

Revised Draft
September 1, 2006

The Graduate Education Initiative: Description and Preliminary Findings

by

Ronald G. Ehrenberg, Harriet Zuckerman, Jeffrey Groen, and Sharon M. Brucker*

* Zuckerman and Brucker are with the Andrew W. Mellon Foundation. Groen is at the U.S. Bureau of Labor Statistics and Ehrenberg is with Cornell University. This paper summarizes some of the material that will appear in *The Education of Scholars: The Effectiveness and Quality of Doctoral Programs in the Humanities*. The views expressed here are solely our own.

I. The Graduate Education Initiative

In 1991 the Andrew W. Mellon Foundation launched the Graduate Education Initiative (hereafter GEI) to improve the structure and organization of PhD programs in the humanities and social sciences. Such changes were seen as necessary to combat high rates of student attrition and long times-to-degree in these programs. While attrition and time-to-degree were deemed to be important in and of themselves, and of great significance to degree seekers, they were also seen more broadly as indicators of the effectiveness of graduate programs. Several characteristics of doctoral programs were earmarked as contributing to high attrition and long degree time, including: unclear expectations, a proliferation of courses, elaborate and sometimes conflicting requirements, intermittent supervision, epistemological disagreements on fundamentals and not least, inadequate funding. Projections that faculty shortages would occur in the late 1990s in the humanities made the goals of reducing student attrition and time-to-degree particularly timely if an adequate number of PhDs were to be available.¹

This was far from the first such effort to reduce times-to degree-and rates of attrition. Earlier programs, which provided grants in aid to individual students or to graduate schools to distribute as they saw fit, had failed conspicuously.² Based on data which showed that there were marked differences among departments and on a great deal of experience on the ground, the architects of the GEI determined that to improve graduate education would require departments to make changes in their PhD programs. As such, the Foundation shifted much of its support for doctoral education, which had

¹ Bowen and Sosa (1989)

² Bowen and Rudenstine (1992)

previously gone directly to students, to block grants that would be awarded to departments within major universities.

Thus ten universities – Berkeley, Chicago, Columbia, Cornell, Harvard, Michigan, Pennsylvania, Princeton, Stanford and Yale – were invited to nominate 4 to 6 departments at their universities to be recipients of GEI awards. These universities had previously attracted the largest number of fellowship winners of the Mellon Foundation’s portable doctoral dissertation awards.³ To be eligible for participation and funding, each department had to develop a plan to improve its doctoral program that would be consistent with the objectives of the Foundation. Departments were encouraged to review their curricula, examinations, advising, and official timetables with an eye towards facilitating timely degree completion and reducing attrition (especially late attrition), while maintaining or increasing the quality of doctoral training they provided. There was no requirement that the departments be in need of help – that is, that they had low completion rates and long times to degree – or that they had already made substantial progress in these respects. Universities made their own selections with the result that participating departments had a variety of profiles with respect to completion rates and times to degree. They did, however, share one major characteristic and that is a general reputation for turning out high quality PhDs.

An important goal of the designers of the GEI was to encourage departments to establish incentive structures that would encourage the timely progress of students through requirements they had to complete to earn the PhD, such as foreign language examinations, comprehensive examinations and dissertation proposal reviews. So, for

³ The recipients of the portable Mellon Fellowships were thought to be among the best and brightest new humanities PhD students and their decisions to enroll in these PhD programs provided a market test of their quality.

example, rather than guaranteeing incoming students multi-year financial aid packages, the intent was that financial aid would be made contingent on the timely completion of requirements. Funding for dissertation-year fellowships was encouraged, but only for students who had completed all of their other requirements prior to their 6th year of doctoral study and who were judged to be within one year of completing their dissertation writing.

The Foundation understood that it would take time for proposed changes in programs to be agreed on and implemented, that program changes would evolve over time, and that the changes that occurred would differ across the departments. As such, the GEI began with the expectation that the program would run for 10 years, but left open the possibility of providing support for only 5 years if the evidence indicated that little of value was resulting. The program did in fact run for 10 years, from 1991-92 to 2000-01.⁴ During that time period, approximately \$58 million was provided by the Foundation to the 10 universities and 54 participating departments and programs, an average in the range of \$113,000 per department per year.⁵ Moreover, to help the participating universities sustain the progress that had been made with the help of GEI grant funds, endowment grants were made to each participating university as the GEI ended and subsequently each university received an additional challenge grant; the Foundation spent a total of \$22.5 million on these two types of grants. The challenge grants were contingent on proposals submitted by the universities that indicated how they would use

⁴ At the end of 5 years, it became clear that two departments did not really care to participate in the GEI or were judged to have made too little progress to continue. These departments ceased participating in the GEI and were replaced by three new departments and programs in 1996 and 1997. One additional program was previously added in 1993.

⁵ Although there were 54 participating departments and programs, some departments ended their participation before 10 years and some started after 1991-92. In total there were 515 department/years of support provided under the GEI.

such funds to continue improving their PhD programs in the humanities throughout the university; there was no requirement that the funds be used in their participating GEI departments. In all, the Foundation devoted almost \$85 million dollars to support the GEI.⁶

Because the programmatic changes that the GEI induced would likely differ across departments and within a department over time, the framers of the GEI understood that it was important to learn not only whether on average the GEI led to an improvement in desired outcomes, but also to learn which programmatic changes were actually associated with the desired changes. Such information was thought to be essential if the successful innovations in the GEI were to be emulated by other departments. So along with collecting detailed data on student outcomes and the financial support they received, as part of the GEI information would also have to be collected on the characteristics of each department's doctoral program.

