Graduate Economics Teaching of Core Microeconomics: Diversity, Knowledge Clusters, and Job Placement

Arthur Campbell

Uber &

Jonathan S. Feinstein

John G. Searle Professor of Economics and Management
Yale School of Management

Soonwook Hong

Yale School of Medicine

Sharon Qian

Citi Group

Trevor C. Williams

Yale University

April 2016

We are grateful for summer research support from the Yale School of Management. We thank Ben Gifford and Tiffany Lin for their research assistance. Soonwook, Sharon and Trevor all worked on the project during their time as students at Yale College. We are responsible for the content herein.

1. Introduction

It has been informally recognized for many years that different graduate economics programs teach different materials and present different perspectives on the field of economics. This kind of product differentiation in the field of higher education can have important downstream implications. For example, graduates of programs with significant curriculum differences may approach policy questions differently, generate and pursue different lines of creative research, and bring different conceptual lenses and knowledge to bear in applied careers.

In this paper we study empirically the diversity in what is taught in core microeconomics courses at a set of eight top U.S. economics doctoral programs. Our empirical analysis is based on detailed information we gathered from course syllabi and reading lists, and also in some cases direct interaction with professors and teaching assistants, for academic year 2009-10. We used this information to code the concepts covered in the required readings assigned in these courses, supplemented in certain cases by class notes; our coding is based on close reading of required chapters and sections of textbooks, articles and notes, identifying all concepts that are clearly defined, discussed, and explicated, including via worked out examples. The set of concepts we identified is detailed and we believe reasonably comprehensive, including more than 1300 concepts, including definitions, models, theoretical results, heuristics, and empirical methods and important empirical findings. We also develop a list of broader topics based on grouping concepts, thus creating a concept hierarchy.

Our empirical analysis reveals substantial diversity across programs in what concepts are taught in core microeconomics. Specifically, we present results describing the coverage of each program compared to the universe of concepts covered over all eight programs in our sample. Every program teaches some concepts not taught by any of the others. Conversely, only a very small percentage of concepts - 12% - are covered by all 8 programs. Most programs cover half or less of the universe of concepts covered by all 8 programs collectively. At the broader topic level there is more similarity of coverage but still considerable diversity, with most programs covering around 75 percent of the universe of broad topics, where we count a topic as covered if any concepts it includes are covered, a quite generous measure of coverage. We compute the overlap (correlation in concept coverage) between pairs of programs and show that in general it is modest, typically around forty percent. Our results are consistent with what we find listed on course syllabi. For example, some programs in our sample spend teach more general equilibrium, others teach more information economics, others focus more on empirical examples.

As a further part of our empirical analysis we perform a cluster analysis and show that the programs in our study cluster into two main "schools of thought" for the year we study. The grouping is reasonably intuitive in terms of informal discussion of programs. Finally, we use the cluster analysis to explore the impact of the knowledge clusters on job placement. For this analysis we build a two-stage econometric model in which the first stage is used to control for which program a student attends and the second stage is restricted to students graduating from one of the programs who have a job placement at one of the other 7 programs. The first stage results are interesting in their own right, revealing how a student being of foreign origin in particular influences which cluster he or she attends. The second stage is our main focus. While our data sample for the second stage is small, our results indicate that students are more likely to place at other schools in their own cluster, rather than the other cluster, indicating that schools of thought may replicate themselves through their hiring practices; we find this result in models that control for two other potentially important factors, program rank and geographical closeness of programs.

Our data and empirical analysis can both inform and spark debate within the profession of economics about what is taught in graduate economics programs, the degree to which standardization and diversity coexist, the existence and role of schools of thought in economics education, and the impact of schools of thought on hiring practices. More broadly, our analysis contributes to the broader discussion on higher education and how best to train students for careers in research, education, and policy.

We focus on core microeconomics for several reasons. Practically, microeconomics is part of the core program of study in every program in our sample and is taught as a full one year course sequence. This enables a quite clean, direct comparison across programs. It also means that, in studying job placement of graduating students, we can be certain that every student has taken (and passed) the core micro courses in their program and that programs considering hiring these students are also aware of this fact.

Second, microeconomics is arguably the most foundational field in economics. The topics covered in core microeconomics courses are likely to be viewed by most economists as foundational: they include utility functions and consumer demand, general equilibrium, market structure, uncertainty, imperfect information, basic game theory, social welfare, externalities, and standard empirical applications. Core microeconomics will form the basis for students' understanding, analysis, dicussion, and policy formulation of a wide range of economic issues, even if they go on to develop more specialized knowledge. Given how foundational these topics are, this is an area of study for which it might be expected that there would be a considerable degree of standardization of curricula across programs. Further, it is an area for which an argument can and indeed has been made in favor of standardization, on the grounds that all professionally trained economists should have an understanding of a certain set of core concepts; on this see our discussion below of several earlier studies and reports on graduate economic education. Taking these different factors into account, if standardization is to be argued for, or emerge in graduate economics, it is especially likely in the context of core microeconomics. Thus, a finding of a lack of standardization - and our results demonstrate that there is certainly far from complete standardization in core micro material – is especially powerful. In contrast, as an example, macroeconomics, while also part of the core of all the programs in our sample, is perhaps somewhat less foundational and less likely to be standardized, though that is an empirical question that would also be worthy of exploration.

