped with TTD of 35.

Figure 5.2.2: Total PhDs Produced over BA-Year, by Field and by Carnegie Classification
 (A)

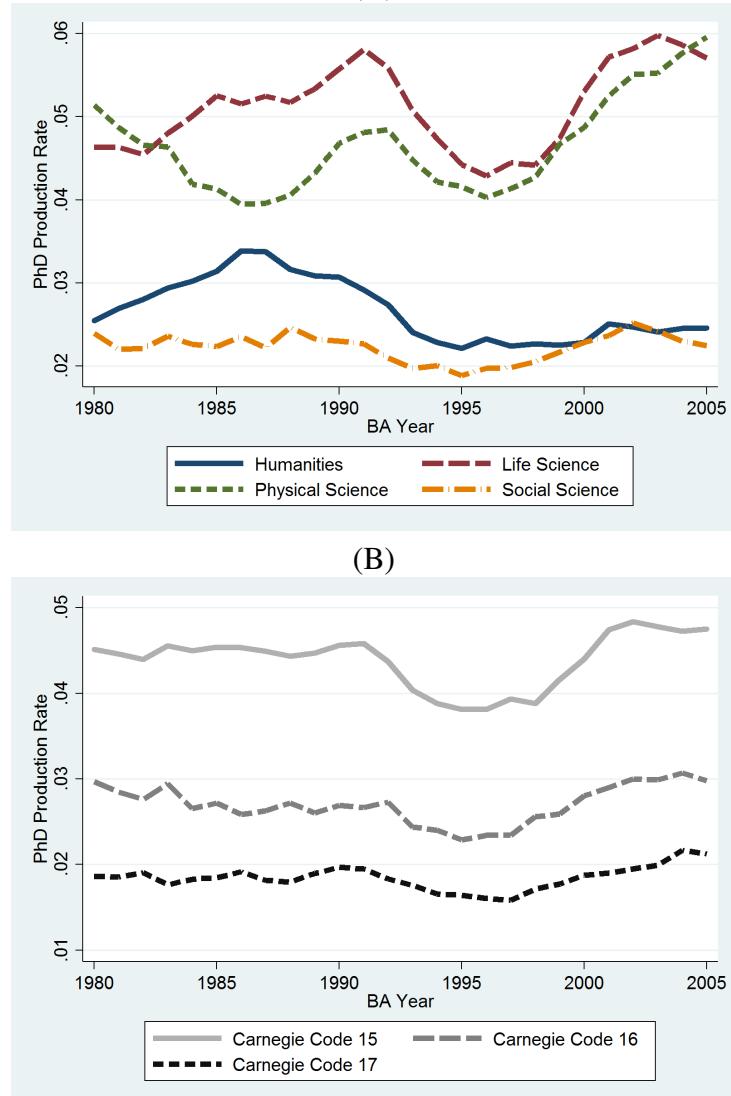


(B)



Notes: Panel A plots the total number of PhDs produced each BA-year, separately by PhD field (Humanities, Life Science, Physical Science, Social Science). Panel B plots the same, but splits by the Carnegie category of the institution. Carnegie categories are: 15—Doctoral Universities: Highest Research Activity, 16—Doctoral Universities: Higher Research Activity, and 17—Doctoral Universities: Moderate Research Activity. Maximum TTDs for fields are as described in Section 2.4.1.

Figure 5.2.3: PhD Production Rate over BA-Year, by Field and by Carnegie Classification



Notes: Panel A plots the PhD production rate each BA-year, separately by PhD field (Humanities, Life Science, Physical Science, Social Science). Panel B plots the same, but splits by the Carnegie category of the institution. The PhD production rate in Figure A is computed by dividing the total number of PhDs of a given field (for a given BA-year) by the total number of BAs from the corresponding field and BA-year. The PhD production rate in Panel B is computed by dividing the total number of PhDs (for a given BA-year and across all four fields) of a Carnegie Classification by the total number of BAs across all four fields. Carnegie categories are: 15—Doctoral Universities: Highest Research Activity, 16—Doctoral Universities: Higher Research Activity, and 17—Doctoral Universities: Moderate Research Activity. Maximum TTD and truncation correction are as described in Section 2.4.

CHAPTER 6

APPENDICES FOR DO PEERS INFLUENCE OCCUPATIONAL PREFERENCES?

6.1 Rank Agreement Example

This section gives a stylized example of how to calculate rank agreement for one cadet for one year in relation to other cadets in the company. Suppose there are only three cadets in a company—Cadet 1, Cadet 2, and Cadet 3. Suppose also that the cadets rank only 5 branches (instead of the usual 16)—Aviation, Chemical Corps, Engineering, Field Artillery, and Finance. Cadet 1 has the following preferences, where the numbers correspond to the rank given to them. Aviation is the top choice, Field Artillery is the second choice, etc.:

$$Pref_1 \begin{pmatrix} Aviation \\ Chemical Corps \\ Engineering \\ Field Artillery \\ Finance \end{pmatrix} = \begin{pmatrix} 1 \\ 5 \\ 4 \\ 2 \\ 3 \end{pmatrix} \quad (6.1)$$

Cadet 2 has the following preferences:

$$Pref_2 \begin{pmatrix} Aviation \\ Chemical Corps \\ Engineering \\ Field Artillery \\ Finance \end{pmatrix} = \begin{pmatrix} 2 \\ 4 \\ 1 \\ 3 \\ 5 \end{pmatrix} \quad (6.2)$$