Initially, the impact of the GEI on attrition rates and times-to-degree was to be assessed by comparing outcomes for students who had enrolled in these departments' doctoral programs eight years prior to the GEI with outcomes for students who enrolled in these programs during the GEI⁷. However, after the GEI began, it was quickly realized that even highly satisfactory changes in, say, attrition rates could be caused by factors other than the GEI (such as changes in the labor market for humanities and related social science PhDs) that could not adequately be controlled for in such a simple "pre-post" design.

⁶ Included in this total were funds provided in the form of planning grants and funds expended for data collection and data management.

⁷ Seven institutions had data going back a full eleven years but those available for 1980-82 were not used in order to achieve comparability for all ten institutions.

As a result, the Foundation went back to the universities and asked them if they would provide similar data on student outcomes for their departments in the humanities that were not participating in the GEI. This would enable the Foundation to estimate the impact of the GEI by comparing changes in outcomes over time at the departments that received funding from the GEI to those departments that did not receive funding, while holding constant other variables that differed across departments and over time that might be expected to influence the outcomes.

Five of the universities had sufficiently detailed information on hand for other departments and they agreed to provide information, but others were unable to do so.⁸ So the Foundation turned to a set of other more or less similar universities and asked if they would provide data for a set of their departments in the humanities and related social sciences to help serve as a comparison group. The University of North Carolina, the University of California at Los Angeles and the University of California, San Diego all agreed to do so. No funding was provided to comparison departments at any of the 8 universities; they and their universities nonetheless provided data requested by the Foundation to aid in the Foundation's evaluation of the GEI.⁹

Table 1 lists the 54 treatment and the 47 comparison programs that participated in the GEI. Because the universities, rather than the Foundation, selected the treatment and comparison departments, the sample is not balanced across fields of study. For example, there are only three East Asian studies programs (only one of which is a treatment

⁸ That they were unable to do so says a great deal about the importance that the universities placed on collecting information about their PhD students' progress prior to the development of the GEI. One major benefit of the GEI is that departments now more regularly collect data to monitor the progress of their students.

⁹ Two of the universities at which there were no treatment departments (UNC and UCLA) received small amounts of funding to help them collect and provide student-level data to the Foundation.

department), one Ethics program and one Medieval studies program among the 101 treatment and comparison programs. These three fields had to be eliminated from the analyses that underlie many of the findings that we summarize below because there was no comparison department for the latter two fields and the sample size was too small to obtain meaningful results for East Asian Studies. So our empirical analyses made use of data from 51 treatment and 46 comparison departments.¹⁰

II. Data Collection: Departmental Student Level Data Bases

A major goal of the GEI was to use empirical data to inform decisions regarding departmental practices and to evaluate progress towards GEI goals. Thus participating universities were required to collect extensive data that would be submitted to the Foundation annually. Data collection requirements included reporting information about all entrants to the relevant PhD programs, their demographic characteristics at entry, and their progress through the program and the financial support that they received until completion or attrition occurred. This information was to be reported for entry cohorts at both treatment and comparison departments starting in 1980 (ten years prior to the start of the GEI) and continuing through 2006 (six years after the completion of the GEI).¹¹ Qualitative information was also collected from the treatment departments annually about the characteristics of their PhD programs and how the programs were evolving over time. This section describes the student-level data collection; we describe the departmental reports in the next section.

Great care was taken by the Foundation to establish a well-defined and standardized set of formats for data collection. It was feared that if each institution made

¹⁰ Some analyses use a smaller number of departments; for example the *Graduate Education Survey*, which we describe below, only collected information for students from 44 treatment departments.

¹¹ In some cases these data were available only starting in 1982.

its own statistical calculations of times-to-degree and completion rates that definitions might vary across departments and thus that comparisons across departments or institutions might be difficult. Thus, the Foundation asked that only raw data for each student and each department be reported, and the universities agreed to provide the data to the Foundation. The Foundation then accepted the responsibility to transform these data into consistently defined measures and to provide summary tables annually to the universities and departments for their own use.

To preserve confidentiality, the records were provided anonymously to the Foundation, with the universities assigning identification numbers to each student for the purpose of creating longitudinal records. Student names were not included on these records and the Foundation agreed that any publication of the data or analyses would be presented in a way that neither individual students nor individual departments could be identified.¹²

The first set of data requested by the Foundation provided information on student characteristics at the time of their entry to the PhD programs. These included demographic data, such as gender, citizenship, race and ethnicity; information on educational background including institution of undergraduate degree, undergraduate degree year and whether the student had a master's degree upon entry to the PhD program; and scores on the verbal and mathematical portions of the Graduate Record Examination (GRE) if these were available.

¹² Later, in 2002 when the *Graduate Education Survey* (GES) which is described below was undertaken, each university agreed to provide a file with student names and addresses matched to the identification numbers for the purpose of locating their students and former students so that they might participate in the GES. This file, by agreement with each institution's human subjects review board, was accessible only to one Mellon staff member, her assistant, and the staff at Mathematica Policy Research Institute (which undertook the survey), all of whom had completed human subjects protocol training. At no time were these names associated with the data that the institutions had provided to the Foundation.

The second set of data was reported annually on each student's progress through the program and the types of financial support that the student received that year. Each year institutions reported whether the student received the PhD, continued to study in the program, or terminated graduate study during the year. If a student advanced to candidacy or graduated during a year, the dates that those events occurred were reported. Some universities reported data for another milestone, when the student completed all the coursework – but unfortunately not all institutions collected and were able to report this information.

