

# CHANGING THE EDUCATION OF SCHOLARS

## An Introduction to the Andrew W. Mellon Foundation's Graduate Education Initiative

RONALD G. EHRENBERG, HARRIET ZUCKERMAN,  
JEFFREY A. GROEN, AND SHARON M. BRUCKER

In 1991 the Andrew W. Mellon Foundation launched the Graduate Education Initiative (GEI) to improve the structure and organization of PhD programs in the humanities and social sciences and to combat the high rates of student attrition and long time to degree completion prevailing in these fields. While attrition and time to completion 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. An array of characteristics of doctoral programs was earmarked as likely contributors to high attrition and long degree-completion time. These included unclear or conflicting expectations of the academic performance of students, a proliferation of specialized courses, elaborate and sometimes conflicting requirements, intermittent supervision, epistemological disagreements on fundamentals, and—not least—inadequate funding.<sup>1</sup> In short, the intention was to improve doctoral education and make it more efficient.

This was far from the first such effort to reduce times to degree completion and rates of attrition. Earlier interventions, which provided grants in aid to individual students or to graduate schools to distribute as they saw fit, had conspicuously failed.<sup>2</sup> Based on data that showed marked differences among graduate departments in the time it took to earn degrees; data about attrition rates among and within the sciences,

social sciences, and humanities; and a great deal of experience on the ground, the architects of the GEI concluded that graduate education could be improved only if departments would change their PhD programs. The Mellon Foundation then shifted much of the support it provided for doctoral education away from fellowships for individual students and moved to block grants that would be awarded to major universities and the departments they selected.

Ten institutions—the University of California, Berkeley, the University of Chicago, Columbia University, Cornell University, Harvard University, the University of Michigan, the University of Pennsylvania, Princeton University, Stanford University, and Yale University—were each invited to nominate four to six departments to participate in the GEI. These universities were chosen because as a group they had attracted the largest number of fellowship winners of the Mellon Foundation's portable doctoral dissertation awards.<sup>3</sup> 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 carefully review their curricula, examinations, advising, and official timetables with an eye toward facilitating timely degree completion and reducing attrition (especially late attrition), while maintaining or increasing the quality of doctoral training they provided.<sup>4</sup> There was no requirement that the departments named by the universities be in need of particular help—that is, the departments did not need to have low completion rates and long times to degree completion, nor were they necessarily well-organized, thus meriting additional support. 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 completion. They did, however, share one major characteristic: a general reputation for turning out high-quality PhD holders.

The designers of the GEI encouraged departments to establish incentive structures that would promote students' timely progress through requirements they had to complete to earn the PhD, such as meeting foreign language requirements, passing comprehensive examinations, and completing dissertation proposals. For example, rather than guaranteeing incoming students that they would receive multiyear financial aid if they met departmental standards, the GEI sought to make annual financial aid contingent on the timely completion of a series of requirements. Funding for dissertation-year fellowships was encouraged, but only for students who had completed all other requirements before their sixth year of doctoral study and who were judged to be within one year of completing their dissertations.

The Mellon Foundation understood at the outset 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 ten years, but left open the possibility of providing support for only five years if the evidence indicated that the effort had been ineffective. The program did in fact run for ten years, from 1991–92 to 2000–2001.<sup>5</sup> Approximately \$58 million was provided by the Foundation to the ten universities and fifty-four participating departments and programs, an average in the range of \$113,000 per department per year.<sup>6</sup> 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 \$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 such funds to continue improving their PhD programs in the humanities throughout the university; there was no requirement that the funds be used in the participating GEI departments. In all, the Foundation devoted almost \$85 million in support of the GEI.<sup>7</sup>

Because the programmatic changes that the GEI induced would likely differ across departments and within each department over time, the framers of the GEI understood that it was important to learn not only whether on average the GEI led to improvements but also to identify programmatic changes associated with general changes that occurred. Understanding the mechanisms of change was essential if the successful innovations the GEI introduced were to be emulated by other departments. This led to the decision to collect evidence on characteristics of each department's doctoral program along with detailed data on student outcomes and the financial support students received.