It is also likely that core microeconomics curricula reflect programmatic values. A microeconomics core curriculum is likely to be determined, at least in considerable part, more centrally, rather than controlled by individual faculty; it is also arguably likely to be more stable over time. Further, given that it is so integral to graduate economics, the core micro curriculum is likely to reflect the ideological committments and values of a program. Due to these factors, diversity in what is taught across core micro courses is particularly likely to reflect systematic differences in perspectives across programs, as opposed to diversity for example in elective offerings, which may vary greatly from year to year. The fact that programs cluster into distinct schools of thought is particularly interesting in this regard, as it confirms the view that there are distinct traditions in terms of perspectives on economics and its teaching.

Of course, many students, especially those who focus on microeconomics fields, will be exposed to additional microeconomics concepts in electives and through research beyond what they are exposed to in their core curriculum. However, for all students, including even students who pursue microeconomics, what they are exposed to in their core classes may be important in sparking creative interests and preliminary ideas that in turn influence their subsequent creative development (Feinstein, 2006). Further, students who go into other fields of economics undoubtedly draw significantly upon what they have learned in their core micro courses.

This paper fits within both the particular stream of literature exploring what is taught in doctoral economics programs, as well as the broader literature on institutions, education and innovation. The literature on gradudate economics education includes two well known reports commissioned by the American Economics Association, Bowen (1953) and the report issued by the Commission on Graduate Education in Economics (COGEE, 1991, with summary of findings in Krueger, et. al. 1991). Both reports survey and evaluate curricula in general terms, not in detail as we do, and report on job placement and wages of graduates. The COGEE report makes two central points in regards the core curriculum taught in graduate economics programs. One is that there is an overemphasis on "mathematical technique" and tools over "economic substance" (p. 1044) (see Barone 1991 for a comparable statement made around the same time). We have not explored this in depth though our data could be used to shed light on this issue. The other is the finding of a large amount of apparently idiosyncratic diversity across programs. This is not inconsistent with what we report, but our findings are more readily interpreted as showing diversity across programs reflecting differences in programmatic values. Bowen also reports significant differences among programs and this seems to have troubled him. Indeed he states that it may be desirable to establish a "common core" and that he believes the core course in economic theory - probably the closest to the core microeconomics courses we study – should teach essentially uniform, standard material [Bowen (1953), pp. 42, 104, 109). In his preface he states that he found "substantial difference among the various graduate departments of economics not only in their practices, but even in their objective" (p. iv).

There have also been a series of papers on graduate economic education. McDonald (2009) discusses differences in philosophy and conceptual teaching between the University of Chicago and Yale during the 1950's and early 1960's and corroborates Bowen, showing that the views of economic systems and the kinds of frameworks viewed as central and taught were quite different between these two programs at this time. Interestingly, our findings fifty years later corroborate these earlier findings, albeit for more programs and different clusters of schools. In a well known study Colander and Klamer (1987) present results from a survey of economics graduate students in six programs conducted in 1985. Their results show significant differences across programs, with Chicago a notable outlier in terms of a higher level of perceived commitment to neoclassical assumptions (a follow-up study by Colander (2005) reports smaller differences). Certain broad curriculum topics are discussed, including rational expectations, imperfect competition, and price rigidities, but not in the level of detail we bring to the study of core micro courses. Lastly, the recent paper by Abito, Borovickova, Golden, et. al. (2011) present findings from a workshop held with recent graduates of economics graduate programs. They state that among workshop particpants, "the consensus was that the core should be designed to teach graduate students those aspects of economics that should be understood by all economists graduating today" but that most felt that was not being fully accomplished. This seems to indicate a desire for uniformity similar to Bowen's call for a common core. They do not provide details about the concepts covered in existing core courses.