The financial support data that the departments were asked to report annually include the types of support allocated to the student during the academic year – fellowship, teaching assistant, research assistant, and tuition stipend – and summer support. Universities were asked to provide the dollar amount of the types of support received. However, these amounts were not reported for the early years in some departments.¹³ These dollar amounts were to include funds from all sources, internal and external, but in more than a few instances, information on external fellowships was incomplete. In addition, treatment departments were asked to report which students received academic-year or summer fellowships from the Foundation under the GEI program and the dollar amount of each of these awards.¹⁴

Initially, Sarah Turner, then on the Mellon staff, worked on designing and coordinating the data collection. Then, another Mellon staff person, Sharon Brucker, one

¹³ Treatment departments were more likely than control departments to provide dollar amounts.

¹⁴ Treatment departments were also asked to provide placement information for the students six months after they had received their PhDs. For the years prior to the start of the GEI, their responses were uneven; some institutions provided such information for over 90 percent of their graduates, while others slipped down to 65 to 70 percent in some years. Once the GEI was underway, response rates tended to be even lower and the comparison departments were never asked to provide this information. We have not used these data in the research that follow; focusing instead on the placement information that students reported when they participated in the *Graduate Education Survey*, which is described below.

of the coauthors), took over and has worked for more than a decade with data contact people at each university. It was she who made sure the data were submitted to the Foundation each year. In order to assure comparability of the data across institutions, a two-day conference with all of the institutional contact people and Mellon staff was held in 1993. The nature of each datum was discussed and if its definition differed across institutions, definitions were adjusted by some institutions to make the data as comparable across institutions as possible. Mellon staff then better understood the few cases where they needed to make modifications to data they received to make them comparable to those submitted by other institutions.

Each year as new data were uploaded into the database several consistency checks were done to assure that the data were consistent with information submitted in earlier years.¹⁵ This consistency checking process improved the accuracy of the database considerably. This checking required vigilance on the part of Mellon staff (and benefited greatly from the fact that the same staff member at Mellon was in charge of the database throughout the entire period), as well as cooperation and time from the data contact person at each institution. These individuals often spent long hours with staff at the individual departments to identify a student's status when the graduate school database did not capture the student's status accurately. The data contact peoples' efforts were essential to maintaining the quality of the database and, as a byproduct of their efforts

¹⁵ For example, if a student had been reported as graduating or leaving the program in year t-1, no new data should have been reported for the student in year t. If a new record were reported in year t, Mellon staff contacted the data person at the university and requested a check be made as to which year's data accurately reflected the facts. Often this resulted in a revision of a prior year's record; a student who had been reported as having left the program may actually have been "on leave" and was retrospectively "enrolled" in the program in the current year. Sometimes students who had been reported as continuing in the program in a prior year would retrospectively be reported as having left the program that year, based upon the report from the current year.

there have been many improvements in data collection and maintenance at participating graduate schools.

III. Data Collection – Departmental Program Data

One of the original goals of the GEI was to encourage departments to examine their programs and identify areas where change would improve both the efficiency of resource use as well as the effectiveness of the program. Once needs were identified, changes were to be designed and implemented. The request for departmental introspection and examination was intended to encourage departments to consider their degree programs as a whole (piecemeal reviews had been more common) and to give them incentives to make changes they deemed necessary. A commitment to tracking how these changes affected some key departmental outcomes was designed to engender accountability. Furthermore, by maintaining a record of the changes that were made and the impact of those changes on the outcomes, the Foundation hoped to identify those “innovations” which were the most efficacious in improving graduate education in the humanities. For all of these reasons, the Foundation required that treatment departments submit annual reports to it of how their programs were evolving.

The department reports were responses to questions that Mellon staff posed each year in an effort to learn more about what was going on, as well as what was going right.¹⁶ Furthermore, Mellon staff went to considerable effort to try to identify all

¹⁶ The departments also prepared financial reports each year that showed how they spent the money they received from the Foundation and all of the other funds that they spent supporting graduate education. These data were required because the Foundation did not want their dollars to “crowd out” funds that the departments would otherwise be spending on graduate education. So the Foundation used these reports to check that the departments and universities were maintaining their own efforts in support of graduate education in these fields.

innovations that were being made and then to try to summarize how they differed across departments.

Table 2 illustrates the types of innovations that were made in five exemplary departments that span English, History, and two other fields.¹⁷ In this table, a “y” indicates that the department adopted that innovation, an “n” indicates that it did not and a “y(plan)” indicates that the department indicated that it planned to adopt the innovation but that Foundation staff could find no mention in the department’s annual narrative reports that it actually had done so.

The first row asks whether the departments had clarified their expectations about students’ timetables for satisfactory progress towards the degree during their orientation of new students; four of the five departments reported that they had done so. The next two rows ask more specifically if the departments had clarified their time to degree expectations and their deadlines for advancement to candidacy and submission of a dissertation prospectus. Only one department did the first and four did the second.

The next two innovations related to advising. Four of the five departments required students and their advisors to schedule regular meetings and two of the five instituted formal group advising meetings. Another set of innovations related to faculty mentoring and reviews of students. Four of the five departments instituted earlier reviews of students’ progress and three of the five required faculty members to submit regular reports on students’ progress on their dissertations.