Cadet 3 has the following preferences:

$$Pref_3 \begin{pmatrix} Aviation \\ Chemical Corps \\ Engineering \\ Field Artillery \\ Finance \end{pmatrix} = \begin{pmatrix} 3 \\ 4 \\ 2 \\ 5 \\ 1 \end{pmatrix} \quad (6.3)$$

We compute rank agreement for Cadet 1 with respect to Cadets 2 and 3. We first compute the leave-out mean of the company's preferences, the mean position of each branch of all cadets other than Cadet 1:

$$\overline{Pref}_{2,3} \begin{pmatrix} Aviation \\ Chemical Corps \\ Engineering \\ Field Artillery \\ Finance \end{pmatrix} = \begin{pmatrix} 2.5 \\ 4 \\ 1.5 \\ 4 \\ 3 \end{pmatrix} \quad (6.4)$$

For Cadet 1, we calculate the absolute differences of choices with those of the remain-

der of the company, where $abs(\cdot)$ denotes absolute value:

$$DiffPref_1 \begin{pmatrix} Aviation \\ Chemical Corps \\ Engineering \\ Field Artillery \\ Finance \end{pmatrix} = abs(Pref_1 - \overline{Pref_{2,3}}) = abs \begin{pmatrix} 1 \\ 5 \\ 4 \\ 2 \\ 3 \end{pmatrix} - \begin{pmatrix} 2.5 \\ 4 \\ 1.5 \\ 4 \\ 3 \end{pmatrix} = \begin{pmatrix} 1.5 \\ 1 \\ 2.5 \\ 2 \\ 0 \end{pmatrix} \quad (6.5)$$

Until now, we have sorted the vectors alphabetically by branch name. To make the next step clearer, we resort the difference vector such that the branches are in the order in which Cadet 1 ranked them:

$$DiffPref_1^{Sorted} \begin{pmatrix} Aviation \\ Field Artillery \\ Finance \\ Engineering \\ Chemical Corps \end{pmatrix} = \begin{pmatrix} 1.5 \\ 2 \\ 0 \\ 2.5 \\ 1 \end{pmatrix} \quad (6.6)$$

We weight by observed probabilities that cadets are assigned to their 1st, 2nd, 3rd, etc. preference. Suppose that the realized branch that cadets obtain is given with the following

probabilities over the preference ranking:

$$\begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{bmatrix} = \begin{bmatrix} 0.70 \\ 0.15 \\ 0.10 \\ 0.04 \\ 0.01 \end{bmatrix} \quad (6.7)$$

70% of cadets are assigned to their most-preferred branch, 15% are assigned to their second, etc. We use these probabilities to weight the absolute differences and calculate a rank agreement for Cadet 1:

$$RankAgreement_1 = (.7) * 1.5 + (.15) * 2 + (.10) * 0 + (.04) * 2.5 + (.01) * 1 = 1.46 \quad (6.8)$$

We perform analogous steps to compute rank agreements for Cadets 2 and 3.

6.2 Tables

Table 6.2.1: Point Estimates for Branch-Specific Regressions; Second Company Same-Cohort Peer Preferences on Third Round Ranking

AD	-0.137 (0.184)	AG	0.237 (0.203)	AR	-0.027 (0.204)	AV	0.165 (0.158)
CM	0.304 (0.387)	EN	0.066 (0.170)	FA	-0.205 (0.091)	FI	-0.169 (0.136)
IN	-0.242 (0.146)	MI	-0.088 (0.118)	MP	-0.124 (0.198)	MS	-0.043 (0.133)
OD	-0.082 (0.400)	QM	0.237 (0.230)	SC	-0.156 (0.177)	TC	0.302 (0.189)

Notes: This table shows point estimates of 16 individual regressions, one for each branch, of third round top preference on the average preferences of a cadet's second company peers, and a cadet's own initial preference for the branch. These last two objects were measured during the first round, before assignment to the second company. Each observation contains an indicator for whether a cadet chose the branch as top choice during the third round, the same for the first round, and the leave-out fraction of cadets in the second round company that chose that branch as top choice. Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. Year fixed effects are included. Cyber was an option during the third round, at least for one cohort, but it was unavailable as a choice during the first two rounds. We exclude it from this analysis. See Table 4.1.1 for the meaning of branch initials. Confidence intervals refer to an alpha of 0.05. Standard errors are clustered at the company-level.

Table 6.2.2: Second Company Same-Cohort Peer Preferences on Third Round Ranking, Excluding Cyber

	(1) First Pref, R3	(2) First Pref, R3	(3) First Pref, R3	(4) First Pref, R3
% of 2nd Company Peers, R1	0.4046*** (0.0223)	0.4046*** (0.0223)	-0.0021 (0.0676)	-0.0018 (0.0676)
Own First Pref, R1	0.3867*** (0.0137)	0.3867*** (0.0137)	0.3750*** (0.0135)	0.3745*** (0.0134)
Controls	No	Yes	Yes	Yes
Branch FE	No	No	Yes	No
Branch \times Year FE	No	No	No	Yes
Observations	22160	22160	22160	22160
R^2	0.209	0.209	0.218	0.219
Adjusted R^2	0.209	0.208	0.216	0.216

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table shows estimates of a stacked regression of third round top preferences on average preferences of a cadet's second company peers and a cadet's own initial preference. These last two objects were measured during the first round, before assignment to the second company. The data were constructed such that each cadet reported 17 observations, one for each branch. Each observation contains an indicator for whether the cadet chose the branch as top choice during the third round (*First Pref, R3*), the same for the first round (*Own First Pref, R1*), and the leave-out fraction of peers in the second company that chose that branch as top choice (% of 2nd Company Peers, *R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. *Branch FE* means that an indicator variable for the branch to which the observation refers is included. *Branch x Year FE* is a fixed effect for Branch times Year. Cyber is excluded as a choice; all cadets who chose Cyber as top choice in Round 3 are also excluded. Standard errors are clustered at the company-level.

