Economics of Education Review 36 (2013) 153–165

Contents lists available at SciVerse ScienceDirect

Economics of Education Review

journal homepage: www.elsevier.com/locate/econedurev

The male–female gap in post-baccalaureate school quality☆

Adam Stevenson *

Department of Economics, University of Michigan, United States

ARTICLE INFO

Article history:
Received 8 February 2012
Received in revised form 10 April 2013
Accepted 17 June 2013

JEL classification:
I21
I23
J16

Keywords:
Graduate school
Professional school
Gender
Ability
Program quality

ABSTRACT

Women are less likely than men to earn degrees from high quality post-baccalaureate programs, and this tendency has been growing over time. I show that, aside from the biomedical sciences, this cannot be explained by changes in the type of program where women tend to earn degrees. Instead, sorting by quality within degree program is the main contributor to the growing gap. Most of this sorting is due to the initial choice in which program type to apply to. No gender differences in selection with respect to ability or program quality arise as students progress through the admissions, enrollment or persistence choices.

1. Introduction

The number of post-baccalaureate (PB) degrees granted in the U.S. has grown explosively in the last half-century, and a major driving force behind this expansion has been the rapid increase of women's representation in higher education. By 2006, women earned around 361,000 PB degrees, compared to the 236,000 earned by men. The number of master's degrees granted to women has grown over 16-fold in the last half-century (274,000 master's degrees to women in 2000, compared to only 17,000 in 1960), with even greater proportional increases at the professional and doctoral levels (Snyder, Tan, & Hoffman, 2006).

This relative female success in PB education is a "homecoming" in the sense that Goldin, Katz, and Kuziemko (2006) use the term. Over the twentieth century, women's attainment of master's and doctoral degrees grew until 1930, fell through the Great Depression, and then began an unbroken increase in the 1970s.[1] Women's relative attainment rates reached an historic high among doctoral degrees by 1990, and approached the historic high among master's degrees in 2000 (Snyder et al., 2006).

While these gains are have brought women's educational attainment to a historic peak, I show that once educational quality enters the analysis, we must conclude that women's relative gains in PB education are overstated. While women invest to a much greater degree than men in the quantity of education, their average investment in quality is substantially smaller. This disproportionate

☆ Portions of this research were written while I wassupported by a grant from the American Educational Research Association which receives funds for its "AERA Grants Program" from the National Science Foundation and the National Center for Education Statistics of the Institute of Education Sciences (U.S. Department of Education) under NSF Grant #DRL-0634035. Opinions reflect those of the author and do not necessarily reflect those of granting agencies.

* Tel.: +1 734 615 8156.
E-mail address: adamstev@umich.edu.

[1] Berelson (1960) presents a detailed history of the early days of formal graduate education.

0272-7757/$ – see front matter
http://dx.doi.org/10.1016/j.econedurev.2013.06.005

concentration of women in low-quality programs is little known, it is an important component of U.S. postsecondary educational growth.

This gender-quality gap exists at every PB award level (master's, professional, or doctoral). Women are 40% less likely than men to earn a degree from a high-ranked doctoral program, but are equally likely to earn a degree from a low-ranked doctoral program. Women are 40% more likely than men to earn a degree from a low-ranked master's degree program. In addition, low-quality programs are the main driver of the increased share of degrees granted to women over the last 20 years, explaining 74% of the growth in the gender-quality gap. After I establish these facts, I begin to document the sources of the gender-quality gap and its growth. Two sets of results come out of this analysis.

The first set of results is negative. I show that the increasing over-representation of women in low-ranked programs can not be explained by changes in gendered patterns of sorting across fields of study. For example, master's programs in education are popular among women, they are growing rapidly, and there are hundreds of low-ranked education programs. But if anything, women are decreasing their representation in these programs relative to men. The same holds for other female-intensive programs like master's programs in nursing or the arts. I also show that there is very little evidence of gender differences in educational continuation choices across quality once students self-sort into applicant pools. Past studies of PB continuation (Millett, 2003; Montgomery, 2002; Mullen, Goyette, & Soares, 2003; Schapiro, O'Malley, & Litten, 1991) find, consistent with my own results, that women are less likely to continue into professional and doctoral studies. Most papers in this literature do not address PB educational quality. The few that do typically focus on narrow fields of study. Montgomery (2002) shows that women are less likely than men to apply to top-tier MBA programs, a result that this paper generalizes across fields and award levels.

The other set of results are constructive. I show that to the extent gender sorting across fields can explain the bottom-driven growth in women's PB attainment, the explanation lies almost entirely in the biomedical fields. MD and PhD bioscience programs are increasingly popular choices among women, they are growing quickly, and their growth is disproportionately in low-ranked programs. In the top 10% of MD programs, women's attainment rates caught up to men's at a rate of 2.9% per year, while in the bottom half of the rankings, they converged at a rate of 4.8% per year. To the extent that educational continuation choices create the gender quality gap, this effect comes entirely from women's tendency to select into the applicant pools of low-pay, low-selectivity programs. While men prefer to apply, controlling for other factors, to programs where expected income after graduation is higher, women exhibit no significant tendency to apply to higher-paying versus lower-paying PB programs.

I also provide preliminary but novel evidence that, holding ability constant, there are positive and large returns to PB quality (between a 1.1 and 1.3 percentage point increase in salary per one percentile increase in the quality rankings) for women, but not men, in professional and doctoral degree programs. Song, Orazem, and Wohlgemuth (2008) investigate the returns to the three major PB award types (masters, professional, and doctoral), and demonstrate substantial the existence ability bias in the estimates, but I am aware of no other paper to estimate the mean value of quality in PB education.

The gender-quality gap is a puzzle. Women are over-represented in the top 10% of their high school classes (Goldin et al., 2006) and graduate from top undergraduate programs in numbers equal to men, with increasingly greater representation below. Choosing to attend a low-quality PB program can have a major impact on future earnings. Since the labor market returns to quality appear to be larger for women than for men, and women who are otherwise similar to men are being placed into lower-ranked programs, there may be substantial economic costs to the gender-quality gap.

This paper proceeds as follows. In Section 3, I establish the basic descriptive facts of the PB gender quality gap and how it has changed over time. I describe the changes in terms of completion rates and in terms of the share of degrees obtained by women. In Section 4, I model student progression towards a PB degree as a series of discrete choices. This allows me to study the relationship between gender, student ability, and program quality as students proceed through their formal schooling. I also present some estimates of how the monetary returns to quality differ by gender across award levels of PB schooling.

2. Data

I use three types of data in this paper: institutional-level data, program-level data, and individual-level data. I present a brief description of each data source here, and give a more detailed description of the data and its limitations in the online appendix. My source of institutional data is the Integrated Postsecondary Education Data System (IPEDS). I use completions data from 1985–2006, disaggregated by gender, award level and field of study.[2] I exclude all students of certification programs (such as the CPA), and I also exclude non-citizens under temporary permission to be in the country (typically an educational visa).[3]

The second type of data is program-level PB quality data. I use the 1994 Study of Research Doctorate Programs (SRDP) (Goldberger, Maher, & Ebert Flattau, 1995) to measure the quality of master's and doctoral programs in arts and sciences. For non-arts-and-sciences PB programs, I use a recent edition of the U.S. News and World Report's "America's Best Graduate Schools" survey (USNWR, 2005). The USNWR universe is almost entirely complementary to the SRDP, focusing on professional and service-based programs. All rankings are stated as centiles within the field and award level. In all cases, rankings are at the institution

[2] I also use IPEDS data on the 75th percentile of undergraduate SAT scores from 1993 as my measure of undergraduate educational quality.