Finally, the topic of this paper is part of the large literature on institutions, especially institutions supporting innovation and the knowledge economy (see for example Mokyr 2002). Graduate education trains individuals to pursue a career in their chosen field of study, as researchers, as teachers and administrators, and in the case of economics as important participants in public policy debates. Viewing core teaching as important for all of these, what is the appropriate balance in breadth and depth of coverage and between older, enduring models and perspectives and newer frontier ideas? Our data sheds some light on current practices, and theoretical modeling building on our approach can further illuminate this issue. Many economists seem informally to take the view that the core teaching in microeconomics should err on the side of teaching enduring frameworks. This perspective associated with the belief that it is important for the profession collectively to be confident that its graduates speak and use the "language of economics" correctly and effectively, with an assumption that this language is shared in common and agreed upon, a "common core." In fact we that in the education curricula of an academic field, like any industry, diversity will naturally emerge in product offerings and persist and that in fact some degree of diversity is socially appropriate. We do not expect the socially optimal configuration to have all producers producing the same product, in this case teaching the same material. Indeed as the concept of schools of thought emphasizes, there is unlikely to be complete agreement on what concepts are most important to a field and should be taught. Futher, core teaching is also important for future research which is inherently an activity that is more specific to the individual. While students who go on to research in microeconomics will learn much beyond the core, core teaching can spark particular creative interests that set them off on a path of develoment. Feinstein (2006) provides many examples of individuals whose main creative interests were sparked during their first year in graduate school, often in core classes. For all these reasons, we would not expect a common core to be optimal and certainly our empirical results show that it is not current practice.

The remainder of the paper is organized as follows. In the next section we provide a conceptual overview of forces lying behind diversity and overlap in graduate education. In section 3 we describe our construction of the concept map for describing what is taught in core curricula, and in section 4 we describe our statistical methodology. In section 5 we describe our data and in section 6 we present our empirical findings. Lastly, section 7 presents concluding thoughts.

2. Conceptual Overview of Diversity & Overlap

Graduate programs must choose what topics to teach in their core curricula. While we do not construct a formal theoretical model of program choices it is useful to ask the question, What factors lead a program to choose a given set of topics to teach? A related question also very relevant for our empirical analysis is, why do different programs choose to teach different topics?

Scare resoures of time and attention are the root cause why not all topics in a subject are covered in a core course. Core microeconomics typically contains one year of coursework balanced with a suite of other courses students take, and this effectively limits how much can be covered. Indeed microeconomics, with its long tradition, dense mathematical formalism for many topics, and very wide range of applications, offers a vast amount of material. Thus programs are forced to make choices regarding core micro: they do not teach many topics, and even within topics do not teach many specifics.

Multiple factors are likely to influence a graduate program in developing its curriculum, including supply and demand factors as well as commitments to programmatic values. Demand side factors are those that drive the goal of attracting the best possible students and placing students in good jobs, which in turn influences attractiveness of the program to future students. Graduate economics is preparing students for a range of careers, including research, careers as educators, and careers of practical engagement, including private sector careers and public management. Programs will differ in their pool of applicants, driven by their relative attractiveness to students as well as their attractiveness to different kinds of students. For example, higher ranked programs can be expected to be more attractive to all students and in particular to students hoping to pursue a research career, whereas lower ranked programs may be more likely to enroll students who hope to pursue a career in the private sector or government. As our empirical results document, foreign students may be attracted to somewhat different kinds of programs, possibly because they hope or expect to return to their home country and pursue a career in the public/political sector. In turn differences in the kinds of students programs attract may influence their choices of curriculum; for example, a program that attracts more research-oriented students may teach more mathematical theory than a program that attracts more students interested in private sector careers.

Supply side factors will also be important for curriculum choices. It will be less costly as well as more enjoyable for a faculty member to teach topics he or she knows well and that connect in some way with his or her research or other outside-of-teaching activities such as consulting or public service. Since different faculty have different backgrounds in terms of knowledge and different research and outside activities, it is natural to expect this to generate diversity in coverage. At the same time, a program selects which faculty teach in its core courses which could somewhat mitigate this factor and lead the core more to reflect programmatic values. Further, a program will provide guidelines and possibly incentives for faculty to teach what is viewed as an appropriate set of core topics, in principle taking into account faculty costs & utility, student learning costs, attracting the best students, and student placement.

Program traditions and committments to certain values may also have an important role in what is taught in core courses. Many programs have a tradition of teaching microeconomics from a certain perspective. For example, the University of Chicago is well know for its committment to certain foundational principles and approaches in teaching microeconomics, a legacy still followed in the year we study, 2009-10, during which Gary Becker continued to teach in the micro core. Yale, Harvard and MIT also all have strong, distinctive traditions and committments as McDonald documents for Yale. Senior faculty may have strong views, based on program traditions and their own personal values, about what it is appropriate to teach, which may influence what is taught by their junior colleagues. Overall, programs may adhere to certain traditions or “schools of thought” in their curriculum. We explore this issue in our empirical analysis through a cluster analysis, and the clusters we generate to accord reasonably well with informal views about programs. Programmatic committments may be balanced by a field-level perspective advocating for a common core of topics to be taught across programs, promoting the development of a common framework for basic economic analysis. Again, our empirical analysis allows us to explore this issue. We note finally, that, as opposed to the supply-side factor of an individual faculty member teaching what is easiest or best from his point-of-view, programmatic traditions tend to suggest stability over time in what a program teaches and in differences across programs.

The three main factors we have identified, taken together, predict that graduate programs will not teach the same set of topics. Our empirical analysis shows this and also documents the extent of overlap and diversity among programs, revealing that the overlap is in fact just moderate.