Table 6.2.3: First Company Same-Cohort Peer Preferences on 2nd–6th Round Preferences

	(1)	(2)	(3)	(4)	(5)
	1st Pref, R2	1st Pref, R3	1st Pref, R4	1st Pref, R5	1st Pref, R6
% of 1st Cpy. Peers, R1	-0.0212 (0.0515)	-0.0099 (0.0654)	-0.0070 (0.0456)	-0.0236 (0.0594)	-0.0047 (0.0588)
Own First Pref, R1	0.4940*** (0.0132)	0.3687*** (0.0135)	0.3402*** (0.0162)	0.3166*** (0.0146)	0.2942*** (0.0154)
Controls	Yes	Yes	Yes	Yes	Yes
Branch FE	Yes	Yes	Yes	Yes	Yes
Observations	24327	24327	23171	23205	24327
R^2	0.306	0.209	0.182	0.163	0.148
Adjusted R^2	0.305	0.208	0.181	0.162	0.147

Standard errors in parentheses

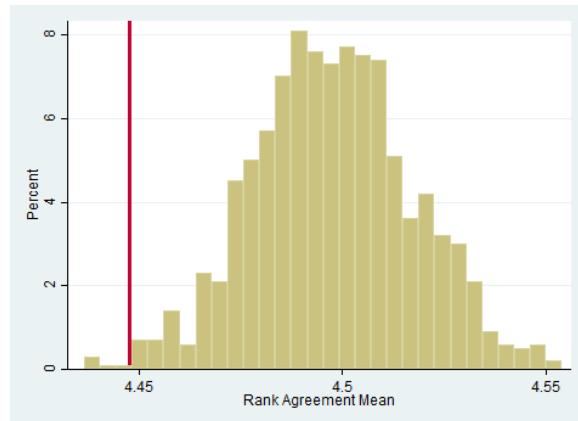
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table presents estimates of a stacked regression of top preference on average preferences of a cadet's first company peers, and a cadet's own initial preference. These last two objects were measured during the first round. Each column corresponds to a different preference round, beginning in Round 2 and ending in Round 6. The data are constructed such that each cadet reported 17 observations, one for each branch. The outcome variable is an indicator for whether the cadet chose the branch as top choice in the final (sixth) round (*First Pref, R6*), the same for the first round (*Own First Pref, R1*), and the leave-out fraction of seniors in the first company that chose that branch as top choice (*% of 2nd Company Peers, R1*). Controls are racial/ethnicity variables (i.e., Black, Hispanic, Other Race), CEER Score, Whole Candidate Score, NCAA Athlete, Prior Enlisted, Parent Served in the Army, and Attended West Point Preparatory School. All columns have year fixed effects. All columns also have Branch FEs, which means that an indicator variable for the branch to which the observation refers is included. Cyber was an option during the third round for one cohort, but it was unavailable as a choice during the first two rounds; it is included in this table for all rounds. Standard errors are clustered at the company-level.

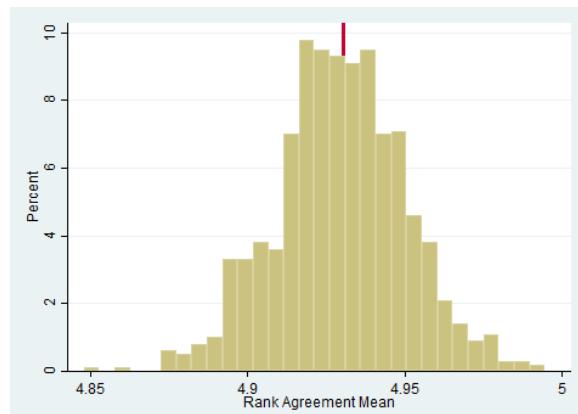
6.3 Figures

Figure 6.3.1: Rank Agreement, Company 1, Round 1 and Round 2

Panel A: Preference Round 1
P-Value: 0.01



Panel B: Preference Round 2
P-Value: 0.48



Notes: These figures show the rank agreement distribution of 1,000 synthetic companies, with the red line indicating the rank agreement from the actual data, as described in Section 1.8. Panel A is for preference round 1, calculated in relation to company 1. Panel B is for preference round 2, also calculated in relation to company 1. P-values are also indicated.

CHAPTER 7

APPENDICES FOR ARE HIGH-QUALITY PHD PROGRAMS AT UNIVERSITIES ASSOCIATED WITH MORE UNDERGRADUATE STUDENTS PURSUING PHD STUDY?