[3] I exclude these individuals primarily for expositional and analytical clarity, and because I can not observe students of this type in the individual-level data.

by award level by field combination: different programs within the same university have different rankings.

This single dimension of "quality" aggregates disparate factors like research productivity, consumption value and labor market returns to students. In some fields, there may be multiple types of differentiation (for example "prestige" versus "vocational preparation"). Even if there is a single unique ranking of programs, the measurements will always be imperfect (Black & Smith, 2006; Brooks, 2005; Dichev, 2001). With these caveats in mind, using a univalent field-specific measure of quality is a substantial contribution to the literature on PB educational continuation.

The third type of data is individual-level data of potential PB students (bachelor's degree holders) from the 1993 Baccalaureate and Beyond (B&B) survey. The B&B sample of applicants is not representative of all PB applicants from their cohort. The B&B only has detailed school information for students who apply within a year of their bachelor's completion. Students who apply within a year of earning a bachelor's degree have higher academic achievement than others. This group of applicants is important because they are among the most highly skilled and ambitious of their undergraduate class. The B&B is the best dataset for my purposes because it is the only nationally representative microsample with any substantial level of detail about PB schooling outcomes.

The B&B collects student SAT and/or ACT test scores. I define "ability" to be the student's test score, measured in standard deviation units above or below the population mean. I discuss some limitations of this measure in the online appendix. In the context of PB admissions, Stevenson (2012) shows that test scores do make women appear to be less appealing candidates in some fields, relative to measures like undergraduate GPA or school characteristics. In my estimates of the admissions choice equation, I control for these other measures of academic promise.

3. The growing gender-quality gap

This section documents the existence of the PB gender-quality gap, and how it has changed over the last thirty years. Using my institutional completions data, in Section 3.1 I present the basic time-trends in terms of women's relative completion rates by award level and program quality. In Section 3.2 I measure the sources of the changes in women's attainment by studying the share of degrees granted across time.

3.1. Attainment rate trends by gender and quality

I classify academic programs by either award level a, or by sixteen broad "degree programs" d.[4] The traditional academic divisions of the humanities, social science, and natural science are a useful baseline for classification. At the doctoral level, I use these categories, breaking up the natural sciences into biomedical and "hard" sciences. At the master's level, each of these academic divisions has a relatively applied portion where the master's is typically terminal, and a relatively academic portion where the master's is not typically terminal. I split the fields in each academic division at the master's level accordingly. I list each degree program with its primary component fields of study in the online appendix, and they can be seen listed in Table 1.

As the number of bachelor's granted increases, the number of subsequent PB degrees granted will grow more or less mechanically. Following the literature (Bowen, Turner, & Witte, 1992; Groen & Rizzo, 2004), I define PB completion rates to be the number of PB degrees granted as a fraction of the number of bachelor's degrees granted in the past.[5] Bachelor's completion numbers are stated as a proportion of the 18-year-old cohort five years previous.

Formally, let C_a^{st} be total number of completions of award a in year t to individuals of gender s. The mean time-to-completion for d is defined as τ_d so that, where B denotes bachelor's degree completion in all fields, I define[6]

$$p_a^{st} \equiv \frac{C_a^{st}}{C_B^{s(t-\tau_d)}}. \tag{1}$$

The completion rate p_a^{st} is most properly interpreted as the probability that a randomly selected individual of gender s who obtained an undergraduate degree in year $t-\tau_d$ will earn award a by year t. Undergraduate attainment trends are interesting primarily as a baseline comparison to PB education.

In Fig. 1, I plot the gender attainment rate ratio p_a^{Ft}/p_a^{Mt} over time to describe relative changes in PB investment by gender. The number represents how much more likely a woman is than a man to earn the given degree, conditional on holding a bachelor's degree. A value greater than one implies that women are more likely, less than one implies men are more likely.

Women's relative attainment has grown steadily in all award levels. This implies gender divergence in rates among bachelor's and master's degrees, where women started out more likely to complete in the 1980s, and convergence in rates among professional and doctoral degrees.[7] Since the relative rates are increasing sharply for bachelor's degrees, the upward-sloping PB lines imply female completions have grown at an even greater rate in these award levels. This (proportional) growth is strongest

[4] I use the term "degree program" to define these overall field groupings (for example, "law programs"), and I use the word "program" to refer to specific institution and degree program combinations (for example, "Harvard law").

[5] The relevant length of time in the past varies by degree program. I take the average enrolled time-to-degree and add one year as an approximate average time between degrees. For most master's degrees I use the average size of the graduating bachelor's degree cohorts 2–4 years prior. For the MBA, I use a lag of 3–4 years. For the JD 4–5 years, for the MD, 5–6 years. For doctoral degrees, I use the average size of the bachelor's degree cohorts 8–10 years prior.

[6] For the bachelor's degree completion rate, substitute H for B and B for d in Eq. (1), where H is the population of 18-year-olds, and let $\tau_B = 5$.

[7] Bootstrapping (over institutions) to recover standard errors implies that we can reject equivalence of men's and women's attainment rates in all years for bachelor's, master's and professional degrees, and for all years prior to 2004 in doctoral programs

Table 1
Changes in the female–male attainment rate ratio, by degree and quality category.

	% of all PB degrees granted in initial year	% Female in initial year	Total growth, % of college grads earning degree	Annualized growth in female–male attainment rate ratio				
				Top 5%	Top 5–10%	Top 10–25%	Top 25–50%	Bottom 50%
Masters	73.3%	53.8%	43.0%	−0.07%	−0.27%	−0.26%	0.24%	1.22%
Liberal Arts	3.6%	51.4%	19.5%	−1.61%	−1.44%	−0.18%	−0.77%	−0.37%
Comm and Media	4.5%	63.4%	20.5%	−0.82%	−0.90%	−1.33%	−0.39%	−0.71%
Social Science	5.1%	56.9%	37.9%	1.45%	0.27%	0.39%	1.61%	1.85%
Public Service	5.1%	62.2%	53.0%	0.73%	2.61%	0.42%	3.55%	2.68%
Engineering	7.3%	20.6%	−1.8%	1.85%	0.37%	0.66%	−0.31%	−0.61%
Physical Science	2.0%	31.2%	−23.1%	1.01%	0.67%	1.71%	1.12%	1.98%
Biological Science	2.9%	43.0%	9.0%	0.37%	−0.57%	0.74%	0.12%	0.53%
Health Science	5.0%	80.1%	98.1%	−0.68%	−0.38%	−1.78%	−1.12%	−1.19%
Education	20.7%	74.8%	64.1%	−0.72%	0.41%	−0.64%	−0.77%	−1.13%
Business	17.1%	34.0%	43.1%	−1.22%	−1.40%	−1.07%	−1.00%	0.66%
Professional	18.8%	38.1%	−3.7%	0.94%	0.69%	1.14%	1.25%	1.91%
JD	10.4%	40.7%	−9.5%	0.40%	−0.07%	0.28%	−0.27%	0.82%
MD	8.4%	33.6%	3.3%	3.37%	2.47%	3.20%	4.18%	4.76%
Doctoral	7.9%	44.3%	7.7%	1.37%	0.65%	0.62%	0.42%	0.62%
Humanities Ed	2.9%	52.6%	−8.2%	0.83%	−1.08%	0.08%	−1.94%	−0.34%
Social Science	2.0%	52.2%	1.2%	3.57%	0.00%	0.01%	0.79%	1.60%
Hard Science	1.6%	18.9%	−17.9%	0.78%	2.37%	3.48%	2.43%	0.51%
Biomedical Science	1.4%	46.1%	79.9%	2.58%	1.72%	0.87%	3.69%	2.51%

Note: the changes (and "initial year") are computed over the same time period as plotted in Figs. 1 and 2, which varies by award level.