Initially, the impact of the GEI on attrition rates and times to degree completion was to be assessed by comparing outcomes for students who had enrolled in these departments eight years before the instigation of the GEI with outcomes for students who enrolled in these same departments during the time the GEI was in place.<sup>8</sup> However, after the program began, Mellon Foundation staff quickly realized that even highly satisfactory changes in, say, attrition rates or times to degree completion could be caused by factors other than the GEI (for example, changes in the labor market for humanities and related social-science PhDs) that could not adequately be gauged using the original “pre/post” design.

As a result, the Foundation decided that comparative data would be needed on departments that were in the same disciplines as those in the GEI but which had not benefited from it. These departments would serve as “controls” while “treatment” departments would be those participating in the GEI and receiving financial aid—following the conventional terms used in evaluation studies of all kinds. A first step was to ask the universities participating in the GEI to provide similar data on student outcomes for departments that were not receiving GEI funding. This would make it possible to estimate the impact of the GEI over time and to hold constant other variables that might be expected to influence the outcomes.

Five of the universities had on hand sufficiently detailed information about other departments, and they agreed to provide it; the other five were unable to do so.<sup>9</sup> To increase the number of departments in the comparison group, the Foundation turned to a set of other universities with highly rated graduate programs for data on their departments in the humanities and related social sciences. The University of California, Los Angeles, the University of California, San Diego, and the University of North Carolina at Chapel Hill generously agreed to do so.<sup>10</sup> None of the departments designated as controls received funding, a condition necessary to make the planned comparisons as valid as possible.<sup>11</sup>

The 54 treatment and 47 comparison programs that participated in the GEI are listed in table 1.1. Because the universities, rather than the Foundation, selected the treatment departments and the comparison departments were selected if data were available, the two groups do not contain the same number of departments in each discipline. For example, there are only three East Asian studies programs, one ethics program, and one medieval studies program among the 101 treatment and comparison programs. To improve the comparability of the treatment and comparison groups, these three fields were eliminated from the analyses that underlie many of the findings we summarize below, in part because the sample sizes for East Asian studies were too small to obtain meaningful results and in part because no data were available for comparison departments for the other two fields. Ultimately, our empirical analyses drew on data from 51 treatment and 46 comparison departments.<sup>12</sup>

It is essential to underscore the fact that the treatment and comparison departments were not randomly assigned and thus there are differences between them on a number of dimensions, including program size, selectivity (as measured by entering student test scores), and doctoral program rankings. These differences are controlled for to the extent possible in our empirical analyses.

TABLE 1.1

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

Field (number of treatment, comparison programs)	Berkeley	Chicago	Columbia	Cornell	Harvard	Michigan <sup>f</sup>	Penn	Princeton	Stanford <sup>g</sup>	Yale	UCLA	UCSD	UNC
Anthropology (6, 4)			T	C	T	T <sup>e</sup>		T	T	T	C	C	C
Art History (6, 3)	T		T <sup>a</sup>	C	T	T <sup>ab</sup>		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		
East Asian Studies (1, 2)				C				T <sup>e</sup>			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 <sup>c</sup>	T	T	C	C	C
Medieval Studies (1, 0)				T									
Music (3, 6)			T	C		C	T	C	C	T	C	C	
Philosophy (4, 5)		T	T	C		T <sup>a</sup>			T	C	C	C	C
Politics/Government (4, 5)		T		T	T			T	C	C	C	C	C
Religion (2, 3)			T		T			C	C	C			
Romance Languages (2, 4)				C		C	T	T <sup>ad</sup>	C		C		
Total (54, 47) <sup>f</sup>	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

<sup>a</sup>Added as a treatment program in 1996.<sup>b</sup>Includes classical art and archaeology.<sup>c</sup>Includes history of science.<sup>d</sup>Includes German and Slavic languages.<sup>e</sup>Ended treatment department status in 1995–96.<sup>f</sup>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 program (which began as a treatment department in 1993) and Princeton's ethics program, have been excluded from the evaluation of the GEI because of a lack of any control programs in these fields.<sup>g</sup>Stanford departments started treatment status one year later.