Given the often solitary nature of the research process in the humanities, another set of innovations related to whether the department encouraged students to work in

¹⁷ Per agreement with the participating departments and universities, the university names have been suppressed in table 2.

groups to reduce isolation by establishing group workshops or colloquia. For example, two departments established such workshops to aid students in preparing their dissertation prospectuses, while one established a workshop to aid students in getting their dissertations started.

A number of potential changes related to the nature of the curriculum. Three departments altered their coursework requirements as might have been expected given the complex requirements many departments had (English Department A actually increased the number of seminars it required students to take, changing a number of courses from lecture to seminar formats). A number changed the timing of advancement to candidacy, required a dissertation prospectus for advancement to candidacy, required students to identify the fields in which they would take their preliminary examinations earlier or established deadlines for submitting dissertation chapters. One changed course grading practices and four modified incomplete policies.¹⁸

Departments also reported how they planned to use the funding that they received from the GEI. Some reported that they planned to use some funds to support summer study for language examinations. Some planned to use funds to support summer travel and fieldwork. None of these five departments used funds for fellowships for students prior to their working on the dissertation; some used funds for fellowships at the start of the dissertation process, while all five of these departments used funds to help support dissertation writing. Indeed, four of the five departments used most of the funds that they received from the GEI for the latter purpose.

¹⁸ The number “97” in an entry means the department did not make the change until 1997; as noted above the departmental programs evolved over time.

One department, presumably in conjunction with a change in its university's policies, increased the tuition that students had to pay after their sixth year of enrollment in the program. Four of the five programs, presumably in response to competitive pressures, increased the number of years that they guaranteed students funding at the time of admission; as we noted earlier this was not a change intended by the architects of the GEI.

Department policies also included efforts to increase enforcement of rules to tie funding to satisfactory progress through the program. For example, two of the programs set a limit on the maximum years of funding that students could receive and two gave a final year's support only to students who had already satisfactorily completed two chapters of their dissertation.

In response to concerns that students were spending too many years as teaching assistants (TAs) and that work as a TA was slowing down their progress, three departments limited the amount of time students could spend working as TAs.¹⁹ One department tried to enhance the teaching assistant experience with more mentoring and a program that permitted students to design their own courses.

Finally, these departments instituted a number of other structural changes in their programs. Two more explicitly defined what they expected students to accomplish during summers, one added a placement advisor, two (as directed by their universities) said that they planned to improve their programs by reducing the size of their entering classes and one said that it would try harder to match students with mentors with similar interests.

¹⁹ Reducing the time students spent as TAs did not always reduce the amount of time they spent teaching because students at sample institutions in urban areas could often easily find part-time teaching opportunities at other academic institutions in the urban area.

Of course, what departments said that they planned to do, and what they actually did, often changed over time for at least four reasons.²⁰ First, some departments designed a long list of intended changes and discovered it was difficult to enforce them, while others made only a few changes but took great pride in enforcing them. So simply stating that an innovation was made doesn't mean that it persisted over time.

Second, many of the changes departments made “morphed” as they were implemented; sometimes we know about these changes, but sometimes we do not. This suggests that the nature of the PhD program facing PhD students at a given institution may change from cohort to cohort and the departmental reports do not provide enough detail to allow us to capture the details of all these changes.

Third, to be effective, many of the program changes require that students understand and respond to these changes. However, when a department says that it made a change, it does not follow that students understood the change and responded to it. Even if a change is made, it may take time for the change to be fully implemented and then to have an effect. Similarly, changes in the individual faculty member(s) administering a department's graduate program may influence the effectiveness of any program change.

Finally, two departments may have introduced the same innovations, but their implementation may have occurred in different graduate school environments. In one, the graduate school may have dictated the outlines of the program changes in a top-down fashion, while in the other the changes may have come directly from departmental faculty. Faculty members in the second department are much more likely to have “bought

²⁰ We present vignettes in *The Education of Scholars* gleaned from the same department reports, discussions with departmental representatives and the individual student records that describe why this occurred.

into” the changes and worked harder to implement them than faculty members in the first department did.

Foundation staff realized that to understand fully the changes in treatment departments, students would have to be surveyed about their experiences in their doctoral programs and their understanding of the nature of the program curriculum, expectations and culture. Furthermore, to obtain an understanding of which program changes actually had resulted from the GEI and which were simply responses to competitive pressures (see the discussion of five years of guaranteed financial support for students below), it would be necessary to obtain similar information for students who had studied in comparison departments.

In large part, the *Graduate Education Survey* (GES) grew out of this need. But Foundation staff also saw the survey as an opportunity to understand more about these humanities students themselves in addition to their experiences in their departments. They wanted to know about the students’ demographic profiles; the nature and extent of time commitments – both in their roles as teaching assistants and in employment outside their departments; the reasons for and the timing of their leaving their programs or graduating; their patterns of publishing; and their route to tenure if they chose that goal.

IV. *The Graduate Education Survey*²¹

The *Graduate Education Survey* (GES) queried all PhD students who had entered treatment and control departments’ PhD programs during the 1982-1996 period in order to obtain their views about the nature of their graduate programs and graduate departments, their experiences while in graduate school, their post-degree (or post-drop

²¹ A more complete description of the GES, including how it was administered and how Mathematica achieved such a high response rate, is found in Kalb and Dwoyer (2004).

out) labor market and educational experiences, and their publication experiences while in graduate school and after graduation. Many of the questions that related to the characteristics of their graduate experiences were designed by Foundation staff to reflect the characteristics of graduate programs that appeared in table 2.