3. The Concept Map

Our analysis of core graduate microeconomics programs centers on the construction of a detailed empirical map of the concepts that are taught in each program in our study. We take a broad perspective on what the term "concept" covers. Concepts can be theoretical or practical/empirical in nature, though for core graduate micro they are mainly theoretical. Theoretical concepts include definitions, assumptions, principles, models, theorems, key steps for proofs (sometimes linked to auxiliary mathematical concepts), theoretical implications, standard examples or functional forms, and specific worked out examples linked to foundational concepts that are flushed out by the example. Empirical concepts include applications of theoretical concepts, links to practical real world examples or topics, tests of theories and especially important or general findings. Concepts are also grouped under broader topics; this generates a concept hierarchy.

We built our concept map through identification and close reading of the set of concepts presented in the required readings that are assigned students in their core microeconomics sequence. We focused on required readings for several reasons. First, required readings include concepts that all students in a class are expected to read and learn to the best of their ability. We expect that a typical, reasonably conscientious student will read the required readings, though of course we are not gauging how well any particular student, or even a class as a whole, learns the material covered in a reading.[12] Second, every course in our sample has assigned required readings, typically for every class session, making this a good foundation for building a map of what is covered in courses and a natural basis for comparison among programs.

There are two main other kinds of class materials that could be incorporated in a study like ours. One is material covered in class, such as material written on blackboards and oral class discussion. We chose not to focus on in class material mainly for practical reasons. Much classroom material cannot be obtained without sitting in on the class, something we have not done and which might not be welcomed by some instructors.[10] Classroom learning material also varies widely, from typed handed out notes - which in fact we do include as required readings - to hand-written notes a professor uses, to material written on chalk or white boards, to oral discussion; much of this would be difficult to code in a systematic fashion for our study, particularly discussion. In contrast, information about required readings is readily obtainable from class syllabi or in some cases instructors or teaching assistants. It can also be argued that required readings provide a more comprehensive view of what students in a course are exposed to than classroom material. In particular, while required readings typically cover similar concepts to what is covered in class

[12] In this regard it is important to make clear that our focus is on the concepts students are exposed to, not on determing what students learn or how well they learn it, which would require very different sorts of information, such as test results: We are concerned with what programs choose to teach and focus on, not on their effectiveness in teaching the material.

[10] We have found that some instructors do not want to share their lecture notes or do not have them in a form that is readily shared; so that if we were to attempt to gather this information it would be difficult to build a complete sample. However, we did obtain lecture notes from professors or students in a few cases where they were essentially the required readings; see our further discussion in the next Section.

(lecture notes we have access to confirm this in the specific cases we have compared), it is not unusual for the required readings to go into more depth than what is covered in class and in that regard the readings are actually a better measure of what students are exposed to - at least conscientious students. A second kind of material we could add is supplemental readings. We chose not to do this for two reasons. First, students are not required to do this reading and it seems a stretch to imagine that most students do so. Including supplementary readings would provide additional information about what more highly motivated students may read; but a typical student most likely does only a fraction of the supplementary reading, making this too inclusive a measure of what such a student is exposed to. Given we are not certain how much of the supplementary material students read, we believe this is not a good way to compare across courses. Second, courses differ a great deal in the amount of supplemental readings they list, from none to very long lists of articles, and we do not think it appropriate to take the view that this properly reflects true differences in knowledge exposure across programs. However, we acknowledge that programs that provide more supplemental materials might be viewed as following a somewhat different teaching philosophy and therefore exploring differences among programs in these materials might be useful in future work.

We built our concept map in 3 steps. First, we identified the universe of required readings across all courses for all programs. As an example, if a given text is used in two courses (whether from the same or different programs), and the first course assigns Chapter 4, sections (i), (ii) and (iii) as required reading and the second assigns Chapter 4, sections (ii), (iii) and (iv), we code Chapter 4, sections (i), (ii), (iii) and (iv). Likewise if any school codes an article, we add the article to the universal map; we code articles in their entirety. Second, the main step, we went through each required reading, section by section, and coded the concepts that are defined and explained therein in a spreadsheet which we call the concept map. We also built a conceptual hierarchy identifying broader topics that subsume narrower concepts. We discuss this step in more detail next. Third, we mapped each course to the set of concepts in the map that are covered in its required readings, then aggregated over all courses in a program to enumerate the set of concepts covered in that program.

We created the concept map meticulously, working through each required reading, section by section. We focused on identifying important definitions, principles, and theoretical results, which are the heart of core microeconomics courses. In addition, we included theoretical models and examples and empirical models and results in cases in which these provide significant additional content. Identifying relevant concepts is most straightforward for textbooks, since explicating concepts is the principle aim authors have in textbooks. In fact the vast majority of the concepts we code come from textbooks, as the vast majority of required readings that are assigned are from textbooks. For articles, we identified in general only one or a few concepts in each article; these are typically either advanced theoretical concepts or worked out empirical principles or important empirical findings viewed as significant in the field. In addition we also constructed a link mapping that links each concept to all readings that cover that concept.