7.1 Data

We make use of several datasets in our analysis. The primary dataset is the SED, a survey administered to new PhD recipients that has a very high response rate.¹ These individual-level data contain information on those who received their PhD by 2016. We consider the PhD recipients who received their BA between 1980 and 2005, where 1980 refers to the academic year 1980-1981.² We drop observations with missing year of BA or PhD receipt as well as those missing BA institution. We construct TTD, used in correcting for truncation (Section 2.4.2), by computing the length of time between BA and PhD receipt. Our measure of TTD thus includes any time pursuing other degrees, working, etc., before completing the doctoral degree.

As our focus is separately on the humanities, life sciences, physical sciences, and social sciences, we consider the number of PhDs earned in each of these fields; Appendix 7.2 describes how these fields are classified. We focus our analysis on doctoral institutions as

¹For example, the response rate in 2016 was 91.8%. See National Science Foundation (2017).

²The SED data indicates the month and year of the completion of both the BA and the PhD. For BAs, we consider completions occurring between September-December of year x-1 and January-August of year x as corresponding to school year x. Because of this, our data includes individuals who graduated in the latter half of 1979. If the month of completion is missing, we use the year variable as the year of completion. For PhDs, we do not attempt to determine the school year, and only use the year of completion.

defined in the Carnegie 2015 classifications, those with Carnegie categories 15 (Doctoral Universities: Highest Research Activity), 16 (Doctoral Universities: Higher Research Activity), and 17 (Doctoral Universities: Moderate Research Activity). We exclude for-profit institutions, institutions that are only graduate schools, and institutions not in the 50 states and Washington DC. We include only institutions that produced at least five PhDs over the sample timeframe. This restriction is made without taking into account TTD and only includes observed PhDs, not those filled in by truncation correction. This is done separately by field, meaning that one institution may be represented in one field and not another. For this reason, the life and physical sciences have a somewhat larger sample than the humanities and social sciences.

We use the IPEDS *Completions Survey* to obtain the number of BA degrees granted by each undergraduate institution each year, by field. See Appendix 7.2 for classification details. We consider only first reported BA major, which is problematic if the student received two (or more) majors in different fields. The number of BA degrees is combined with the SED data on PhDs produced to compute the humanities PhD production rate, described more fully in Section 2.4.3. We also calculate the percentage of total BAs earned in each field from this dataset. The Completions Survey data includes a FICE code (from the older HEGIS survey) for all years and a UnitID code (from the IPEDS survey) for 1987 and after. There are a small number of instances where a FICE code maps to more than one UnitID; this may happen, for example, when there are satellite campuses. If the first year that we observe an institution having more than one UnitID is 1987, we aggregate all UnitIDs together for this institution for all years. Because 1987 is the first year that we observe UnitID, it is likely that the multiple campuses were included in the FICE in 1986

and before. This results in an inflated number of BAs for some institutions. If the first year we observe a FICE mapping to multiple UnitIDs is after 1987, we exclude the UnitIDs associated with institution names that do not appear to correspond to the main campus. In general, we will observe too many BAs for observations where the satellite UnitIDs were not broken out into their own UnitIDs.³

We obtain additional explanatory variables, which are matched at the undergraduate institution-BA year level, from a number of sources. The public institution indicator comes from the SED. We use the IPEDS *Fall Enrollment* dataset to obtain the total number of students, including graduate and professional students. Specifically, we sum full-time students, having a weight of 1, with part-time students, having a weight of 0.4, regardless of institution type or level of student⁴ This dataset is also used to compute the share of students who are undergraduates, the share of undergraduates who are female, and the share of undergraduates who are from underrepresented minority groups.⁵ Each of these is measured as percentage (out of 100). These variables also use the same weights. We divide total students by full-time faculty counts from the IPEDS Fall Staff Survey to compute the student-faculty ratio.⁶

³In one case, the FICE code, which appears to sometimes include satellite campuses and sometimes not, began in 1987; we used one of the several other FICE codes for prior years which correspond to the main campus and which stops in 1986.

⁴The weight of 0.4 for part-time students is very similar to the weight used for undergraduate in the IPEDS “Calculation of FTE Students” (IPEDS, 2018). The weight listed for Public 4-year institutions is .403543 and for private 4-year institutions is .392857. The weights are somewhat further from 0.4 for part-time graduate students: the weights at public 4-year and private 4-year institutions are .361702 and .382059.

⁵Due to problems with subcategories not adding up correctly in early years, when computing certain of these variables, we did not use certain subcategories such as FT or PT UG Unclassified. When computing race, the sum of the race categories were often less than expected; we assume that the race categories do not include the non residence variable.

⁶Full-time faculty counts come from the sum of the “Faculty, full-time men” and “Faculty, full-time women” variables.

Data from the *Delta Cost Project* (Lenihan, 2012) allow us to compute instructional expenditures per student and research expenditures per full-time faculty. The former is obtained by dividing instructional expenditures by full-time students; the latter comes from dividing research expenditures by full-time faculty.⁷

We construct the 75th percentile score variable from the Annual Survey of Colleges *Standard Research Compilation*. We combine the incoming freshmen 75th percentile ACT scores and 75th percentile SAT scores together; we do not use just one of the tests because institutions vary widely in the propensity of their students to take these tests.⁸ This dataset includes the percentage of students with reported SAT scores and reported ACT scores. We weight the ACT score by the percentage of the total that reported ACT scores, and do similarly for SAT scores. For example, if 70% reported a SAT score and 35% reported an ACT score (a student can take both tests), then the SAT will get twice the weight as the ACT. We do this after converting the SAT to be on the ACT scale.⁹ We note that after

⁷Instructional expenditures come from the instruction01 variable—“Expenditures for instruction—current year total.” Research expenditures come from the research01 variable—“Expenditures for research—current year total”.