**DATA COLLECTION: DEPARTMENTAL DATABASES—EVIDENCE ON STUDENTS**

One condition universities were required to meet in order to receive GEI support was the provision of both quantitative data on students and their progress and qualitative data on departmental educational practices. Such data were needed if lessons were to be drawn from the GEI. Thus, participating universities were required to collect extensive data that would be submitted to the Foundation annually. Data were collected about all entrants to the relevant PhD programs, their demographic characteristics at entry, their progress through the program, and the financial support they received until completion or attrition occurred. This information was to be reported for entry cohorts in both treatment and comparison departments, starting with entrants from 1980 onward (ten years prior to the start of the GEI) and continuing through 2006 (six years after the completion of the GEI).<sup>13</sup> This design allowed for comparisons of treatment departments before the GEI was instituted and during its tenure and the comparison of GEI treatment with non-GEI control departments in the same time periods. Qualitative information was also collected annually from the treatment departments about the characteristics of their PhD programs and how those programs were evolving over time. This section describes the student-level data collection; we will describe the departmental reports in the next section.

The Foundation established standardized formats for data collection to assure comparability among institutions. In order to avoid variations in reporting due to the use of different measurement procedures, institutions were not asked to calculate their own times to degree completion, attrition rates, and completion rates. The Foundation asked that only raw data for each student and each department be reported, and it transformed these data into consistently defined measures.<sup>14</sup> To preserve the confidentiality of the data, the universities assigned identification numbers to each student for the purpose of creating longitudinal records, but the records they provided to the Foundation were anonymous. Further, the Foundation promised that when analyses of the data were published, neither individual student records nor individual departments would be identified.<sup>15</sup>

Two classes of data were routinely collected on students. The first consisted of students' demographic and educational characteristics at the time they entered PhD programs, including gender, citizenship, race, and ethnicity; their educational backgrounds (where and when they received their undergraduate degrees and whether each student had a master's degree upon entry); and scores on the verbal and mathematical

portions of the Graduate Record Examination (GRE), if these were available.

The second class of data was reported annually by the institutions and provided information on each student's progress through the doctoral program<sup>16</sup> and the types of financial support that each received that year—whether each held a fellowship, teaching assistantship, or research assistantship; received a tuition stipend; and/or received summer support—as well as the dollar amount of each allocation.<sup>17</sup> 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 indicate which students received academic-year or summer fellowships from the Foundation under the GEI program and the dollar amount of each of these awards.<sup>18</sup>

Initially, Sarah Turner, then on the Mellon Foundation staff and now on the faculty at the University of Virginia, designed and coordinated the data collection. In 1991 Sharon Brucker (one of the coauthors of this chapter) took over these responsibilities and has worked as data manager and analyst since then. During this time she has been in continuous touch with data representatives at each university to make sure the data were submitted annually and in proper form.<sup>19</sup> Each year, as new data were uploaded into the database, checks were made to ensure that the new data squared with information submitted in earlier years.<sup>20</sup> Consistency checking vastly improved the accuracy of the database. This required continuing vigilance on the part of Mellon Foundation staff, and the fact that the same staff member was in charge of the database throughout the entire period was of great assistance in this vigilance. The exceptional cooperation of the institutional representatives also helped to improve accuracy.<sup>21</sup>

#### DATA COLLECTION: DEPARTMENTAL DATABASES—EVIDENCE ON PROGRAMS

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 quality of the education students received and the effectiveness of the programs. 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 periodic reviews, for example, of language requirements or qualifying exams were common) and to give them incentives to make changes they deemed necessary.

Tracking how these changes affected some key departmental outcomes was intended to promote accountability, while maintaining a record of the changes that were made and their subsequent effects would serve—the Foundation hoped—as a means of identifying those innovations that proved useful. For all of these reasons, the Foundation required treatment departments to submit annual reports on how their programs were evolving.

The reports were not free-form narratives but instead were responses to questions Mellon Foundation staff posed each year in an effort to learn more about what was going on in the departments, both right and wrong.<sup>22</sup> In the meantime, Foundation staff went to considerable effort to try to identify each innovation that was being tried and then to summarize how such innovations were distributed among departments.