The key judgement in developing the concept map is identifying relevant concepts. In general we followed the lead of the text, especially for textbooks, which generally explicitly state what the important concepts or takeaways are in a section. As a result for many concepts the identification of the concept is straightforward, particularly in textbooks, because the text identifies the concept explicitly, generally as a new concept being introduced. For worked out examples we again followed the lead of the text, which for a textbook will often state what the main point or takeaway is from the example. For example, a worked out example might be used to make a general point about how two variables commonly relate to one another and the text will state this. Rarely, we added additional concepts, such as a key diagram, if we believe this is general knowledge a student would be expected to take away from the reading. Identifying concepts is more challenging for articles, for which a high level of background knowledge may be assumed. We focused on identifying concepts that are defined and worked with extensively in the article, not concepts assumed as background knowledge. If the main point of an article is to establish an empirical claim we coded that as a concept; if the main point is to show the properties of a model and its solution we coded these properties.

In building the concept map we did not include some material that we view as less significant. We generally did not include any concepts discussed only in introductory passages in a chapter or article, preferring to focus on material in the body of the chapter or article. Also we did not include concepts or examples mentioned only in passing, or in footnotes. We were also parsimonious with regards extended examples, which in some cases run to several pages. Here we typically coded only one or two items, the main points of the example. Overall, our approach casts quite a fine net, as we document below when we discuss the number of concepts in our map.

For purposes of illustration of what our concept map looks like we discuss the entries for risk aversion. In our concept map we have 5 distinct concepts under this topic. The first is the basic definition of risk aversion and the statement that for a risk averse person his utility for the expected monetary value associated with a gamble is larger than his expected utility for the gamble itself. The second entry is the statement (and explication) that the utility function for a risk averse person (over the relevant range) is concave. The third entry is for diagrammatic illustrations, typically showing concavity and how it relates to risk aversion; we include diagrams because they are present in many treatments and in our judgement are important in helping a student grasp the nature of a utility function for a risk averse person. The fourth entry is the definition of risk-neutrality as a linear utility function, the expected value. The fifth entry is for the definition of risk-seeking and associated utility function. Our entries for risk fit within a larger section of the concept map coding concepts referring to utility functions defined over monetary values. In addition to the entries for risk described above, the broader topic includes entries for measures of risk aversion, including absolute and relative risk aversion, entries for comparing two lotteries to determine if one is riskier than the other according to some measure of risk (stochastic dominance), entries for certainty equivalents, including the basic definition, definitions for the risk premium and probability premium and for how to construct certain equivalents and do consistency checks on constructed utility functions, and a further set of entries for applications to financial assets and portfolio choice, including mean-variance analysis, allocation of wealth between a risky and riskless asset, and insurance.[11] In total there are about 50 concept entries under this

[11] Not all of these concepts may seem natural core concepts to readers. For example, in the area of financial portfolio analysis it is clearly a judgement call an instructor must make as to how much to cover in a first-year core microeconomics course. Most texts mention this problem and the leadings texts work out the example in which there is one risky and one riskless asset and use first-order conditions to show that a risk-averse investor will always invest some amount in the risky asset as long as its expected return exceeds the expected return on the riskless asset. But texts (and class notes) diverge beyond this. One text in our empirical data, Kreps, works out an example with two risky assets and one riskless asset to show that in this case the analysis is considerably more complicated and depends on covariance terms (we note that this example is however in smaller print, though a conscientious student may be viewed as likely to attend to it). In contast the leading text by Mas-Colell, Whinston and Green does not discuss this in such length in the chapter on choice under uncertainty, however has a relatively lengthy section on pricing of financial assets later in their book in the chapter on general equilibrium under uncertainty.

topic heading of utility over money.

An advantage of defining a large number of concepts as we do is that we can aggregate concepts up to define broader, encompassing concept levels. To this end we defined a set of topics such that each topic encompasses a set of concepts, as well as a top level of very broad topics each of which includes a set of topics. For example the concepts above fall under the broad topic "risk" which in turn lies under the very broad topic "choice under uncertainty."

The final step in our data set-up was identifying the set of concepts covered by the required readings for each program in our study. For each program we identified the required readings for each core micro course in the program. We then used our link mapping from readings to concepts to identify the set of concepts covered by these readings. Aggregating over courses in the program generates a set of rows in our concept map which defines the set of concepts covered in the program.

4. Empirical Measures & Models

We use our concept map to evaluate the degree of overlap and diversity in concepts students are exposed to among the programs in our study. We assess overlap and diversity through three distinct approaches: (i) correlations in coverage between pairs of programs; (ii) centrality of concepts, measuring how many programs cover each concept; and (iii) clustering programs based on the concepts they cover.