⁸The score data include two identifiers, ID and FICE. However, there is not always a 1:1 mapping between the two. After cleaning the test score data, in cases where the FICE code maps to multiple IDs, we select one ID to use based first on the one that seems to match the institution name the best, second based on the one that has non-missing test scores, and finally based on the one with the highest test scores. If ID maps to multiple FICE, we use the FICE that is in the main dataset for matching. Finally, there are cases where there is no match between datasets via FICE, but the institution name is associated with a different FICE code in the score data. We manually make these matches. Finally, we manually match a number of cases by institution name if the FICE codes do not match between datasets. We dropped the data from 1983 as it is unreliable. In one case, the campus the ID referred to apparently changed over time; we used all years. In addition, there were no reported percentage taking ACT or SAT for 1984-1987, so we backfilled the 1984 values with the earliest observation before 1993. We then interpolated the percentage taking ACT and SAT.

⁹The composite SAT score, the sum of an institution’s verbal and math 75th percentile scores, was converted to its ACT equivalent using the conversion found at <https://web.archive.org/web/20160130223549/http://act.org/solutions/college-career-readiness/compare-act-sat/>—accessed April 27, 2016. We rounded up to the ACT score if the SAT score was outside one of the listed ranges. For example, if the range was 1290-1320 for a 29 and 1330-1350 for a 30,

conversion and among observations with both scores, the mean ACT is 1.5 points lower than the mean SAT; the standard deviation for both tests is approximately 3.

We obtain doctoral program rankings from the “scholarly quality of program faculty” ranking from the 1995 National Research Council (NRC) doctoral program rankings (Goldberger et al., 1995). The newer, 2000s rankings do not provide point estimates of program rankings and are thus not used (Ostriker et al., 2011). We also do not use the 1980s rankings, which contain a smaller set of institutions (Jones et al., 1982). These rankings have a high correlation with the 1995 rankings.¹⁰ Because we only use the 1995 data, the ranking variables are constant within institution-field. We first classify the different subjects into the four broad fields; classification details appear in Appendix 7.2. For a given subject, such as comparative literature (in the humanities), we construct percentile rankings over all departments, including departments at institutions not in our sample. We then count the number of departments a given institution has across all subjects in the field that are in each of the following percentile ranges: 0-10, 11-25, 26-50, and 51-100.¹¹ Im-

we converted a score of 1325 to 30. We made note of the re-centering of the SAT scores in 1995 and have converted all scores prior to 1995 to their current day equivalent (separately for math and verbal) using the SAT I Individual Score Equivalents conversion table from the CollegeBoard (see <http://research.collegeboard.org/programs/sat/data/equivalence/sat-individual>—accessed April 27, 2016). The chart presented single conversion numbers. We rounded to the higher number. For example, a verbal score of 690 is re-centered to 750 and 700 is re-centered to 760. We re-centered 695 to 760. For cases where there was a test score for SAT, but not ACT, we used the SAT score entirely and vice versa. As there was a great deal of missing data on earlier years of percentage taking each test, we backfill years missing this data with the earliest reported percentages. We also fill in some of the missing score data using linear interpolation.

¹⁰The 1980s ranking include fewer institutions than the 1995 rankings, but for departments that appeared in both rankings (and after re-ranking them to reflect only the departments in both rankings), the correlations of the number of departments at an institution in each rankings interval are: 1 to 10 (.942), 11 to 25 (.888), 26 to 50 (.867), and 51+ (.805). Note that these are not percentiles, but the number in the top 10, etc. If one does not restrict to departments in both, somewhat smaller, but still high correlations are observed: 1 to 10 (.917), 11 to 25 (.842), 26 to 50 (.797), and 51+ (.593).

¹¹Using the set of institutions in the life sciences, correlations within percentile ranges and across fields

portantly, the smaller the percentile, the higher ranked the department. We use percentiles as opposed to the number in the top 10, etc., because there is a large amount of heterogeneity in the number of departments ranked in each subject. For example, many more departments are ranked in English Language and Literature than in French Language and Literature, and a ranking being ranked, e.g., in the 11-25 range may mean something very different for the former subject than the latter.

We face the challenge of missing data. For many of the control variables, we use a linear extrapolation to fill in data for years that are bookended by two other years with non-missing data. When variables do not contain bookended data due to the survey starting after 1980, we backfill using the first possible year of data (for example, we fill in 1980-1983 with the 1984 value).¹² After this process, if an observation is still missing, we code it as 0 and create a missing data dummy that takes on value 1 for missing and 0 otherwise. Universities not included in the 1995 NRC are coded as having zero programs in each of our quality intervals and we do not include a missing dummy variable for such institutions.

range from 0.33 to 0.75, with progressively higher correlations for the 0-10 range than lower ranges in general. These calculations include 0's.

¹²More specifically, we fill in full-time faculty of years 1980-1986 with 1987; test scores of 1980-1983 with 1984; student-faculty ratios of 1980-1986 with 1987; and instructional and research expenditures of 1980-1986 with 1987. We note that this is not ideal as it does not take into account trends, only takes care of cases that have a non-missing value for the first year of data, and results in a lack of variation over these several years. We compute student-faculty ration, instructional expenditures per student, and research expenditures per faculty after backfilling.