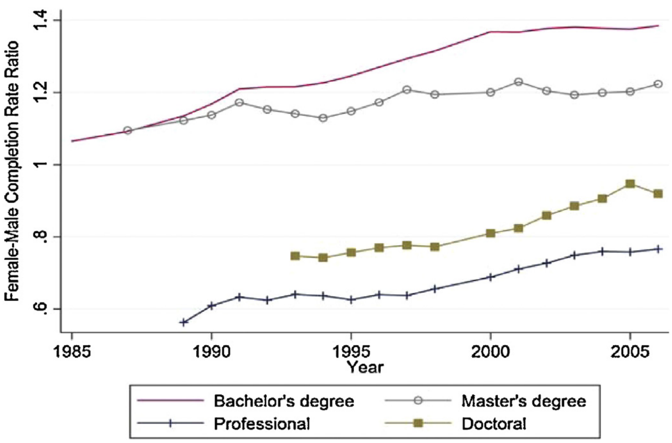

Fig. 1. Gender ratio in attainment rate trends, by award level. *Note*: The attainment rate calculation for synthetic cohorts from the IPEDS is described in the text.

in doctoral degrees, a 40% increase in the female–male attainment rate ratio.

One might expect that as women increasingly earn bachelor's degrees, they become increasingly negatively selected with respect to ability. If so, then even as more women earn PB degrees, the likelihood that a given female bachelor's degree holder earns a PB degree could fall. Fig. 1 shows that this is clearly not so. Among professional and doctoral degrees, the fraction of bachelor's degree holders obtaining these degrees was constant over time. As women increased their attainment likelihoods, the rate fell among men. This stands in contrast to the master's degree, where both male and female attainment rates increased but female rates rose faster. By the most recent cohort, more than 40% of female undergraduates go on to earn a master's degree.

The gender attainment rate ratio by award level and quality is presented Fig. 2. Except among professional degrees, women were always less likely (relative to men) to earn a degree from a high-quality program than a low-quality program. In 2006, women were 40% less likely than men to earn a doctoral degree from a top program, but equally likely to earn a degree from a low-ranked program.[8] Relative to men, women's attainment rates were 60% higher at low-ranked schools than at high ranked

[8] As in Fig. 1, confidence intervals have been suppressed for visual clarity, but all statements of equivalence (or lack of equivalence) of attainment rates are based on bootstrapped standard errors and evaluated on a 95% confidence interval.

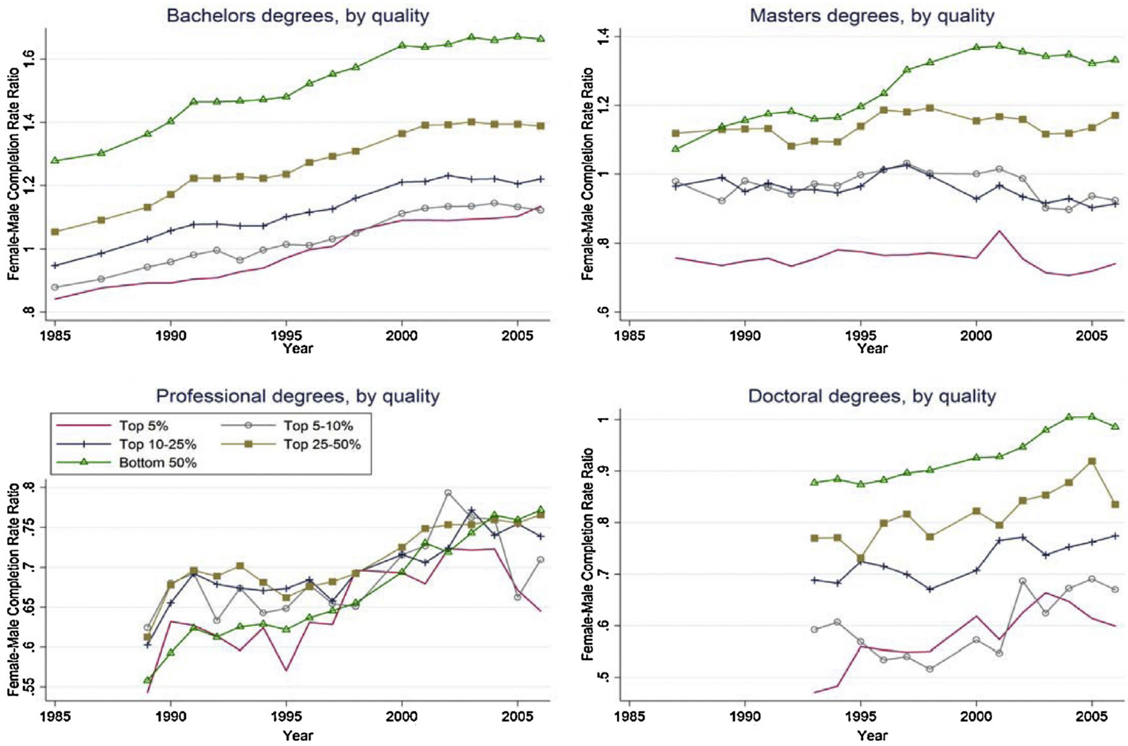

Fig. 2. Gender ratio in attainment rate trends, by award level and quality. *Note*: The attainment rate calculation for synthetic cohorts from the IPEDS is described in the text. Quality categories are mutually exclusive.

schools at the bachelor's and master's degree levels. Women were the majority group at every undergraduate quality level by 2000. Yet in master's degree programs, women's attainment rates were significantly lower than men's in top programs, and were statistically indistinguishable from men's in the top 5–25% of master's programs. Women lost ground somewhat relative to men in these programs over the sample period. Women's completion rates are always significantly higher than men's in lower-ranked master's programs.

Virtually all of women's gains in relative attainment rate among master's degree programs have come from low-quality programs. Women's relative attainment grew more uniformly across quality among doctoral degrees. Women are less likely to enroll in professional degree programs, but in a way that is largely unrelated to program quality.

Differential gender growth by degree program (rather than award level), with a few exceptions can not explain women's relative success in PB education. I illustrate this argument in Table 1. The main component of Table 1 is the description of the (annualized) changes in the female–male attainment rate ratio, by degree program d and quality category q. To put these growth rates in context, I also provide the relative popularity of d in the form of the fraction of all PB degrees granted in that degree program, the fraction female in the first year for which I can calculate p_d^t for the indicated degree, and the overall change in popularity of the degree in the form of total change in attainment likelihood p_d^t over the sample period.

Those degree programs where women's relative attainment is rising fastest among top programs are largely the slowest-growing overall – master's programs in engineering, and doctoral programs in the humanities and social sciences. In the three most popular degree programs – education, business, and law – the changes in female relative attainment are either negative or very small in all quality levels.

The strongest contributors to the large female relative attainment rate growth among master's degree programs are programs in public service (social work and public administration, primarily) and professional and doctoral programs in the biomedical sciences. Public service is a large, quickly-growing field with large relative gains among women. MD programs are slow-growing, but are large, and they are the place where women's relative gains are largest. Doctoral biomedical programs are small, but fast-growing with large relative female gains. In all three of these program types, women's gains are largest in lower-ranked programs.