Correlation. Define $\mathcal{A}$ to be the set containing the universe of concepts taught across all programs. Let N be the cardinality of this set, the total number of concepts taught across all programs. For program j, for concept i define $X_{ij} = 1$ if program j covers concept i and $X_{ij} = 0$ otherwise.

Our first approach for measuring overlap between programs is via computing the correlation coefficient between pairs of programs. Define r_{jk} to be the correlation between programs j and k. Note that $r_{jk} = r_{kj}$. r_{jk} is defined as:

$$r_{jk} = \frac{\sum_{i \in \mathcal{A}} (X_{ij} - \theta_j)(X_{ik} - \theta_k)}{\sqrt{\left(\sum_{i \in \mathcal{A}} (X_{ij} - \theta_j)^2\right) \left(\sum_{i \in \mathcal{A}} (X_{ik} - \theta_k)^2\right)}}.$$

Here θ_j is the fraction of concepts covered in program j:

$$\theta_j = \frac{\sum_{i \in \mathcal{A}} X_{ij}}{N}.$$

We note that the correlation can be negative, but is not for our data sample. Also we note that correlation can also be computed at the level of broad topics.

In the Appendix we define two additional, related measures of overlap. One is the ratio of the number of concepts both schools teach (intersection) to the set of concepts taught by either school (union). The other is an asymmetric measure that measures, for program j in comparison with program k, the ratio of the number of concepts that are covered at j that are also covered in k divided by the total number of concepts covered at j.

Centrality. Our second approach for asssessing overlap among programs is by computing the centrality of concepts. To do this we compute how many schools cover each concept, and then covert these numbers into a table showing the distribution of how many concepts are covered by all programs, by some programs, or by just a single program. This provides a different way to assess convergence and diversity across programs as well as providing insight into how many concepts and which concepts are central to core microeconomics teaching and how many and which are more peripheral in the sense of being covered in only one or a few programs.

Clustering. Finally, we apply a hierarchical clustering algorithm to cluster programs based on similarity in the concepts students are exposed to. Hierarchical clustering algorithms start with all programs separate and successively combine the closest two clusters, either individual programs or previously combined clusters of programs (Agrawal, et. al. 2005). The output from a hierarchical clustering algorithm is the order in which programs and clusters are combined until reaching a single cluster containing all programs. Consider an example with 5 schools (Schools A, B, C, D, and E). A potential output from a hierarchical clustering algorithm is

$$\begin{aligned} & \{A, B, C, D, E\} \\ \rightarrow \quad & \{A + C, B, D, E\} \\ \rightarrow \quad & \{A + C, B, D + E\} \\ \rightarrow \quad & \{A + B + C, D + E\} \\ \rightarrow \quad & \{A + B + C + D + E\} \end{aligned}$$

Here the first step in the algorithm combines A and C. The second step combines D and E, the third cluster $A + C$ with B, and finally $A + B + C$ with $D + E$. In our empirical analysis we show the cluster sequence for our data via a dendogram graph.

The clustering alogrithm proceeds at each step by calculating a measure of similarity between each of the elements at that step - individual elements or clusters – then combining the two elements that are the most similar. As an example of how the cluster sequencing above may be generated consider the following data table which shows the set of concepts covered by each of the five programs:

$$\begin{array}{c} \\ A \\ B \\ C \\ D \\ E \end{array} \begin{array}{c} \begin{array}{ccccc} A & B & C & D & E \end{array} \\ \begin{pmatrix} & 5 & 2 & 6 & 7 \\ & & 4 & 5 & 5 \\ & & & 5 & 6 \\ & & & & 4 \\ & & & & \end{pmatrix} \end{array}$$

Based on this table, the first cluster step combines A and C since the distance between them is 2, and this is the smallest distance between any pair of programs. The data table with this cluster now looks as follows:

$$\begin{array}{c} \\ A + C \\ B \\ D \\ E \end{array} \begin{array}{c} \begin{array}{cccc} A + C & B & D & E \end{array} \\ \begin{pmatrix} & 4\frac{1}{2} & 5\frac{1}{2} & 6\frac{1}{2} \\ & & 5 & 5 \\ & & & 4 \\ & & & \end{pmatrix} \end{array}$$

Based on this table, D and E are combined. In turn this generates the table

$$\begin{array}{c} \\ A+C \\ B \\ D+E \end{array} \begin{array}{c} \begin{array}{ccc} A+C & B & D+E \end{array} \\ \begin{pmatrix} & 4\frac{1}{2} & 6 \\ & & 5 \\ & & \end{pmatrix} \end{array}$$

At the next step cluster $A + C$ is combined with B because these are the closest. The last step combines clusters $A + B + C$ and $D + E$.