The GES was designed by the Foundation and conducted by Mathematica Policy Research Institute. Between November 2002 and October 2003, the 18,320 who had matriculated at the treatment and control departments were surveyed and 13,552 responded, producing a response rate of 74%, which is remarkably high, particularly for a retrospective survey. As might be expected response rate was better for individuals who had completed their PhDs (81.3%) as compared to the rate for students still enrolled in their programs (75.8%), which in turn was higher than the rate for those who had left their programs (62.8%). The lower response rate for the last group was partially due to the fact that 20% of program leavers, many of whom had departed graduate study 15 to 20 years earlier, could not be located. For the same reason, response rates differed by entry cohort, with the response rates declining the further back in time the recipients had been graduate students. The 1991-1996, 1986-1990, and 1982-1985 response rates were 77%, 74% and 70%, respectively.

The first section of the questionnaire asked students about their entry to their graduate programs (including the type of financial aid they were offered), their department's academic expectations and requirements, and the means by which these expectations and requirements were conveyed to them. The second section asked questions about their interactions with their dissertation advisors and their departments, the overall learning environment in the department, the time it took them to complete

different phases of their programs, and their publications, if any, while in graduate school. The third section asked questions about their types of financial support they received while in graduate school and how they believed this support influenced their progress towards degree completion or leaving their programs. The fourth section solicited degree completion information (which we indicate below allowed us to check on accuracy of the degree completion data that the institution provided) and information on subsequent educational experiences of those who left doctoral programs. A fifth section solicited information on demographic information including the students' marital status and the number of children in their families during their graduate study years. Finally, the last section solicited information on the respondents' employment status six months after degree completion or departure from their program, three years after that event, and as of the survey date. Information on early career publications was also requested.

The same careful consideration to the data's accuracy that was given to the institutional student data base was also given to the GES by Foundation staff. In some situations, respondents' replies about their enrollment status (whether they had received the PhD, were still enrolled, or had dropped out of the program) were different in the GES than they were in the last institutional report about the respondents' enrollment status. Mellon staff and the institutional representatives worked to resolve these discrepancies, most of the time satisfactorily so.

Concern about the accuracy of the self-reported data on publications led the sole Foundation staff members who had access to the names of survey respondents to check these self-reported publications data for a sample of respondents against publications information obtained from Web pages and bibliographical indices. In almost all cases the

self-reported publications data were sufficiently close to the objective measures that we felt confident in our ability to use the self-reported data for the entire sample.

V. What We Learned from the GEI

Taken together, the institutional data bases and GES provide a literal treasure trove of information. These data bases have been analyzed by researchers at the Cornell Higher Education Research Institute (CHERI) and Foundation staff. Details of our technical analyses and findings are reported in a number of working papers and our forthcoming book.²² Here we summarize briefly some of our major findings concerning the impact of the GEI on attrition rates, completion rates and times to degree; on what characteristics of PhD programs in the humanities and related social sciences influence these outcomes and how the GEI influenced these characteristics; on what happened to students who left their PhD programs prior to receiving their degrees; on the early career job market outcomes of the new PhDs and on the graduate school publications and early career outcomes of these new PhDs.²³

A. Graduate Student Outcomes, Graduate Program Characteristics, and How Graduate Program Characteristics Influence Graduate Student Outcomes

Our analyses suggest that the GEI had modest effects on what students did in the expected directions: attrition rates and times-to-degree were reduced and completion rates were increased. These effects, we find, were driven in part by intentional reductions in the sizes of entering cohort, which in turn produced improvements in student quality, as

²² The working papers include Ehrenberg, Jakubson, Groen, So and Price (2006), Groen, Jakubson, Ehrenberg, Condie, and Liu (2006) and Price (2005).

²³ Our survey of findings here is necessarily selective and in our forthcoming book we discuss a much wider range of findings including the role of gender, marital status, family status, race/ethnicity and citizenship status on graduate school and early career job market and publications outcomes.

gauged by GRE scores and improvements in financial support, in addition to increases in financial support that were attributable to infusions of Mellon funds.

Some of the improvements in financial support came in the form of increases in the probability of guaranteed packages of multi-year of support being awarded at the time students were admitted so as to enhance the their inclinations to enroll. Departments undoubtedly moved in this direction in response to increased competition in the market for new PhD students. However, the framers of the GEI did not anticipate such actions; their hope had been to make financial aid more competitive and based on satisfactory progress through the program. While these multiyear packages reduce the probability of students dropping out after their first year, they appear to have been associated with some increase in the probability of dropping out later in students' graduate career, thus leading to an unintended substitution of later drop outs for early drop outs. This is a finding bearing further exploration and discussion.

Analysis of the data collected by the GES provided us with an understanding of the different routes via which characteristics of graduate program influence student outcomes. We find that improving advising and the clarity of program requirements are associated with reduced attrition. It is also the case that departmental expectations about the nature of dissertations have strong effects on attrition, even in the early years of the program. In particular, departments that encourage students to finish their dissertations as quickly as possible have lower rates of attrition, while departments that emphasize the importance of students polishing their dissertations and publishing things prior to graduation have higher rates of attrition. Similarly, graduation probabilities are higher

when advising is improved and when departmental expectations about the dissertation are of the first type rather than the second.

The data also reveal that there is a trade off here; the GES indicates that students who publish while in graduate school have higher probabilities of obtaining tenure track appointments at four-year institutions upon graduation. Students who publish while in graduate school also are more likely to publish early in their careers. To the extent that faculty members are concerned with the career success and publications of their students, advising students to try to publish while in graduate school, even if this increases the likelihood that some students will drop out of their programs and increases the time to degree of others, may be good advice (but see below). Put simply, although the explicit goals of the GEI designers were to reduce times-to-degree and attrition rates, there is evidence that some faculty members did not subscribe to these goals and did not act in accord with them. The inclination for faculty members to do what they think best for the intellectual development of their students and their fields should not be underestimated when efforts in the future are made to change doctoral programs.