The two degree programs where women are most concentrated, master's programs in education and health, grew very quickly, but women's relative attainment rates fell across all quality levels, with the fastest decrease at the bottom. The rapid growth of these female-intensive programs can not explain the substantial increases in women's degree attainment at the PB level.

The discussion of Table 1 is an informal way to reconcile aggregate degree program popularity, female clustering in certain fields, and changes in female relative attainment by degree program and quality. In the next subsection, I present a framework to perform this sort of analysis, which allows us to decompose the share of degrees earned by women into within-cell and between-cell changes in completions.

3.2. Degree shares and sorting by field and quality

Another measurement of relative attainment (aside from the attainment rate ratio) is the share of all degrees granted to women in a given year. This statistic makes it easier to measure the contribution of various program types in explaining women's growth in PB education. As a matter of accounting, changes in degree share can be broken into shifts in a program's popularity among women ("within" changes) and shifts in a program's popularity relative to other programs ("between" changes). I present a basic analysis of the change in degree shares by gender and quality here. A fuller decomposition of the change is given in the online appendix, where I show that 72% of the change in the share of degrees granted to women comes from increased "within" variation – a higher proportion of women in each field. Most of this "within" growth is due to entry in bottom-ranked or unranked programs. The faster-than-average growth of female-intensive fields (or shrinkage of male-intensive fields) accounts for only a small portion of women's degree share gains.

Consider the bachelor's degree cohort in time t. Given mean times-to-degree τ_d, I define all PB degrees d earned in time $t+\tau_d$ to be granted to bachelor's degree cohort t. Specify the share of all PB degrees earned by cohort t held by gender s as

$$S^{st} = \frac{C^{s(t+\tau_d)}}{C^{(t+\tau_d)}} = \frac{\sum_{d,q} C_{dq}^{s(t+\tau_d)}}{\sum_{s,d,q} C_{dq}^{s(t+\tau_d)}}. \tag{2}$$

The value S^{st} fraction of all PB degrees granted to individuals of gender $s \in \{M, F\}$, and not the probability that an obtains a certain degree, and so the numbers below are not directly comparable to those in Section 3.1.

The female–male gap in PB degree shares for a given cohort is $S^{Ft} - S^{Mt}$. In the second row of Table 2, the men and women of the 1982 bachelor's degree cohort earned almost equal numbers of PB degrees, and the degree share gap is negative one percentage point for this group. The change in this gap between two periods is denoted $(S^{F2} - S^{M2}) - (S^{F1} - S^{M1})$. By the 1997 bachelor's degree cohort, women earned many more PB degrees, and the change in the share gap between the two cohorts is 19.69 percentage points.

The summations in Eq. (2) are over d and q, allowing us to partially sum over q to describe the contribution of quality to the overall change in degree share. I do this in Table 2. The top row gives the representation of each gender in each cohort (from Census tabulations of 23-year-olds in the indicated year). The next row gives the proportion of all PB degrees earned by each gender in the 1982 and 1997 bachelor's degree cohorts (the first and last cohorts with full IPEDS data). The remaining entries state the fraction of all PB degrees that were granted in that award level, quality, and gender cell. For example, 38.65%

Table 2
Share of all post-baccalaureate degrees granted by sex, for two bachelor's degree cohorts.

	1982		1997		Change in F–M gap
	Female	Male	Female	Male	
Fraction of cohort pop	49.99	50.01	50.74	49.26	1.50
Fraction of all PB degrees	49.50	50.50	59.35	40.65	19.69
Masters	38.65	33.67	47.21	29.30	12.93
Top 5%	1.58	2.06	1.37	1.59	0.25
Top 5–10%	1.86	1.55	1.69	1.25	0.13
Top 10–25%	4.62	4.38	4.63	3.59	0.80
Top 25–50%	5.59	5.02	5.89	4.10	1.22
Bottom 50%	7.21	7.30	10.16	5.76	4.49
Unranked	17.79	13.36	23.47	13.01	6.03
Professional	7.34	12.38	7.87	7.98	4.93
Top 5%	0.31	0.55	0.30	0.33	0.22
Top 5–10%	0.50	0.75	0.41	0.39	0.28
Top 10–25%	1.10	1.65	0.98	1.03	0.50
Top 25–50%	1.64	2.46	1.60	1.58	0.84
Bottom 50%	2.46	4.36	2.74	2.79	1.83
Unranked	1.33	2.61	1.84	1.86	1.27
Doctoral	3.50	4.45	4.26	3.38	1.83
Top 5%	0.20	0.31	0.17	0.18	0.10
Top 5–10%	0.26	0.36	0.23	0.21	0.12
Top 10–25%	0.67	0.91	0.61	0.59	0.27
Top 25–50%	0.71	1.04	0.71	0.68	0.36
Bottom 50%	0.70	0.93	0.69	0.59	0.33
Unranked	0.98	0.90	1.86	1.13	0.66

Note: All numbers are measured in percentage points. The first line is the share of the entire birth cohort that is either male or female in that age group. The next line indicates the fraction of all PB degrees earned males or females from a given cohort. The remainder of the entries are the fraction of all degrees that were granted in that award level, quality, and gender combination in each cohort. The table includes only fields where at least some programs were ranked.

of all PB degrees granted to the 1982 bachelor's degree cohort were master's degrees earned by women. 2.79% of all PB degrees granted to the 1997 cohort to men in bottom-ranked professional degree programs. Since I include unranked programs in this tabulation, Table 2 only includes fields where I have quality data.

The righthand column describes the change over time in the female–male difference in degree shares. In the 1982 cohort, more master's degrees were granted to women, and men earned an almost exactly offsetting greater number of professional degrees. By the 1997 bachelor's cohort, the share gap grew by 19.7 percentage points. Two-thirds of this gap change comes from master's degree growth, and three quarters of this gap change, in turn, comes from low-ranked and unranked programs. Women still earn slightly fewer professional degrees than men, but reversed the share gap among doctoral degrees. Women have a higher share of doctoral degrees than men but lower attainment rates because they are more likely than men to hold bachelor's degrees.

In the online appendix, I decompose the share gap change into "within" and "between" effects. A fuller exposition of the results are presented there, but the main results are as follows. Among professional degrees and doctoral degrees taken as a whole, the increased proportion of degrees going to women comes largely from within-cell growth, accounting for 83% of women's share growth. The within effect also strongly dominates the between effect among low-ranked and unranked master's programs. Among ranked master's degree programs, women's gains are roughly in equal parts due to the secular increase in attainment likelihood and the fact that female-intensive fields like social work, education, and nursing are the fastest growing.

This section documents the existence and growth in the gender-quality gap in PB degree completions. There are at least two important aspects of the trends that merit further scrutiny. First, it is important to understand why programs in the biomedical fields are so appealing to women relative to other fields (Goldin and Katz (2012) explores this question in the context of pharmacy). Second, we must understand why women are (increasingly) less likely to graduate from top programs and more likely to earn degrees from lower-ranked programs. In order to guide future work for this second item, I present a more detailed multivariate analysis of the sources of the gender quality gap for one cohort in the next section.

4. Student selection by ability and quality: the path to a PB degree

In this section, I use the B&B to follow the 1993 bachelor's degree cohort from the completion of undergraduate education, through to graduate degree attainment. The goal is to identify gender differences in selection across quality that are consistent with the gender-quality gap described in Section 3. Previous studies of PB educational continuation rarely consider the dimension of quality, and virtually always treat the continuation decision as a single choice, most often enrollment (Mullen et al., 2003; Schapiro et al., 1991) or completion (Baker, 1998). Occasionally the process is modeled in two steps (Millett, 2003; Montgomery, 2002).