Table 1.2 identifies ten classes of innovations or changes in five sample departments,<sup>23</sup> spanning English, history, and two other fields, and details specific innovations within each class.<sup>24</sup> These innovations include:

- clarification of deadlines to be met and expectations for time to degree completion
- improvements in advising, such as required schedules, matching the substantive interests of students and their advisers, and formal group advising
- increased monitoring, such as early review of students' performance, requirements of faculty to submit reports on students' dissertations
- introduction of workshops and colloquia on dissertation prospectuses and writing and on job seeking and placement; increased collective activity among graduate students
- curricular changes, such as changes in coursework requirements, advancement to doctoral candidacy, examination formats, incomplete policies
- focused use of Mellon Foundation funding for summer study, predissertation research, and field trips, and particularly for dissertation-year fellowships
- changes in financial aid policies and tuition charges, introduction of guaranteed multiyear financial aid; increases in tuition past the sixth year of enrollment
- enforcement of rules already in place, including limits on the number of years of funding; prohibition on registration if deadlines are missed; and limits on funding if doctoral defense not scheduled
- changes in the timing of teaching assistantships and training for them
- structural changes, including reductions in the size of entering cohorts; and establishing department placement officers.

Of course, what departments said that they planned to do did not square perfectly with what was actually done and what survived over time.<sup>25</sup> Some departments designed long lists of intended changes and discovered it was difficult to make them, whereas others made only a few changes but took great pride in maintaining them. Stating that an innovation was put in place does not necessarily mean that it remained there.

(Table 1.2—cont.)

Innovation	English A	History B	History C	Other D	Other E
Final-year support conditional on completing specific chapters	y	n	n	n	y
Can't register if deadlines are missed	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
Postdoctorate position available if defense is scheduled by end of write-up fellowship year	n	y	y	n	n
TA Changes					
Reduced time as a teaching assistant	y	y	y	n	n
Enhance TA experience/design personal 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
Add placement advisor	y	n	n	n	n
Reduce size of entering cohort	y	n	n	n	y
Better match with available mentors	n	n	n	y	n

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.

Of the changes departments made, many evolved as they were implemented; sometimes we have data on what occurred, but surely we do not have complete data. This suggests that the PhD program that students encounter at a given institution may differ from cohort to cohort, and that departmental reports do not provide enough detail to allow us to capture all changes that occurred.

To be effective, many program changes required that students understand and respond to them. However, when departments introduced an innovation, students did not necessarily understand it, much less respond appropriately. It thus might take time for innovations to be fully implemented, and subsequently to have an effect. Finally, departments may have introduced the same innovations, but their implementation in different graduate school environments can readily make them incomparable.

Foundation staff realized that to understand the changes treatment departments made and their effects, it would be necessary to ask students directly about their doctoral programs, the curriculum, the expectations

their professors had, and about the prevailing departmental culture. Collecting such information from students, in both the treatment and comparison departments, would be more useful for sorting out the distinctive effects of the GEI than would analyzing the more general changes occurring in many departments. In large part, the Graduate Education Survey (GES) grew out of a need to get students' perspectives on their PhD programs. But Foundation staff also thought that doing a survey would provide an opportunity to learn more about former and current students themselves. To be sure, elementary information was available on students' demographic profiles, but it was clear that learning more about why they chose the graduate school they did; the nature and extent of time commitments; students' own assessment of the advising they received; the extent of competition in their departments; the employment they took outside graduate school; the reasons for and the timing of their leaving doctoral programs, and, if they did so, their publication records and job histories after leaving; whether they married or not; and their route to tenure if they chose that goal would provide a far richer understanding of graduate education in the humanities than has heretofore been available.