Our analyses also suggested which characteristics of doctoral programs were influenced by the GEI. On average, the GEI was associated with increased seminar requirements, increased expectations for summer work related to making progress in graduate programs, increased clarity of program expectations, an increased likelihood in the smaller departments that students were encouraged to finish their dissertations promptly, and improvements in financial support in larger departments (where financial support was tightest prior to the GEI). The effect of the GEI on student outcomes that were produced by these program characteristics was modest; there remains considerable

variation among departments in these characteristics even now. Hence it is still possible that there is still room for changes to affect student outcomes in the future more than the current data indicate they have.

Finally, it is worth emphasizing that financial factors are not the primary reason that students drop out of PhD programs, as many suppose. To be sure, financial support is necessary for graduate students but it is not a sufficient guarantee of degree completion. Even the most generous financial aid packages – for instance, ones that include fellowships for each of the first six years that students are enrolled in their PhD programs – are associated with substantial drop out rates. Amply supporting graduate students but doing nothing else will not solve the problem of attrition.

B. People Who Drop Out of PhD Programs

Stated simply, dropping out of one of the PhD programs in our sample is not the end of a student's life. The unique nature of the GES allowed us to ascertain what actually happened to students who left PhD programs without degrees. Indeed, over 10 percent of the “drop outs” in the GES sample ultimately received PhDs from different departments, with many of these people receiving their PhDs in fields other than the one in which they were initially enrolled.²⁴ Early leavers are much more likely than those who left later on to receive PhDs elsewhere. And almost another 20 percent of the leavers went on to receive professional degrees, including law and MBA degrees. We now know that rates of attrition from individual universities considerably overestimate what appear to be high rates of failure of graduate education.

²⁴ It is worth emphasizing the point that departmental or university level studies of completion rates of PhD programs will understate system-wide completion rates, just as individual institutional level studies of undergraduate completion rates understate the number of first-time freshmen who ultimately receive undergraduate degrees from some institution.

We also find that what might be seen as downward occupational mobility of those leaving their PhD programs is large but only temporarily so. While 10% of dropouts were employed in clerical and administrative positions six months after departure from graduate school, by the time three years had elapsed after their departure, this percentage had been reduced even further and the majority was employed in professional occupations.

C. Job Outcomes After the PhD

The percentage of new PhDs in our sample employed in tenure track positions at 4-year institutions six months after receipt of their PhDs declined slightly during the decade of the 1990s and was about 30% for the cohort that received degrees during the 1998-2000 interval. While a similar decline was observed for new PhDs three years after receipt of the degree, the percentage of them in tenure track positions at 4-year institutions was much higher, 52% for the 1998-2000 cohort. These data suggest that there is considerable early career mobility for these new PhDs; indeed about 50% of the new PhDs who had full-time non-tenure track positions six months after receipt of the PhD had moved to full-time tenure track positions three years later..

As time to degree increases, the probability of obtaining a tenure track position within three years after receipt of the PhD monotonically declines but only for those who took eight years or more to complete. This is an important finding in that it demonstrates that time-to-degree matters in getting much sought after tenure track posts but only if time-to-degree exceeds eight years. As we already noted, publishing while in graduate school enhances the probability of obtaining a tenure track position and it also enhances the probability of attaining tenure within 15 years after entry to graduate school.

D. Graduate School and Early Career Publications

About 40% of our sample published, or had accepted for publication, at least one book or refereed article while in graduate school and about 67% did so during their first three years after receipt of the PhD. Publications while in graduate school are an important predictor of publishing soon after earning the degree and PhDs who reported that their departments expected them to publish while in graduate school have published more early in their careers than did other PhDs in our sample.

As noted above, this may provide one explanation for why our estimates of the effects of the GEI on time-to-degree are so modest. Faculty members at these top programs appear to be less concerned with their students' times-to-degree than they are with preparing the next generation of scholars. As we noted, it is only among those who took more than seven years to complete their degrees that long times-to-degree are inversely correlated with tenure track job probabilities. However, it should also be noted that taking more than seven years is far from unusual; in fact, over 50% of the completers in our sample took that long. Hence advisors appear to have little incentive to press students who will complete in seven years or less to hurry to completion. For students who otherwise would complete in eight years or more, the positive effect of publishing on the quality of jobs degree recipients get must be weighed against the adverse effect of longer times-to-degree on job outcomes.

Having said this, we find that in the GES sample, the probabilities that students will publish while in graduate school are largest for those who complete their degrees within five years. As times-to-degree increase, publication probabilities while in graduate school decline. This is likely to be the outcome of a selection effect; other factors held

constant, the more talented and motivated students are, the higher the likelihood of their publishing and the shorter the time it takes them to finish their degrees. It follows that the students with the shortest times-to-degree are those who have published most often while in graduate school are the ones who are also most apt to be employed in tenure track positions after their degrees.

While the explicit goal of the GEI was to enhance graduate program effectiveness, its framers did not give priority to one important indicator of program effectiveness – that is, encouraging students to publish and thus transforming them into contributors to knowledge. Nonetheless, it is reassuring as we noted that as many as 40 percent of the students in the GEI published while still in graduate school and further, that we estimate that the GEI increased the probability that students would publish while in graduate school by roughly 20% to 25%. This is no small accomplishment and one not typically taken into account. For reasons we do not yet understand, the GEI had a somewhat smaller impact on the propensity of degree recipients to publish early in their careers.