I model the path to a PB degree as having four discrete steps: application, admissions, enrollment, and persistence-to-degree.[9] Each is a place where gender-differential choices according to program quality or student ability could create gender differences in the quality of the program from which students earn their degree. In each step, I present a discrete-choice regression of the influence of gender, program quality, and student ability on the continuation choice. While we should be concerned with selection bias in these sequential estimates, I show in the online appendix that (with one exception noted below) correcting for selection-on-observables in prior steps has no effect on the estimates. This is consistent with the findings of Arcidiacono (2005) at the undergraduate level.

First, I will describe the gender composition by award level and program quality across these four steps. I then show that gender differences in continuation come almost entirely from the initial decision of where to apply. These results should largely be interpreted as being descriptive in nature. Discovering the underlying mechanisms that cause the gender differences at each step is left to future work.

4.1. Gender and the path to a post-baccalaureate degree

From the completion of a bachelor's degree, I assume the path to a PB degree has four steps. First, students face the multinomial choice of whether to apply, and if so, to a program of which award level and quality? After this initial application, a series of binary choices follow. In the admissions step, each program chooses to admit or decline each applicant. Students with admissions offers make an enrollment decision by choosing whether to accept each admissions offer they receive. Finally, given enrollment students make a persistence decision – does she drop out, or continue in school until she earns a degree? In Fig. 3, I plot the fraction female at each of the four steps, by program quality and the highest award level to which the student applied by 1994. The quality categories in this figure are mutually exclusive.

In doctoral programs, there is a clear stratification across quality in terms of gender. Women are substantially under-represented among those in the high-quality programs and over-represented in the low-quality programs. The master's degree data shows a similar pattern. Women are more likely to complete degrees from low-quality programs than from high-quality programs, representing 45% of the applicants to top master's programs, but 60% of the applicants to bottom-ranked programs. Women are almost equally represented across quality among professional degree applicants. Women are a smaller fraction of the graduates from low-quality programs and a greater fraction of the graduates from high-quality programs. This is largely a composition effect. Women disproportionately select out of law programs as they progress, but men disproportionately select out of

[9] Manski and Wise (1983) and Arcidiacono (2005) model the same four-step process at the undergraduate level.

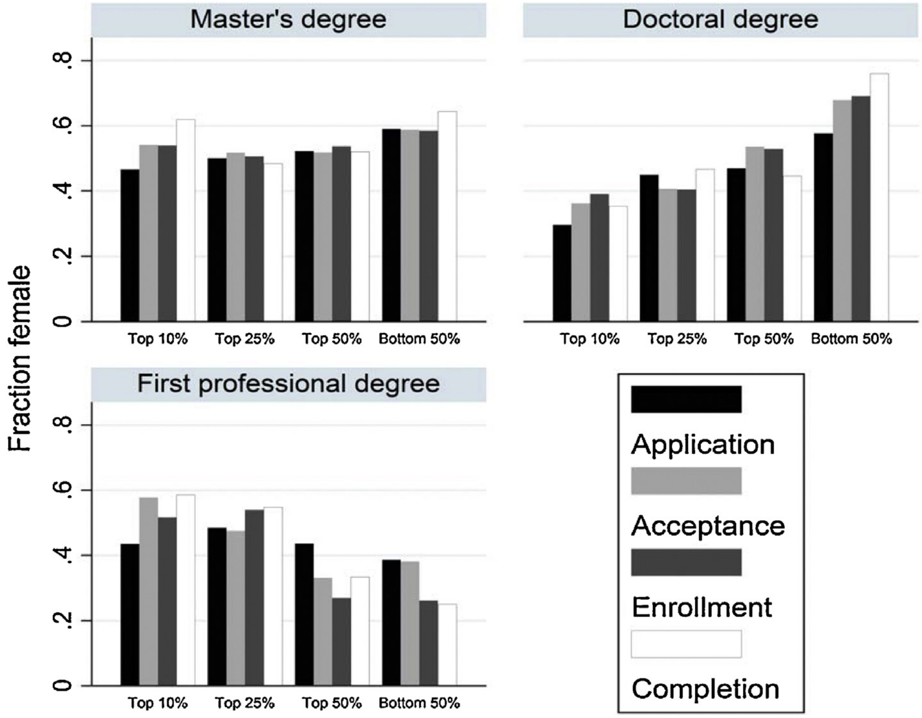

Fig. 3. Fraction female over progression to a post-baccalaureate degree, by award level and quality. *Note*: Fraction female by highest award level applied and program quality among those who applied, were accepted, and enrolled by 1994, and of these, who completed by 2003. Quality categories are mutually exclusive.

medical programs. Top medical programs are larger than top law programs, and vice-versa at the bottom, generating the observed pattern.

The representation of women does not appear to change substantially as a function of who accepts admissions offers and enrolls. We can not reject the hypothesis that, given award level of quality, gender composition does not change along the path to a degree.

4.2. Ability, quality, and the path to a post-baccalaureate degree

I show in this section that there are strong gender differences in the pattern of application choices between men and women. Selection into the applicant pools of lower-income, less selective programs reinforces the over-representation of women among low-ability bachelor's degree holders. In steps other than the application choice, gender differences are small and statistically insignificant. This is largely consistent with the previous literature, which has considered PB educational continuation in a relatively piecemeal manner.

In each step along the path to a PB degree, there are two main effects consistent with a gender quality gap. First, if women systematically apply to lower-ranked programs than men, the gap would arise. Second, taking applications as an example, if low-ability females are more likely to apply to PB education than low-ability men, and this difference is greater than among high-ability students, then a gender-quality gap may emerge if lower-ability students tend to apply to lower-quality programs. As such, the coefficients associated with the interaction of gender and either ability or quality are the focus of this section. Program quality and award level are modeled as a choice variables in the applications choice, and are taken as given throughout the remainder of a student's progress towards a degree.

Each step involves a unique set of variables, constraints, and heuristics that are potentially correlated with gender, ability, or program quality. To control for the effects of these shifters, in each step I include a separate set of relevant right-hand-side X^j variables (where j indexes the four steps). I describe the precise elements of each X^j as I discuss the regression results.[10]

By the nature of the B&B, every individual in the dataset is "at risk" to apply to a PB program. Analytically, the initial application choice is the most complicated stage of the four steps to a PB degree. The bachelor's degree holder must decide whether to continue their educational career or to opt out into the job market. Conditional on deciding to apply to school, they must make a choice of award level and program quality. The estimation problem is difficult

[10] In the selection-corrected models in the online appendix, these X^j variables serve an important role as exclusion restrictions for each step along the path.

because it is not clear *ex ante* whether, for example, a mid-range professional degree program is preferable to a top master's degree program.

To estimate the parameters of this choice over where and whether to apply, I model the application decision as a multinomial choice over award level and quality combinations.[11] This approach differs from that typically taken in the literature (and that I take in the other steps), which models applications choices as a function of student characteristics, rather than choice characteristics.

I take the twelve categories presented in Fig. 3 (three degree levels by four quality categories) plus the outside option of no school and define each as a "program type" *aq* over which students choose. Students draw value from program characteristics, and this value is potentially heterogeneous according to the observable student characteristics Z_i and unobservable characteristics ν_i of the student.