#### THE GRADUATE EDUCATION SURVEY

The Graduate Education Survey was designed by Foundation staff and conducted by Mathematica Policy Research. Between November 2002 and October 2003, the 18,320 students who had matriculated at the treatment and comparison departments from 1982 to 1996 were surveyed. Of these, 13,552 responded, producing a response rate of 74 percent, which is remarkably high in this context, particularly for a retrospective survey. As might be expected, the response rate was higher for individuals who had completed their PhDs (81.3 percent) as compared to the rate for students still enrolled in their programs (75.8 percent), which in turn was higher than the rate for those who had left their programs (62.8 percent). The response rate for the last group was lower in part because 20 percent of those who left programs, many of whom had departed graduate study fifteen to twenty years earlier, could not be located. For the same reason, response rates differed by entry cohort, with the response rates declining the farther back in time the recipients had been graduate students. The response rates of the 1991–96, 1986–90, and 1982–85 entering cohorts were 77 percent, 74 percent, and 70 percent, respectively.<sup>26</sup> Thus, while some self-selection plainly occurred, these response rates are high enough to make us confident that the data do represent the graduate populations of the relevant universities in the time period under discussion.

The first section of the questionnaire asked students about entering their graduate programs (including why they chose the programs they did and 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 dissertation advisers and departments, the overall learning environment in the departments, the time it took them to complete different phases of their programs, and their publications—if any—while in graduate school and during the first three years after graduation. The third section asked questions about their experiences as research and teaching assistants, including the intensity, extent, and nature of those experiences. It also asked about the extent and nature of students' nonassistantship employment at various stages of their programs. The fourth section solicited information on degree completion (which, as we indicate below, allowed us to check the accuracy of the data on degree completion that the institutions provided) and information on the subsequent educational experiences of those who left doctoral programs. A fifth section sought demographic information, including the students' marital status and the number of children in their families during their graduate study years. The final section solicited information on the respondents' employment status six months after degree completion or departure from their programs, three years later, and again as of the survey date. Information on early career publications was also requested.

Just as the accuracy of the institutional database was consistently checked, so too were the data supplied in the GES. In particular, we were intent on reconciling respondents' GES replies about their enrollment status (whether they had received the PhD, were still enrolled, or had dropped out of the program) with the data the institutions supplied. These checks were critical in discovering which respondents had in fact earned PhDs when their institutions had no record of their having done so.

One more comment about accuracy is in order. As long as studies have sought data on the published productivity of scientists and scholars, investigators have been skeptical about the accuracy of self-reported survey data on publication counts. This led the Foundation staff members who had access to the names of survey respondents to compare the self-reported publications data for a sample of respondents to publications information obtained from websites and bibliographical indexes. We can report that in almost all cases the self-reported publications data were close enough to the objective measures that we felt confident in our ability to use the self-reported data for the entire sample. As far

as we know, this is the first time a validity check has been made on self-reported publications data.

## WHAT WE LEARNED FROM THE GEI

Taken together, the institutional databases and the GES are rich sources of information about graduate education. These databases have been analyzed by researchers at the Cornell Higher Education Research Institute and by Mellon Foundation staff. Details of our technical analyses and findings are reported in a number of journal articles, working papers, and a forthcoming book.<sup>27</sup> Here we summarize briefly some of our major findings concerning the impact of the GEI on attrition rates, completion rates, and times to degree completion in the humanities and related social science PhD programs; what characteristics of PhD programs in the humanities and related social sciences influence these outcomes, and how the GEI influenced these characteristics; what happened to students who left PhD programs prior to receiving their degrees; the early career job-market outcomes of new PhD recipients; and their graduate school publications and early career outcomes.<sup>28</sup>

### *Student Outcomes, Graduate Program Characteristics, and Their Interrelations*

Our analyses suggest that the GEI had modest effects on student outcomes in the expected directions: attrition rates and times to degree completion were reduced and completion rates were increased. These effects, we find, were driven in part by intentional reductions in the size of entering cohorts, which in turn permitted departments to become more selective in their admissions, as gauged by GRE scores. Reductions in cohort size also allowed improvements to be made in financial support over and above improvements that were attributable to infusions of Mellon Foundation funds.