VI. More General Lessons

The GEI confirmed that what goes on at the department level does matter. Many of the characteristics that influence graduate students' progress through their programs are controlled by departments, not the graduate dean or other central administrators. While there is strong role for graduate deans to play in improving graduate education, and the paper by Deborah Stewart and Daniel Denecke in this volume discusses this role, future efforts to improve graduate education surely should focus on departments. And one should not underestimate the difficulty of changing departmental culture and getting

faculty to “buy into” program changes. Our analyses suggest that innovations that are initiated at the departmental level are much more likely to be supported than those that are initiated “top down.”

Furthermore, even after innovations take place, departmental programs evolve over time. Sometimes this is due to the departure of a key concerned faculty member or to a key concerned faculty member assuming responsibility for administering a department’s program. Sometimes, as the discussion above on multiyear guaranteed financial aid package policy makes clear, this is due to universities and departments responding to competitive pressures. Because of this it is important for departments and graduate deans to constantly examine their graduate programs to make sure that progress towards desired outcomes does not erode.

To do so requires departments to collect data regularly on their students’ progress through their PhD program. One important impact of the GEI was to encourage departments to collect such data and they now more regularly do so. An important role of graduate deans is to monitor and standardize the collection of such data and where appropriate to respond to the messages it carries. The forthcoming National Research Council evaluation of doctoral programs will similarly press departments to collect such information and we view this effort as very important.

The GES has shown that retrospective questionnaires of current and former students can provide detailed information on multiple characteristics of graduate programs. Our analyses have shown how these characteristics can be aggregated using factor analysis methods into a smaller number of underlying factors and, with data from

multiple departments and multiple entering cohorts, an analysis of which factors influence graduation and attrition probabilities can be undertaken.

We strongly believe that similar analyses may be profitably undertaken for other graduate fields of study (such as science and engineering) where different characteristics of graduate programs may prove important.²⁵ Such analyses require that departments have collected data on student characteristics, their progress through their programs and the types of financial support that they receive each year, as well as a GES-type retrospective study. However, they do not require that a major intervention, such as the GEI, has already taken place.

The GEI did not collect several types of data that, in retrospect, appear useful for evaluations of graduate education. The GEI had no information on the numbers of program faculty in each department and their stability over time. It had no information on whether incentives were present for faculty to mentor doctoral students (for example, workload credit for supervising dissertations). It collected no information on the match of students and faculty by research interests and/or gender and ethnicity, no information on the number of doctoral students that a student's advisor was simultaneously supervising and had supervised in the past, and no information on advisors' past success in placing their students. While collecting such information may prove to be a formidable task, it clearly would be useful.

Finally, we believe it is important to be modest about the likely impacts of any foundation-related efforts to improve doctoral education, because there are a host of

²⁵ We have already suggested that this be done at a National Research Council workshop conducted in June 2005.

factors that may coalesce to make it difficult to achieve a foundation's objectives.²⁶ As we have noted, the objectives of individual faculty members may differ from those of foundations or graduate deans and competitive pressures in the market for doctoral students may push departments to behave in a manner different from what foundations intend. An excellent example of such unanticipated consequences coming into play can be found in the case of the GEI, in which multiyear guaranteed fellowships rather than incentive-based financial aid became the norm.

In addition, changing labor market conditions may cause doctoral students to behave differently than a foundation expected them to behave. For example, the GEI was developed at least partially in response to a projection of forthcoming shortages of faculty in the arts and sciences; the hope of the framers of the GEI was that shorter times to degree and higher completion rates would increase the potential flow of new doctorates into faculty positions. However, these shortages never materialized and increasingly institutions instead turned to the use of part-time and full-time non-tenure track faculty members to hold down costs. With an increased difficulty of finding tenure track employment, which led to concerns about health insurance, doctoral students probably correctly perceived that rushing to finish their degrees quickly might not be in their best interest. Put simply, it is difficult to predict all of the consequences of programs and how these will interact with the changing world.

²⁶ Clotfelter (2005) discusses this issue in much more detail.

References

William G. Bowen and Julie Ann Sosa, *Prospects for Faculty in the Arts and Sciences* (Princeton NJ: Princeton University Press, 1989)

William G. Bowen and Neil L. Rudenstien, *In Pursuit of the PhD* (Princeton, NJ: Princeton University Press, 1992)

Charles Clotfelter, "Patron or Bully? The Role of Foundations in Higher Education," *Terry Sanford Institute of Public Policy Working Paper 05-09* (Durham, NC: Duke University, November 2005)

Ronald G. Ehrenberg, George Jakubson, Jeffrey Groen, Eric So, and Joseph Price, "Inside the Black Box of Doctoral Education: What Program Characteristics Influence Doctoral Students' Attrition and Graduation Probabilities?" *National Bureau of Economic Research Working Paper W12065* (Cambridge, MA: National Bureau of Economic Research, March 2006)

Jeffrey Groen, George Jakubson, Ronald G. Ehrenberg, Scott Condie and Albert Yung-Hsu Liu, "Program Design and Student Outcomes in Graduate Education," *National Bureau of Economic Research Working Paper W12064* (Cambridge MA: National Bureau of Economic Research, March 2006)

Laura Kalb and Emily Dwoyer, *Evaluation of the Graduate Education Initiative: Final Report* (Princeton, NJ: Mathematica Policy Research Inc., February 2004)