7.2 Classification of Subjects to Fields

SED

Our classification of degrees from the SED comes from the subjects listed in Appendix E of the 2011 Paper Questionnaire.¹³ The SED uses fine-grained codes to classify subjects into broader categories. We generally follow these with some exceptions. We list the more general categories, but generally not list the more-specific categories.

Humanities

Codes 700–708, 718–770, 775–799. History; Foreign Languages & Literature; Letters; and Other Humanities.

Life Sciences

Codes 0–299. Agricultural Sciences/Natural Resources; Biological/Biomedical Sciences; Health Sciences.

Physical Sciences

Codes 300–599. Engineering; Computer & Information Sciences; Mathematics; Astronomy; Atmospheric Science & Meteorology; Chemistry; Geological & Earth Sciences; Ocean/Marine Sciences; Physics.

¹³See pages 6 and 7 of https://www.nsf.gov/statistics/srvydoctorates/surveys/srvydoctorates_2011.pdf.

Social Sciences

Codes 650–699, 710, 773. Social Sciences. 710 is History, Science & Technology & Society. 773 is Archaeology.

Other

Codes 600–649, 800–989. Psychology; Research & Administration (Education); Teacher Education; Teaching Fields; Other Education; Business Management/Administration; Communication; and fields not classified elsewhere: Architecture/Environmental Design; Family/Consumer Science/Human Science; Law; Library Science; Parks/Sports/Rec./Leisure/Fitness; Public Administration; Social Work; Theology/Religious Education; Other Fields, NEC.

IPEDS Completions

We classify the subjects, which are at the “Academic Discipline, Detailed”-level, into fields as follows:

Humanities

History; English and Literature; Foreign Language; Other Humanities; Religion and Theology; Arts and Music.

Life Sciences

Agricultural Sciences; Biological Sciences; Medical Sciences; Other Life Sciences.

Physical Sciences

Aerospace Engineering; Chemical Engineering; Civil Engineering; Electrical Engineering; Mechanical Engineering; Materials Engineering; Industrial Engineering; Other Engineering; Astronomy; Chemistry; Physics; Other Physical Sciences; Atmospheric Sciences

Social Sciences

Economics; Political Science and Public Administration; Sociology; Anthropology; Linguistics; History of Science; Area and Ethnic Studies; Other Social Sciences.

Other

Psychology; Science Technologies; Engineering Technologies; Health Technologies, Other Science and Engineering Technologies; Interdisciplinary or Other Studies; Architecture and Environmental Design; Science Education; Mathematics Education; Social Science Education; Other Science/Technical Education; Non-Science Education; Business and Management; Communication and Librarianship; Law; Social Science Professions; Vocational Studies and Home Economics; Other Non-Science or Unknown Disciplines.

NRC

We classify the subjects into fields as follows:

Humanities

Art History; Classics; Comparative Literature; English Language and Literature; French Language and Literature; German Language and Literature; Linguistics; Music; Philosophy; Religion; Spanish and Portuguese Language and Literature; and History.

Life Sciences

Biochemistry and Molecular Biology; Cell and Developmental Biology; Ecology, Evolution, and Behavior; Molecular and General Genetics; Neurosciences; Pharmacology; and Physiology.

Physical Sciences

Aerospace Engineering; Biomedical Engineering; Chemical Engineering; Civil Engineering; Electrical Engineering; Industrial Engineering; Materials Science; Mechanical Engineering; Astrophysics and Astronomy; Chemistry; Computer Sciences; Geosciences; Mathematics; Oceanography; Physics; Statistics and Biostatistics.

Social Sciences

Anthropology; Economics; Geography; Political Science; Psychology; and Sociology.

Other subjects are not listed here. We do not include institutions that are listed as a medical school in assigning the number of departments an institution has in the percentiles of the fields.

CHAPTER 8

BIBLIOGRAPHY

Abowd, John M, Francis Kramarz, and David N Margolis, “High Wage Workers and High Wage Firms,” *Econometrica*, 1999, 67 (2), 251–333.

Arcidiacono, Peter and Sean Nicholson, “Peer Effects in Medical School,” *Journal of Public Economics*, 2005, 89 (2–3), 327–350.

Blau, Francine D. and Lawrence M. Kahn, “The Gender Wage Gap: Extent, Trends, and Explanations,” Working Paper 21913, National Bureau of Economic Research January 2016.

Bowen, William G and Derek Bok, *The Shape of the River: Long-Term Consequences of Considering Race in College and University Admissions*, Princeton University Press, 2016.

Caju, Philip Du, Gábor Kátay, Ana Lamo, Daphne Nicolitsas, Steven Poelhekke et al., “Inter-Industry Wage Differentials In EU Countries: What Do Cross-Country Time Varying Data Add to the Picture?,” *Journal of the European Economic Association*, 2010, 8 (2-3), 478–486.

Carrell, Scott E. and James E. West, “Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors,” *Journal of Political Economy*, 2010, 118 (3), 409–432.