There are two important choices that are required to implement the cluster analysis. The first is a distance measure; we denote this by $d(A, B)$, which describes how close two elements A and B are together. We use the Euclidean distance:

$$d(A, B) = \frac{1}{N} \sum_{i=1}^{N} (X_{Ai} - X_{Bi})^2$$

Here recall that X_{Ai} is 1 if program A covers concept i and 0 otherwise.[10]

After the first clustering step a second choice is required for how to treat an element such as $A + B$ which was formed at a previous step when calculating its similarlity to either an individual program or another cluster. We use the average distance method which calculates the average distance between the programs in each cluster. For instance the calculation of the similarity between two clusters $A + B$ and $D + E$ would be $\frac{d(A,D)+d(A,E)+d(B,D)+d(B,E)}{4}$.[11] . The sequence of calculations above implements this methodology. We see that in step 2 the calculation of the distance between the element $A + C$ and the remaining elements B, D, E is the average of the pairwise distances from A and C to each of these elements from step 1.

Job Placement: Link With Concept Coverage Map

The second branch of our empirical analysis focuses on student job placements and how they are related to conceptual closeness between programs. We focus on placements of students *from* one of the programs in our study *to* one of the other schools (we omit students who place in their home school because there may be many factors at play in this outcome that are outside our analysis). The hypothesis we are interested in exploring is that students will be more likely to place at a school for which the econ grad program is relatively close to their own graduate program in terms of core micro concept coverage.

As part of our empirical analysis as robustness checks we consider several other factors that may be important in job placement. First, we consider two additional explanatory variables in some specifications: (i) the role of geography, reasoning that a student may be more likely to place at a program in the same metro area; and (ii) the importance of ranking, considering that students may be more likely to place at programs having a similar ranking. We divide schools into two groups for this, the top 4 programs and the bottom 4 in our sample, based on standard rankings from US News & World Report and the National Research Council for placement years for our job candidates.

[10] Another widely used measure is the city-block measure, which adds up the absolute differences between A and B on each concept/dimension. For two individual schools this is simply the number of concepts which are taught at one school but not the other.

[11] Some other methods are: the centroid method which calculates the centroid of all the schools in an element before calculating the distance between elements based on the centroid of each; Maximum or minimum linkage methods which respectively use the maximum or minimum distance between any school in either element as the measure of distance.

A further concern is that there is a possible selection effect based on which program a student attends. Students have not been randomly assigned to programs, but rather have chosen to apply to certain programs, been accepted and then chosen to attend the program they attend. Thus a student from one program whose job placement ends up at a program that teaches conceptually similar micro material may not be a good match only because they have learned microeconomics in a way that fits with the program they place at, but also because of the type of scholar they are and personal attributes they possess that have made them a good match at both their graduate program and the program they place at. We control for this possible selection match effect in our analysis via estimation of a two-stage model that explicitly incorporates as its first stage a model of graduate program placement.

Given the selection concern, we focus on the cluster analysis for grouping schools because it allows us to fit a tractable two-stage model. In particular, we use the two main clusters for our analysis, denoted a and b, with m_a schools in cluster a and m_b in cluster b. The structure of the model is then as follows. The first stage models which program a student attends and is estimated over all students who attend one of the programs in our study and enter the job market in the years for which we collect job placement data. The second stage again is restricted to students whose job placement is at one of the programs in our study (but not their home program). The pool of students is substantially larger in the first stage, which helps with identification to control for the impact of the first stage graduate program placement on the second stage job placement model. We specify the model as a bivariate random effects model. Define random variables u, v to be drawn from a bivariate normal distribution with means 0, standard deviations 1 and correlation ρ. The parameter ρ is the key parameter that models the link between first stage graduate program placement and second stage job placement. We define u to be a shifter that increases the probability a student will enter a graduate program in cluster a, and similarly define v to be a shifter that increases the probability, for a student who places at one of the other programs in our study, that the student places at a program in cluster a.

The model specification is then as follows. For a student who attends a graduate program in cluster a but does not place at one of the other programs in stage 2 the likelihood integrates just over u and is:

$$\int_{u=-\infty}^{\infty} \frac{m_a e^{\beta_{1a}+\pi z+u}}{m_a e^{\beta_{1a}+\pi z+u}+m_b}\ \phi(u)du$$

where β_{1a} is the intercept for cluster a, z are explanatory variables that influence the probability of attending a program in cluster a, π are the associated coefficients, u is the random effect with ϕ the standard normal density function, and cluster b is the baseline cluster. The probability the student attends a program in cluster b but does not place at one of the other programs is then:

$$\int_{u=-\infty}^{\infty} \frac{m_b}{m_a e^{\beta_{1a}+\pi z+u}+m_b}\ \phi(u)du.$$

The factors m_a and m_b are present because each school in each cluster has an equal likelihood in the model and there can be different numbers of schools in the two clusters.