Some improvements in financial aid also occurred as students became more likely to receive guaranteed packages of multiyear support upon admission. Universities undoubtedly moved in this direction in order to enlarge their chances of successful recruitment of students in response to increased competition in the market for new PhD students. The framers of the GEI did not anticipate intense competition; their hope had been to make financial aid conditional on satisfactory progress through the program. It is clear that market forces intervened and strict adherence to the GEI's conditional regime was replaced by the

inclination to make attractive awards to applicants. Although multiyear packages reduced the probability of students' dropping out early in their graduate careers, the same packages appear to have been associated with an increase in the probability of dropping out later on, thus leading to an unintended substitution of later attrition for early attrition. This is a finding bearing further exploration and discussion.

Analyses of the data collected in the GES identified different routes through which the characteristics of graduate programs in our sample influenced student outcomes. We find that improving advising and clarifying program requirements are associated with reduced attrition. Departmental expectations about the nature of dissertations also have strong effects on attrition, even when students are in the early years of their doctoral programs. In particular, departments that encourage students to finish their dissertations as quickly as possible have lower rates of attrition, whereas departments that emphasize the importance of students polishing their dissertations and publishing their work prior to graduating have higher rates of attrition. Similarly, graduation probabilities are higher when advising is improved and when departments expect that the dissertations will be completed promptly.

The GES data also reveal that students and their faculty advisers confront a trade-off. Students who publish while in graduate school are more likely to obtain tenure-track appointments at four-year institutions upon graduation. Those who publish while in graduate school are also more likely to publish soon after receiving their degrees. To the extent that faculty members are concerned about their students' career success and are eager for them to publish, advising students to publish while in graduate school may be good advice, even if doing so increases the likelihood that some students will drop out and it also increases the time it takes for others to finish.<sup>29</sup> Put simply, although the GEI designers had the explicit goals of reducing times to degree completion and attrition rates, it is not self-evident that both could be pursued at once, nor that they are consistent with promoting students' later academic careers. There are also indications that some faculty members did not accept the legitimacy of reducing time to degree completion and did not encourage students to finish quickly.<sup>30</sup> Faculty members' inclination to do what they think best for their students should be recognized when future efforts are made to change doctoral programs.

Our analyses also helped identify the effects of the GEI on important characteristics of doctoral programs. On average, the GEI seems to have prompted increases in seminar requirements, higher expectations for summer work keyed to students' progress, and clarification of program expectations. In smaller departments the GEI was associated with students being encouraged to finish their dissertations promptly, whereas

in larger departments (in which financial support prior to the GEI was especially scarce) it was accompanied by improved financial support. The effect these GEI-related program characteristics had on student outcomes was modest; there remains considerable variation among departments in the presence of these characteristics even now. Hence, it is possible that there is still room for changes to be made that will have beneficial effects on student outcomes in the future.

Finally, it is worth emphasizing that financial factors are not the primary reason that students drop out of PhD programs, as some suppose. To be sure—except in the case of the rare individual of independent means—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, those that include fellowships in each of the first six years that students are enrolled in their PhD programs—are associated with substantial rates of attrition. Amply supporting graduate students but doing nothing else will not solve the attrition problem.

### *Those Who Leave PhD Programs*

Stated simply, leaving a PhD program without a degree does not spell failure, at least as far as the reports our sample members provided. The unique nature of the GES allowed us to ascertain what actually happened to former students who dropped out. Indeed, over 10 percent of the “dropouts” in the GES sample ultimately received PhDs from different departments, with many of these people receiving their PhDs in fields other than the ones in which they were initially enrolled.<sup>31</sup> Individuals who drop out early in their graduate programs are much more likely to receive PhDs elsewhere than those who leave later on. In addition, almost 20 percent of those who dropped out went on to receive professional degrees, including, among others, law and MBA degrees.

We also find that the incidence of what appears to be downward occupational mobility among those leaving their PhD programs is large, but only temporarily so. Although 10 percent of dropouts were employed in clerical and administrative positions six months after departure from graduate school, by the time three years had elapsed this percentage had been reduced and the majority was employed in professional occupations. This is far from the popular imagery of the long-term results of attrition.