- Carrell, Scott E, Bruce I Sacerdote, and James E West**, “From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation,” *Econometrica*, 2013, 81 (3), 855–882.
- , **Frederick V Malmstrom, and James E West**, “Peer Effects in Academic Cheating,” *Journal of Human Resources*, 2008, 43 (1), 173–207.
- Carrell, Scott E., Marianne E. Page, and James E. West**, “Sex and Science: How Professor Gender Perpetuates the Gender Gap,” *Quarterly Journal of Economics*, 2010, 125 (3), 1101–1144.
- , **Mark Hoekstra, and James E. West**, “Is Poor Fitness Contagious?: Evidence from Randomly Assigned Friends.,” *Journal of Public Economics*, 2011, 95 (7), 657–663.
- , — , and — , “The Impact of College Diversity on Behavior Toward Minorities,” *NBER Working Paper*, 2016, (20940).
- , **Richard L. Fullerton, and James E. West**, “Does Your Cohort Matter? Measuring Peer Effects in College Achievement,” *Journal of Labor Economics*, 2009, 27 (3), 439–464.
- Colarusso, Michael J., Kenneth G. Heckel, David S. Lyle, and William L. Skimmyhorn.**, “Starting Strong: Talent-Based Branching of Newly Commissioned U.S. Army Officers,” *Strategic Studies Institute and U.S. Army War College Press*, 2016, 9.
- Ehrenberg, Ronald G, George H Jakubson, Jeffrey A Groen, Eric So, and Joseph Price**, “Inside the Black Box of Doctoral Education: What Program Characteristics In-

fluence Doctoral Students? Attrition and Graduation Probabilities?,” *Educational Evaluation and Policy Analysis*, 2007, 29 (2), 134–150.

— , **Harriet Zuckerman, Jeffrey A Groen, and Sharon M Brucker**, *Educating Scholars: Doctoral Education in the Humanities*, Princeton University Press, 2009.

Eisenberg, Daniel, Ezra Golberstein, and Janis L. Whitlock, “Peer Effects on Risky Behaviors: New Evidence from College Roommate Assignments,” *Journal of Health Economics*, 2014, 33, 126–138.

Gaviria, Alejandro and Steven Raphael, “School-Based Peer Effects and Juvenile Behavior,” *Review of Economics and Statistics*, 2001, 83 (2), 257–268.

Gittleman, Maury and Brooks Pierce, “Inter-Industry Wage Differentials Job Content and Unobserved Ability,” *ILR Review*, 2011, 64 (2), 356–374.

— and — , “Inter-Industry Compensation Differentials,” *US Bureau of Labor Statistics*, 2012, *Working Paper 453*.

Goldberger, Marvin L, Brendan A. Maher, and Pamela E Flattau, *Research Doctorate Programs in the United States: Continuity and Change*, National Academies Press, 1995.

Goldin, Claudia, “A Grand Gender Convergence: Its Last Chapter,” *The American Economic Review*, 2014, 104 (4), 1091–1119.

Good, Phillip L., *Resampling Methods: A Practical Guide to Data Analysis*, 3rd ed., Boston, Massachusetts: Birkhäuser, 2006.

Goux, Dominique and Eric Maurin, “Persistence of Interindustry Wage Differentials: a Reexamination Using Matched Worker-Firm Panel Data,” *Journal of labor Economics*, 1999, 17 (3), 492–533.

Groen, Jeffrey A, “The Impact of Labor Demand on Time to the Doctorate,” *Education Finance and Policy*, 2016.

— , **George H Jakubson, Ronald G Ehrenberg, Scott Condie, and Albert Y Liu**, “Program Design and Student Outcomes in Graduate Education,” *Economics of Education Review*, 2008, 27 (2), 111–124.

Herbst, Daniel and Alexandre Mas, “Peer Effects on Worker Output in the Laboratory Generalize to the Field,” *Science*, 2015, 350 (6260), 545–549.

IPEDS, “2017-18 Survey Materials, Calculation of FTE students (using fall student headcounts),” <https://surveys.nces.ed.gov/ipeds/VisGlossaryPopup.aspx?idlink=854> 2018. Accessed: 3-19-18.

Jones, Lyle V, Gardner Lindzey, and Porter E Coggeshall, *An Assessment of Research-Doctorate Programs in the United States: Humanities*, National Academies Press, 1982.

Keane, Michael P. and Kenneth I. Wolpin, “The Career Decisions of Young Men,” *Journal of Political Economy*, 1997, 105 (5), 473–522.

Kelly, Inas Rashad, Dhaval M Dave, Jody L Sindelar, and William T Gallo, “The Impact of Early Occupational Choice on Health Behaviors,” *Review of Economics of the Household*, 2014, 12 (4), 737–770.

Knapp, Robert H and Hubert B Goodrich, “The Origins of American Scientists,” *Science*, 1951, 113 (2941), 543–545.

Kofoed, Michael S. and Elizabeth McGovney, “The Effect of Same-Gender and Same-Race Role Models on Occupation Choice: Evidence from Randomly Assigned Mentors at West Point,” *Journal of Human Resources*, 2017, pp. 0416–7838r1.

Lehmann, Erich L. and Joseph P. Romano, *Testing Statistical Hypothesis (Springer Texts in Statistics)*, 3rd ed., Springer, 2005.

Lemke, Robert J, “Generating Doctoral Degree Candidates at Liberal Arts Colleges,” in Ronald G Ehrenberg and C V Kuh, eds., *Doctoral Education and Faculty of the Future*, Ithaca, NY: Cornell University Press, 2009, pp. 93–108.