Each *aq* choice is described by a vector of the mean characteristics of the degree programs within the category, X^1_{aq}.[12] Individual *i*'s expected utility from earning a degree from a program in category *aq* is given by

$$U_i^{aq} = X^1_{aq}(\theta' + \theta_z' Z_i + \theta_\nu' \nu_i) + \eta_i{}^{aq} \tag{3}$$

where θ, θ_z, and θ_ν are vectors of parameters to be estimated. With the further assumptions that ν_i is the same length as $X^1_{aq}, \nu_i \sim N(0, \Sigma)$, and that η_i^{aq} is distributed according to the type-II extreme value distribution, this framework is the standard random-coefficients, or "mixed", logit (Revelt & Train, 1998). An important part of this specification is the covariance matrix Σ. Each individual draws a preference shock for each program characteristic. This allows for realistic substitution patterns across *aq* choices. For example, a student with an unobserved preference shock in favor of shorter programs will be more likely to choose a master's degree program, conditional on all other factors. The mixed logit does not rely on the assumption of independence of irrelevant alternatives, as does the simple multinomial logit. In addition, Σ need not be diagonal, so these preference shocks may be correlated.

Given that the student applied, an admissions committee decides whether to accept or reject the applicant. Suppose that the applicant is admitted whenever the latent variable [13]

$$y_i^j = f^j(Z_i, X_i^j) + \varepsilon_i^j > 0. \tag{4}$$

Once the admissions equation is estimated (along with the likelihood of acceptance), we proceed to the enrollment step of whether the student enrolls in a program given the presence of an admissions offer. Since I model admissions, enrollment, and persistence each with a probit model, the general Eq. (4) applies for all steps $j \geq 2$.

While the framework above controls for sorting across award levels and quality categories, another concern is the presence of heterogeneity with respect to program quality and gender sorting as it relates to field-of-study. For example, most primary and secondary school teachers have a union-mandated item in their contract that guarantees a substantial salary increase once the teacher earns a master's degree, and many states require teachers to earn a master's degree. These regulations are not indexed to the quality of the degree program. The incentive to enroll in a high-quality program is likely to be small (or even negative) in this situation. The programs where these conditions tend to exist are exactly those where females tend to cluster – the "applied" master's degree programs such as education, public service, and health science. I control directly for the fixed effects of field-of-study in the admissions, enrollment and persistence decisions steps. In the application decision, I control for degree program indirectly, via its influence on income expectations.

Before any individual sets along the path to a PB degree, we have the pool of potential applicants, all bachelor's degree holders. In Fig. 4, I plot the frequency of males and females across the ability distribution. The sample is 54.9% female, and the average ability of women, 0.21, is significantly lower than that of men, 0.5. A Kolmogorov–Smirnov test rejects equivalence of the ability distributions across gender both taken across the entire distribution, and when the sample is split above and below the threshold of ability = 0.5. Summary statistics on the evolution of the mean gender ability gap along the path to a PB degree are given in the online appendix.

4.2.1. Applications decisions

First, I estimate the application choice, Eq. (3). The vector X_1^{aq} includes data on the fraction female of enrolled students, the average ability level of enrolled students, the average net annual tuition paid (tuition minus all financial aid), and the average program length, measured in years. It also contains an estimate of the student's expected income from choice *aq*. To estimate this value, I use the 2003 wave of the B&B to regress income on cubic polynomials in ability and age, racial indicators, plus a set of indicators interacting the *aq* level of an individual's degree attainment, undergraduate field of study, and gender.[14]

Table 3 presents the estimated coefficient of Eq. (3), with bootstrapped standard errors. One column presents the fixed parameters θ, θ_x, and θ_ν, and the other column gives the diagonal elements of Σ. The off-diagonal elements were estimated, but are not reported here. These elements are not statistically significant, but some are quite large, particularly the correlation between the random parameters of fraction female, program average ability and program length. In terms of the distribution of

[11] The following exposition closely follows that of Ackerberg, Benkard, Berry, and Pakes (2007).

[12] For the outside option of choosing not to apply, net tuition and program length are set to zero. Income expectations, fraction female, and mean ability are estimated from the pool of people who never apply to PB education.

[13] There is reason to expect that ε_i^2 will be correlated with η_i^{aq} from Eq. (3), and therefore the likelihood of acceptance is correlated with the likelihood of being observed as an applicant. Selection of this nature, as well as between the steps of enrollment and persistence, are addressed in the online appendix.

[14] Since all B&B participants earned their bachelor's degrees in 1993, I can not separately identify the effect of potential experience and educational choice in the data.

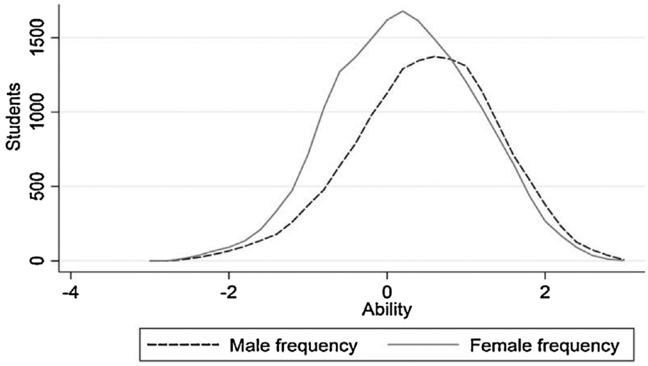

Fig. 4. Ability frequency in the baccalaureate and beyond. *Note*: Data is for all B&B students with ability data surveyed in 1994.

Table 3
Mixed logit estimation of highest level of PB application by 1994.

	Fixed coefficient	S.D. of random coefficient
Fraction female at A-by-Q level (FAQ)	6.621**	2.363
	(1.239)	(1.322)
Female × FAQ	−0.277	
	(0.774)	
Ability × FAQ	−2.641**	
	(0.561)	
Mean ability at A-by-Q level (AAQ)	−3.595**	0.223
	(0.551)	(0.636)
Female × AAQ	−0.872**	
	(0.279)	
Ability × AAQ	1.277**	
	(0.202)	
Mean net tuition at A-by-Q level (TAQ)	0.232	0.110
	(0.175)	(0.158)
Female × TAQ	0.066	
	(0.035)	
Ability × TAQ	−0.015	
	(0.036)	
Mean program length at A-by-Q level (LAQ)	−0.753**	0.031
	(0.073)	(0.251)
Female × LAQ	0.112**	
	(0.040)	
Ability × LAQ	0.029	
	(0.032)	
Expected income from completion at A-by-Q level (IAQ)	0.091**	−0.007
	(0.030)	(0.006)
Female × IAQ	−0.076*	
	(0.034)	
Ability × IAQ	−0.026	
	(0.016)	
N	6472	

Note: Mixed logit regression with five unobservably heterogeneous parameters. The unobservable preference to parameters are allowed are distributed according to the multivariate normal $N(0, \sum)$. The diagonal elements of $\sum$ are reported above, while the off-diagonal covariances have been estimated, but are omitted from the table. Standard errors are generated by bootstrap.

* $p < 0.05$.

** $p < 0.01$.

θ_ν, only the value associated with fraction female exhibits any substantial heterogeneity, but even here the standard deviation is marginally insignificant.

The influence of net tuition, program length, and expected income are all statistically significant in the expected direction, such that lower-cost, higher-benefit programs are preferred. Students with higher ability attach less weight to monetary variables in their decision of where to apply. More important in terms of the gender-quality gap, the interaction between the female indicator and program characteristics is significant for a few variables. Women are more likely than men to avoid programs with high-ability peers – the coefficient associated with mean ability is −3.6 for men but −4.47 for women. This effect primarily tends to select women out of PB education. The marginal effect of this interaction is to reduce the likelihood that women to apply to top master's programs (relative to observably equivalent men) by 0.73 percentage points, from a base probability of 2.05 percent. It reduces the likelihood of applying to top professional programs by 0.32 percentage points, from a base likelihood of 0.56. Women do not choose programs with higher expected income (its effect size is insignificant and close to zero for women), while men do. Women disproportionately select away from the applicant pools of programs with higher-ability students, particularly the professional degrees. These results confirm more generally those of Montgomery (2002), who restricts his analysis to potential MBA applicants.