### *Job Outcomes after the PhD*

Much more so in the humanities and related social sciences than in the science and engineering fields, obtaining an academic job is a near

necessity if one is to work in one's field, and obtaining a tenure-track position at a four-year institution is the prime measure of early-career success for new PhDs. About 30 percent of the cohort that received degrees between 1998 and 2000 found employment in tenure-track positions at four-year institutions six months after receipt of their PhDs. Yet three years later, 52 percent of this same cohort had tenure-track positions at four-year institutions—a considerable increase. To be sure, the 30 percent holding tenure-track jobs right out of graduate school was slightly smaller than those who had graduated earlier in the decade, and so was the 52 percent who had such jobs three years later. These data suggest that there is considerable early-career mobility for new PhD holders; indeed, about 50 percent of new PhD holders who had full-time non-tenure-track positions six months after receipt of the degree had moved to full-time tenure-track positions three years later. The data also suggest that tenure-track jobs were increasingly going to those who had accumulated some post-doctorate experience and had assembled a stronger set of credentials than those of new PhD holders.

As time to degree completion increases, the probability of obtaining a tenure-track position within three years of receiving a PhD monotonically declines, but only for those who took eight years or more to complete their studies. This is an important finding, for it demonstrates that time to degree completion matters in getting coveted tenure-track posts, but only if it exceeds the seven-year threshold. As we have already noted, publishing while in graduate school enhances job candidates' chances of obtaining tenure-track positions and it also enhances their chances of attaining tenure within fifteen years of entering graduate school.

### *Graduate-School and Early-Career Publication*

About 40 percent of respondents to the GES published while in graduate school or had at least one book or refereed article accepted for publication. Within three years of receiving the PhD, about 67 percent had published one or more papers or books. Publishing while in graduate school is an important predictor of publishing soon after earning the degree, and PhD holders who reported that their departments expected them to publish while in graduate school published more often early in their careers than did other PhD holders in our sample.

As we have noted, this may help explain why our estimates of the effects of the GEI on time to degree completion are so modest. Faculty members at these top programs appear to be more concerned about preparing the next generation of scholars than they are about the time it takes for their students to complete their degrees. It is only among those who took more than seven years to complete their degrees that long

time to degree completion is inversely correlated with tenure-track job probabilities. Faculty advisers, in this one sense, are quite realistic about not pressing for shorter degree-completion times. However—and this is important—taking more than seven years to complete one's PhD is far from unusual; in fact, over 50 percent of the degree holders in our sample took more than seven years. Thus, while advisers appear to have little incentive to press students who are apt to complete the degree within seven years, pressing those who are still in course eight years after matriculating seems justified, though the positive effect of publishing on the quality of jobs that degree recipients get must be weighed against the adverse effect of longer degree-completion times on job opportunities for this group.

Having said this much, we find that students who completed their degrees in five years were more likely than others to publish while in graduate school. As times to degree completion increase, probabilities of publishing while in graduate school decline in the GES sample; this is likely to be the outcome of a selection effect. With other factors held constant, the more talented and motivated that students are, the greater the likelihood of their publishing, and the shorter the time it will take for them to finish their degrees. It thus follows that the students with the shortest times to degree completion are those who have published most often while in graduate school, and are those who are most apt to be employed in tenure-track positions after receiving their degrees.

Although the explicit goal of the GEI was to enhance the effectiveness of graduate programs, its framers did not explicitly make transforming students into scholars who would contribute to the extension of knowledge a prime objective or an important indicator of program effectiveness. Nonetheless, it is reassuring, as we have noted, that as many as 40 percent of the students in the GEI published while still in graduate school, and further reassuring that we estimate that the GEI increased the probability of students publishing while in graduate school by roughly 20 percent to 25 percent. This is no small accomplishment, and one not typically taken into account in assessing graduate education. 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.

### *General Lessons*

The GEI has confirmed that the microenvironments of departments matter greatly in doctoral education in the humanities and related social sciences. Many of the departmental characteristics that influence graduate students' progress are controlled by departments, not by the graduate

dean or other central administrators. Although there is a strong role for graduate deans to play in improving graduate education—and the study by Daniel Denecke, Helen S. Frasier, and Kenneth Redd in this volume takes up this matter (see chapter 2)—future efforts to improve graduate education should focus on departments' roles. And one should not underestimate the difficulty of persuading faculty members to “buy into” program changes and, ultimately, of transforming departmental cultures. Our analyses suggest that innovations that are initiated at the departmental level are much more likely to be supported by the faculty than those that are initiated top down.