Joseph Price, "Marriage and Graduate Student Outcomes," *Cornell Higher Education Research Institute Working Paper WP75* (Ithaca, NY: Cornell Higher Education Research Institute, July 2005) (available at www.ilr.cornell.edu/cheri)

Table 1

Treatment (T) and Comparison (C) Programs Participating in the Graduate Education Initiative

Field (number of treatment, comparison programs)	Berkeley	Chicago	Columbia	Cornell	Harvard	Michigan ^f	Penn	Princeton	Stanford ^g	Yale	UCLA	UCSD	UNC
Anthropology (6,4)			T	C	T	T^e		T	T	T	C	C	C
Art History (6,3)	T		T^a	C	T	T^{ab}		T	C	T	C		
Classics (3,5)	T			C		T	T	C		C	C		C
Comp. Lit. (2,4)	T			T		C		C		C	C		
E. Asian Studies (1,2)				C				T^e			C		
English (9,3)	T	T	T	T	T	T	T	C	T	T	C	C	
Ethics (1,0)								T					
History (8,3)	T	T		T		T	T	T^c	T	T	C	C	C
Medieval Std. (1,0)				T									
Music (3,6)			T	C		C	T	C	C	T	C	C	
Philosophy (4,5)		T	T	C		T^a			T	C	C	C	C
Politics/Govt. (4,5)		T		T	T			T	C	C	C	C	C
Religion (2,3)			T		T			C	C	C			
Romance Lang. (2,4)				C		C	T	T^{ad}	C		C		
Total (54,47) ^f	5,0	4,0	6,0	5,7	5,0	8,3	5,0	7,5	4,5	5,5	0,11	0,6	0,5

Where:

^a Added as a treatment program in 1996

^b Includes classical art and archeology

^c Includes history of science

^d Includes German and Slavic

^e Ended treatment department status in 1995-96

^f Two interdisciplinary Michigan programs, Anthropology and History, and American Culture, were also treatment programs starting in 1997-98. They, along with Cornell's Medieval Studies (which began as a treatment department in 1993) and Princeton's Ethics programs, have been excluded from the evaluation of the GEI because of a lack of any control programs in these fields

^g Stanford departments started treatment status one year later

Table 2

Innovations Implemented as Part of GEI: Examples from 5 Sample Departments

Innovation	Description/Purpose	English A	History B	History C	Other D	Other E
Expectations	Clarify/timetable/orientation	y	y	y	n	y
	Clarify TTD expectations	n	n	y	n	n
	Clarify ATC and prospectus deadlines	y	n	y	y	y
Advising	Required schedule	y	y	y	n	y
	Formal group advising	y	n	n	n	y
Monitoring	Earlier review	y	n	y	y	y
	Faculty submit progress reports on dissertations	y	n	y	y	n
Group	To reduce isolation	y	y(plan)	n	y	n
Workshops or Colloquia	To prepare prospectus	y	n	y	y	n
	To get dissertation started	n	n	n	y	n
	Dissertation writing/feedback	y	y(plan)	y	y	y
	Early seminar/fieldwork	n	n	n	n	n
	Collaborative work	n	y	n	n	n
	Job/profession preparation	n	n	y	n	n
Curricular Changes	Coursework requirements	y	n	y	n	y
	Writing requirements	n	n	n	y	n
	Timing for advance to candidacy	n	n	y	y	y
	Nature of advance to candidacy	y	n	n	y	y
	Prospectus required to advance	n	n	y	y	n
	Reduce language requirements	n	n	n	n	n
	Modify incomplete policy	y	y	y	y (97)	n
	Limit length of prospectus	n	n	y (97)	y	y
	Courses graded	y	n	n	n	n
	Prelim fields identified earlier	n	y	y	n	y
Deadlines for submitting dissertation chapters	y	n	y	y	y	
Use of Mellon Funding	Summer language	n	n	8%	y	n
	Summer travel/fieldwork	9%	22%	25%	55%	0%
	Pre advancement to candidacy	n	n	n	n	n
	Post advancement/start-up	13%	26%	0%	20%	2%
	Finishing dissertation	78%	52%	66%	25%	98%

Table 2 (continued)

Innovation	Description/Purpose	English A	History B	History C	Other D	Other E
Tuition Policy	Tuition increases after 6 years	n	n	n	n	y
	Guaranteed multiyear packages	y	y	y	y	n
Enforcement of Rules	Funding conditional on timing	y	y	n	y	y
	Limit years funded	y	n	y	n	n
	Final-year support conditional on completing specific chapters	y	n	n	n	y
	Can't register if miss deadlines	n	n	n	n	n
	No further funding (including TA) if dissertation defense not scheduled by end of dissertation write-up year	n	n	y	n	n
	No TA if prospectus not on time	n	n	y	n	n
	Post-doc available if schedule defense by end of write-up fellowship year	n	y	y	n	n
TA Changes	Reduce time as a TA	y	y	y	n	n
	Enhance TA experience/design own course	n	y	n	n	n
	Improve training for teaching	n	n	n	n	n
Structural Changes	Define summer tasks	n	n	y	y	n
	Added placement advisor	y	n	n	n	n
	Reduce size of entering cohort	y	n	n	n	y
	Match better with available mentors	n	n	n	y	n

where

y (plan) - planned to implement but no report that the department did so in annual reports

y (97) – implemented in 1997

TTD – time to degree

ATC – advancement to candidacy

TA – teaching assistantship