This first step determines the award level and program quality that the student has chosen to progress through. For the rest of the analysis, I take these characteristics to be fixed. Choice decisions are modeled separately (allowed to vary freely) according to award level. All further choice regressions will include quality as a continuous control variable, to investigate whether women are disproportionately into or out of high-quality programs, conditional on their sunk application choice.

4.2.2. The admissions decision

In the first panel of Table 4, I present the results of the probit estimates of gender differences in admissions likelihood. These estimates are performed separately by

Table 4
Progress through post-baccalaureate education, by award level.

	Admissions offer to top choice			Enrollment, given admissions			Completion, given enrollment		
	Master's	Professional	Doctoral	Master's	Professional	Doctoral	Master's	Professional	Doctoral
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	−0.027	−0.145	0.927	0.283	−0.521	1.259	0.593**	−0.05	1.199
	(0.216)	(0.386)	(0.701)	(0.190)	(0.421)	(0.850)	(0.227)	(0.539)	(0.977)
Ability	0.118	0.221	0.282	0.016	0.298	0.444*	0.190*	−0.383	0.111
	(0.080)	(0.161)	(0.174)	(0.082)	(0.197)	(0.224)	(0.094)	(0.269)	(0.231)
Ability × Female	0.199*	−0.063	0.000	0.057	−0.081	−0.403	−0.072	0.405	−0.136
	(0.101)	(0.205)	(0.236)	(0.107)	(0.258)	(0.305)	(0.123)	(0.386)	(0.330)
Program quality	−0.015**	−0.012**	−0.001	0.002	−0.006	0.005	0.004	0.007	0.006
	(0.003)	(0.004)	(0.007)	(0.003)	(0.005)	(0.008)	(0.003)	(0.006)	(0.009)
Quality × female	0.001	0.004	−0.008	−0.005	0.004	−0.01	−0.006	−0.005	−0.016
	(0.004)	(0.006)	(0.009)	(0.003)	(0.007)	(0.012)	(0.004)	(0.009)	(0.013)
Admissions shifters?	Yes	Yes	Yes	No	No	No	No	No	No
Enrollment shifters?	No	No	No	Yes	Yes	Yes	No	No	No
Persistence shifters?	No	No	No	No	No	No	Yes	Yes	Yes
N	1155	315	186	954	229	146	659	153	116
Pseudo-R^2	0.108	0.079	0.087	0.072	0.118	0.241	0.087	0.124	0.124

Note: Each column is a separate probit regression. Acceptance regressions model an admissions offer in 1994 into the student's first-choice program, given application. Enrollment regressions model enrollment by 1994, given any acceptance. Completions regressions model completion of some PB degree by 2003, given any enrollment by 1994. Admissions shifters include the applicant's undergraduate GPA, undergraduate school quality, an interaction between the GPA and undergraduate quality, and an indicator for the presence of honors on the student's undergraduate transcript. Enrollment shifters include the presence of a financial aid offer, the quantity of undergraduate debt, parent's income, whether the student is married, and whether the student has children. Completions shifters are the size of the student's financial aid package, whether the student was working or enrolled part time while in school, and indicators for whether the student married, divorced, or had their first child after enrollment. All regressions control for field-of-study. Standard errors are in parenthesis.

* $p < 0.05$.

** $p < 0.01$.

award level, and allow for the possibility that programs of differing quality attach different weights to applicant gender. Each follows a specification following Eq. (4), where I assume that ε_i^2 is normally distributed, so that I can estimate y_i^2 by probit regression.

Included among the X_i^2 variables are undergraduate GPA, student undergraduate program quality, and whether any honors are reported on the student's undergraduate transcript, since these variables are clearly observable and presumably important to admissions committees.[15] The vector Z_i includes student gender and ability, the program's quality, and gender interactions with student ability and program quality. The dependent variable of the probit is an indicator variable that takes the value of one when the applicant was admitted into her specific first-choice program.

The admissions step is particularly interesting for two reasons. First, it is the only step beyond the initial application choice where I find any significant gender difference, controlling for the relevant shifters. In column (1) of Table 4, we see that high-ability women are significantly and substantially more likely to draw an admissions offer than are high-ability men. Otherwise, no gender differences in admissions are apparent. The second interesting item is that admissions is the one step along path to a PB degree where the selection correction has any "bite". As seen in Appendix Table A3, this coefficient is insignificant in the selection-corrected model. Gender differences in sorting into applicant pools tend to place high-ability women into less selective programs, thereby inflating Table 4's estimate of their appeal to admissions committees.

Among the "applications shifters" X_i^2, the only covariate that contributes significantly to drawing an admissions offer is being the recipient of undergraduate honors, and this effect is only present at the master's degree level.

4.2.3. The enrollment and persistence decisions

Once the student draws an admissions offer, the student must decide whether she expects the costs of attendance to be worth the expected benefits. When and if enrollment occurs, the student must then decide whether to persist through the program, or to drop out. In the enrollment regressions, I include in X_i^3 covariates that may shift a student's willingness-to-pay as of the date of acceptance, relative to the characteristics and expectations modeled in the application choice. I control for family resources and constraints by including parental household income, and whether the student is married or has children. I control for the applicant's financial constraints by including a continuous undergraduate debt variable, and an indicator for whether the offering program included any financial aid as part of the admissions offer.[16] Once enrolled, I assume that the X_i^4 variables include shocks to

[15] I also interact GPA and undergraduate quality in the regression to allow for the possibility that programs weight GPAs differently, depending on the selectivity of the school.

[16] The B&B only asks about the size of financial aid for those who actually enroll. This is variation I will exploit in the completion step.

resources. Time is a particularly relevant resource here, and so I control for academic choices that may delay graduation, including part-time enrollment or if the student changed field of study while enrolled in PB education. I also control for shocks to family status, including entry or exit from marriage, and the arrival of a child. Finally, I control for the total amount of annual financial aid received by the student.

In both steps, the main result is that there are no gender-differential effects. In fact, there are no gender differences detected whatsoever, except for the fact that women are significantly more likely to complete their master's degree, conditional on enrollment. This is consistent with Schapiro et al. (1991) in their post-1985 sample and in Mullen et al. (2003), who also use the B&B, but do not condition on selection though past educational continuation choices.

Women are substantially less likely than men to continue into high-ranked PB programs. The pool of women completing undergraduate education is larger and more negatively selected than the pool of men. Still, women tend to apply to lower-ranked, less selective programs than men, all else equal. This largely determines the gender-quality gap. After this initial negative selection into applicant pools with respect to quality, there are no further substantial forces that select women, relative to men, out of high-ranked programs.

4.3. Post-graduation

Finally, while it is beyond the scope of this paper to engage in a full (and causal) investigation of the monetary value of quality at the PB level, it is nonetheless important here to document the relationship between gender, quality, and income. If women are opting out of high-quality PB programs, what is the economic cost of doing so?

A table of simple means of income across gender, award level, and quality is given in the online appendix. In Table 5, I present the results of an OLS regression of log earnings on PB award level plus its interactions with quality and gender. In addition, I control for potential experience (the time since the last date of educational enrollment), whether the individual was unemployed in 2003, and degree program *d*. Most importantly here, I control for ability. Song et al. (2008) show that there is negative selection according to ability into PB education, particularly with respect to master's programs.