It is also clear that after programmatic innovations are introduced they evolve over time. Sometimes this is due to faculty turnover—for example, the departure of a concerned faculty member and/or the arrival of another. Sometimes shifts occur in response to external competitive pressures, as in the instance of the spread of multiyear guaranteed financial-aid packages. This inclination for practices and procedures to evolve in departments makes it important for departments and graduate deans to keep regular track of indicators of departmental performance by collecting relevant data so that progress toward desired outcomes does not inadvertently erode.

One important benefit of the GEI was to encourage institutions and departments to collect such data, and a number now do so routinely. An important role of graduate deans is to monitor and standardize data collection and, where appropriate, to respond to the messages they carry. The National Research Council evaluation of doctoral programs that is scheduled to be released in late 2008 is similarly pressing departments to collect such information, and we view this effort as very important.

The GES has shown that retrospective surveys of current and former students can provide detailed information on multiple characteristics of graduate programs unavailable in institutional records. Our analyses have shown that these characteristics can be aggregated using factor analysis into a smaller number of underlying factors and, with data from multiple departments and multiple entering cohorts, analysis of factors that most strongly influence graduation and attrition probabilities can be undertaken.

We strongly believe that similar analyses may be profitably undertaken in other fields of graduate study (such as science and engineering) where entirely different characteristics of graduate programs may prove important.<sup>32</sup> 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-style 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 we now think would have been useful in evaluating graduate education. First and foremost, the GEI lacked basic data on the faculty. No usable information was available on the number of program faculty in each department and their stability over time. No information was available on the workloads faculty members shouldered in advising and dissertation sponsorship, nor were data available on the presence of incentives for faculty to mentor doctoral students (for example, workload credit for supervising dissertations). No information was collected on the match (or mismatch) of students and faculty advisers according to their research interests, or on their gender and ethnicity, and no information was sought on advisers' past success in placing their students. We know from experience how time-consuming it is to collect satisfactory information on the number of faculty members in residence in each department. Collecting more thorough information on the faculty would surely prove a formidable task, but nonetheless a useful one. It goes without saying that a survey of faculty members' views about graduate programs, parallel to the GES survey of graduate students, would have been highly desirable. The annual narrative reports the departments provided were highly instructive but clearly no substitute for more systematic data.

Finally, we believe it is important to have modest expectations about the likely effects of foundation-related efforts to improve doctoral education. If the GEI experience can be generalized, a host of factors can emerge and coalesce to make it difficult to achieve all of a foundation's objectives.<sup>33</sup> Not only do the objectives of individual faculty members sometimes differ from those held by foundations or graduate deans, but competitive pressures in the market for doctoral students may also press for the adoption of policies at odds with those advocated by foundations. An excellent example of such unanticipated consequences coming into play can be found in the case of the GEI, in which multiyear guaranteed support rather than incentive-based financial aid became the norm. Similarly, conditions prevailing in the academic market may affect the decisions graduate students make about the timing of dissertation completion and job searches. As the difficulty of finding tenure-track employment has increased and brought with it growing concerns about getting a desirable job and losing health insurance, housing, and library privileges, doctoral students probably correctly perceive that rushing to finish their degrees quickly might not be in their best interests. Put simply, the "law" of unanticipated consequences prevails in efforts to change graduate education, as it does in all other domains in which purposive change is sought. It is difficult to predict all of the consequences of programs and how these will interact with the changing world. To our minds this does not counsel inaction, but it does counsel

the need for regular monitoring of the desired outcomes and the need for continuing alertness to potential second- and third-order effects. The GEI was unique in scale, investment, duration, and departmental focus among efforts to improve graduate education. It was also unique in its intensive concern with monitoring its effects, not simply after the fact but from the outset and while it was under way. All these attributions make it *exemplary*—in the precise meaning of this word.