Lenihan, C, “IPEDS Analytics: Delta Cost Project Database 1987-2010 (NCES 2012-823),” Washington, DC: National Center for Education Statistics. Retrieved from <http://nces.ed.gov/pubsearch>. *America: Its History, Mission and Management*, 2012, pp. 111–122.

Lim, Jaegeum and Jonathan Meer, “The Impact of Teacher-Student Gender Matches: Random Assignment Evidence from South Korea,” *Journal of Human Resources*, 2017, pp. 1215–7585R1.

Lim, Nelson, Jeff Marquis, Kimberly Hall, David Schulker, and Xiaohui Zhuo, “Officer Classification and the Future of Diversity among Senior Military Leaders: A Case Study of the Army ROTC,” *RAND Corporation Technical Report*, 2009.

Lindquist, Matthew J, Joeri Sol, and Mirjam Van Praag, “Why do Entrepreneurial Parents have Entrepreneurial Children?,” *Journal of Labor Economics*, 2015, 33 (2), 269–296.

Lyle, David S., “Estimating and Interpreting Peer and Role Model Effects from Randomly Assigned Social Groups at West Point,” *Review of Economics and Statistics*, 2007, 89 (2), 289–299.

Mani, Anandi and Charles H Mullin, “Choosing the Right Pond: Social Approval and Occupational Choice,” *Journal of Labor Economics*, 2004, 22 (4), 835–861.

Manski, Charles F., “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 1993, 60 (3), 531–542.

Markussen, Simen and Knut Røed, “The Gender Gap in Entrepreneurship—The Role of Peer Effects,” *Journal of Economic Behavior & Organization*, 2017, 134, 356–373.

Marmaros, David and Bruce Sacerdote, “Peer and Social Networks in Job Search,” *European Economic Review*, 2002, 46 (4–5), 870–879.

Mas, Alexandre and Enrico Moretti, “Peers at Work,” *The American Economic Review*, 2009, 99 (1), 112–145.

National Science Foundation, “Survey of Earned Doctorates Survey Description,” <https://www.nsf.gov/statistics/2018/nsf18304/survey-description.cfm> 2017. Accessed: 3-19-18.

Ostriker, Jeremiah P, Charlotte V Kuh, and James A Voytuk, A Data-Based Assessment of Research-Doctorate Programs in the United States, National Academies Press Washington, DC, 2011.

Prenovitz, Sarah J, Gary R Cohen, Ronald G Ehrenberg, and George H Jakubson, “An Evaluation of the Mellon Mays Undergraduate Fellowship’s Effect on PhD production at Non-UNCF Institutions,” *Economics of Education Review*, 2016, 53, 284–295.

Rose, Michael, “Good Deal, Bad Deal? Job Satisfaction in Occupations,” *Work, employment and society*, 2003, 17 (3), 503–530.

Rothstein, Jesse and Cecilia Elena Rouse, “Constrained after College: Student Loans and Early-Career Occupational Choices,” *Journal of Public Economics*, 2011, 95 (1), 149–163.

Roy, A. D., “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 1951, 3 (2), 135–146.

Sacerdote, Bruce I., “Peer Effects with Random Assignment: Results for Dartmouth Roommates,” *Quarterly Journal of Economics*, 2001, 116 (2), 681–704.

— , “Peer Effects in Education: How Might They Work, How Big are They and How Much Do We Know Thus Far?,” *Handbook of the Economics of Education*, 2011, 3, 249–277.

Schapiro, Morton Owen, Michael P O’Malley, and Larry H Litten, “Progression to Graduate School from the “Elite” Colleges and Universities,” *Economics of Education Review*, 1991, 10 (3), 227–244.

Siow, Aloysius, “Occupational Choice under Uncertainty,” *Econometrica*, 1984, 52 (3), 631–645.

Sönmez, Tayfun and Tobias B. Switzer, “Matching With (Branch-of-Choice) Contracts at the United States Military Academy,” *Econometrica*, 2013, 81 (2), 451–488.

Stinebrickner, Ralph and Todd R Stinebrickner, “A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout,” *Review of Economic Studies*, 2013, 81 (1), 426–472.

Tidball, M Elizabeth, “Baccalaureate Origins of Recent Natural Science Doctorates,” *The Journal of Higher Education*, 1986, 57 (6), 606–620.

U.S. News & World Report, “Top 100 Lowest Acceptance Rates,” <https://www.usnews.com/best-colleges/rankings/lowest-acceptance-rate>. Retrieved: 2017-06-10.

Wardynski, Casey, David Stephen Lyle, and Michael J Colarusso, *Towards a US Army Officer Corps Strategy for Success: A Proposed Human Capital Model Focused Upon Talent*, Vol. 1, Strategic Studies Institute, 2009.

Wiswall, Matthew and Basit Zafar, “Determinants of College Major Choice: Identification Using an Information Experiment,” *The Review of Economic Studies*, 2014, 82 (2), 791–824.

Zavodny, Madeline, “Technology and Job Separation Among Young Adults, 1980–98,” *Economic Inquiry*, 2003, 41 (2), 264–278.

Zimmerman, David J, “Peer Effects in Academic Outcomes: Evidence from a Natural Experiment,” *The Review of Economics and Statistics*, 2003, 85 (1), 9–23.

Zöllitz, Ulf and Jan Feld, “The Effect of Peer Gender on Major Choice and Occupational Segregation,” *Unpublished Manuscript*, 2017.