Column 1 controls only for gender, award level, quality, and their interactions. In this base specification, program quality is significantly positively related to earnings for men in master's and professional degree programs. Once I control for ability in column 2, the returns to quality disappear for men and women alike in master's degree programs. The returns to quality in professional degrees remain, and a significant return to quality appears for women in doctoral programs. Controlling further for labor market characteristics and degree program, the positive return to quality for women remains in professional and doctoral degrees, but disappears everywhere for men. These returns are large. A one-percentile move up the doctoral program quality ranking is associated with a 1.3 percent increase in earnings, and only slightly less in professional degree programs. The change in parameters between columns 2 and 3 is almost entirely due to the inclusion of degree program indicators. In this young and highly-skilled sample, selection into the labor market is not a major issue. The unemployment rate among master's degree holders is 2.6% for men and 2.3% for women, and is lower still among professional and doctoral degree holders.

Table 5
Log earnings as a function of award level, program quality, and gender.

	(1)	(2)	(3)
Female	−1.0505**	−1.0365**	−1.0138**
	(0.0593)	(0.0591)	(0.0588)
Ability		0.0797**	0.1045**
		(0.0304)	(0.0321)
Master's × quality × male	0.0014	0.0008	−0.0007
	(0.0013)	(0.0014)	(0.0017)
Master's × quality × female	0.0036	0.003	0.0029
	(0.0021)	(0.0021)	(0.0021)
Professional × quality × male	0.0054**	0.0043*	0.0013
	(0.0016)	(0.0017)	(0.0024)
Professional × quality × female	0.0164**	0.0151**	0.0119**
	(0.0022)	(0.0022)	(0.0028)
Doctoral × quality × male	0.0018	0.0005	0.004
	(0.0014)	(0.0015)	(0.0032)
Doctoral × quality × female	0.0091*	0.0082*	0.0132**
	(0.0042)	(0.0041)	(0.0043)
Labor market variables?	No	No	Yes
Degree program dummies?	No	No	Yes
N		5029	

Note: Dependent variable is the log of labor market earnings in 2003. Quality is measured in field-level percentiles. Labor market variables include an indicator for whether an individual was unemployed in 2003, and the level of an individual's potential experience since the completion of schooling. Degree programs dummies include 16 indicators for field-of-study at the PB level.
* $p < 0.05$.
** $p < 0.01$.

5. Conclusion

There is a substantial gender-quality gap in PB education: women are much more likely than men to enroll in low-quality programs, and this gap has grown over the last 30 years, particularly at the master's degree level. This paper provides an initial investigation into the sources of the post-baccalaureate gender-quality gap.

Low-quality PB programs are expanding quickly, and women are increasing their relative attainment rates (and share of degrees) most quickly in these programs. There is little systematic research on programs of this type, which are clearly an important contributor to the expansion of the system of higher education, in addition to the gender inequalities within that system.

Women select disproportionately into the applicant pools of low-ranked programs. This suggests that intervention in PB admissions or after may be ineffective or counter-productive. On the other hand, understanding the nature and causes of the selection of women into the applicant pools of lower-quality programs than men is important in understanding (and potentially remedying) the PB gender-quality gap. It could be that geographical or familial constraints are more salient to women than to men. This self-sorting could be a rational response to expectations of future discrimination or workplace inequalities. The characteristics of these low-quality programs may be more appealing. If women are more likely to choose part-time schooling or have less leisure time outside of school (to allocate to studying, for example), the presence of these non-academic amenities in low-ranked programs may create the observed self-sorting into the applicant pools of lower-quality programs.

Appendix A. Supplementary Data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.econedurev.2013.06.005.

References

Ackerberg, D., Benkard, C. L., Berry, S., & Pakes, A. (2007). Econometric tools for analyzing market outcomes. In (Series Ed.) & J. J. Heckman, & E. Leamer (Vol. Eds.), *Handbook of econometrics* pp. 4171–4276). (Vol. 6, Pt A, Elsevier).

Arcidiacono, P. (2005). Affirmative action in higher education: How do admission and financial aid rules affect future earnings? *Econometrica, 73*(5), 1477–1524.

Baker, J. G. (1998). Gender, race, and Ph.D. completion in natural sciences and engineering. *Economics of Education Review, 17*(2), 179–188.

Berelson, B. (1960). *Graduate education in the United States*. New York: McGraw-Hill.

Black, D. A., & Smith, J. A. (2006). Estimating the returns to college quality with multiple proxies for quality. *Journal of Labor Economics, 24*, 701–728.

Bowen, W. G., Turner, S. E., & Witte, M. L. (1992). The B.A.-Ph.D. nexus. *Journal of Higher Education, 63*(1), 65–86.

Brooks, R. (2005). Measuring university quality. *The Review of Higher Education, 29*(1), 1–21.

Dichev, I. (2001). News or noise? Estimating the noise in the U.S. news university rankings. *Research in Higher Education, 42*(3), 237–266.

Goldberger, M. L., Maher, B. A., & Ebert Flattau, P. (1995). *Research doctorate programs in the United States: Continuity and change*. Washington, DC: Committee for the Study of Research-Doctorate Programs in the United States, National Research Council, National Academies Press.

Goldin, C., & Katz, L. F. (2012). *The most egalitarian of all professions: Pharmacy and the evolution of a family-friendly occupation*. Cambridge, MA: NBER Working paper 18410, National Bureau of Economic Research.

Goldin, C., Katz, L. F., & Kuziemko, I. (2006). The homecoming of American college women: The reversal of the college gender gap. *Journal of Economic Perspectives, 20*(4), 133–156.

Groen, J. A., & Rizzo, M. J. (2004). *The changing composition of American-citizen PhDs*. Ithaca, NY: Cornell Higher Education Research Institute Working Paper WP48, Cornell University.

Manski, C. F., & Wise, D. A. (1983). *College choice in America*. Cambridge, MA: Harvard University Press.

Millett, C. M. (2003). How undergraduate loan debt affects application and enrollment in graduate or first professional school. *The Journal of Higher Education, 74*(4), 386–427.

Montgomery, M. (2002). A nested logit model of the choice of a graduate business school. *Economics of Education Review, 21*, 471–480.

Mullen, A. L., Goyette, K. A., & Soares, J. A. (2003). Who goes to graduate school? Social and academic correlates of educational continuation after college. *Sociology of Education, 76*(2), 143–169.

Revelt, D., & Train, K. (1998). Mixed logit with repeated choices: Households' choices of appliance efficiency level. *Review of Economics and Statistics, 80*(4), 647–657.

Schapiro, M. O., O'Malley, M. P., & Litten, L. H. (1991). Progression to graduate school from the "elite" colleges and universities. *Economics of Education Review, 10*(3), 221–244.

Snyder, T. D., Tan, A. G., & Hoffman, C. M. (2006). *Digest of education statistics 2005 (NCES 2006-030)*. Washington, DC: US Department of Education National Center for Educational Statistics Government Printing Office.

Song, M., Orazem, P. F., & Wohlgemuth, D. (2008). The role of mathematical and verbal skills on the return to graduate and professional education. *Economics of Education Review, 27*(6), 664–675.

Stevenson, A. (2012). *Admissions policy and the male–female gap in post-baccalaureate school quality. Working paper*. University of Michigan.

U.S. News and World Report. (2005). *America's best graduate schools, 2006*. Washington, DC: US News & World Report.