For a student who attends a program j in cluster a and then places at another program k in a the likelihood is:

$$\int_{u=-\infty}^{\infty}\int_{v=-\infty}^{\infty} \frac{m_a e^{\beta_a+\pi z+u}}{m_a e^{\beta_a+\pi z+u}+m_b}\ \frac{e^{\beta_{2a}+\beta_c+I_r\beta_r+I_m\beta_m+v}}{\sum_{h\in a,h\neq j} e^{\beta_{2a}+\beta_c+I_r\beta_r+I_m\beta_m+v}+\sum_{h\in b} e^{I_r\beta_r+I_m\beta_m}}\ \phi(u,v)dvdu$$

Here β_{2a} is the intercept for placing at a program in a, β_c is the dummy variable for placing at a program in the same cluster as one's graduate program, I_r is a dummy variable set to 1 if the program the individual places at shares a similar rank as program j and has associated coefficient β_r, and I_m is a dummy variables set to 1 if the program the individual places at is in the same metro area as program j and has associated coefficient β_m. In this specification $\phi(u,v)$ is the bivariate standard normal density function with correlation ρ. For an individual who attends a program j in cluster a and places at a program in b the likelihood is:

$$\int_{u=-\infty}^{\infty}\int_{v=-\infty}^{\infty} \frac{m_a e^{\beta_a+\pi z+u}}{m_a e^{\beta_a+\pi z+u}+m_b} \frac{e^{I_r\beta_r+I_m\beta_m}}{\sum_{h\in a,h\neq j} e^{\beta_{2a}+I_r\beta_r+I_m\beta_m+v}+\sum_{h\in b} e^{\beta_c+I_r\beta_r+I_m\beta_m}}\,\phi(u,v)dvdu$$

Analogous likelihoods are specified for individuals who attend a program in cluster b and place either at a program in cluster a or a different program in b.

The key check on our model is the estimate of ρ: If ρ is estimated as positive and statistically significant that indicates that a student who is more likely to attend a graduate program in cluster a is also more likely (conditional on placing at one of the other programs) to place at a program in cluster a due to innate factors that influene his/her placements, not specifically due to conceptual closeness. The variables z are the key exogenous shifters that allow identification of ρ separate from β_c. Thus the model is strongest if one or more of the π coefficients is statistically significant. In addition, z variables that are statistically significant are of independent interest.

In our empirical results presented below we find that ρ is not statistically significantly different from zero. Given this result, we can also simplify the model and run the second stage by itself. In that case the data is restricted to students who place at one of the other schools in our dataset, meaning either at a school in his own cluster or in the other cluster, generating a pair of likelihoods for each cluster or origin. The specification is as follows. For a job candiate who is in a program j in cluster a the likelihood of placing at another program k in cluster a is:

$$\frac{e^{\beta_c+\beta_a+I_r\beta_r+I_m\beta_m}}{\sum_{h\in a,h\neq j} e^{\beta_c+\beta_a+I_r\beta_r+I_m\beta_m}+\sum_{h\in b} e^{I_r\beta_r+I_m\beta_m}}$$

The likelihood of placing in a program in cluster b is:

$$\frac{e^{I_r\beta_r+I_m\beta_m}}{\sum_{h\in a,h\neq j} e^{\beta_c+\beta_a+I_r\beta_r+I_m\beta_m}+\sum_{h\in b} e^{I_r\beta_r+I_m\beta_m}}$$

Analogous likelihoods are defined for candidates who attend a program in cluster b.

We also fit a model based on the pairwise correlations in concept coverage across programs. Due to the fact that there are eight programs, it is not practical to estimate a two-stage model for this analysis. Thus, we specify and estimate a standard multinomial logit model for job placement at one of the other seven programs in our sample, and restrict the sample to only those students who place at one of the other seven schools. The specification is as follows. For a job candidate from program j the likelihood of accepting a position at program k (not equal to j) is:

$$\frac{e^{\beta_{corr}r_{jk}+I_r\beta_r+I_m\beta_m}}{\sum_{i\neq j} e^{\beta_{corr}r_{ji}+I_r\beta_r+I_m\beta_m}}$$

where β_{corr} is the coefficient of interest.

5. Data

We study a set of eight U.S. doctoral economics programs that are widely viewed as among the strongest programs in the U.S. and the world. The 8 universities that house these programs are listed in Table 1. Within the U.S. they are geographically diverse and informally are viewed as embodying somewhat different philosophies towards economics and its teaching.[3] The eight programs occupy 8 of the top 9 positions in the U.S. News & World Report rankings for economics Ph.D. programs for 2009 and again for 2015; the only program in the top not included is Princeton.[4] They are also highly ranked in the more complex multi-dimensional rankings of the National Research Council (2011). These programs spawn many outstanding researchers in the field, though of course outstanding researchers also emerge from other programs, and many individuals who go on to successful, influential careers in the public sector. What students in these programs are taught in the field of microeconomics is important for the further development of the field, for how microeconomics is understood and used in the field of public policy, and for the further transmission of microeconomics concepts to subsequent generations of students. Notwithstanding this, we recognize that this is not a random sample of programs and it would be useful to conduct a study like ours for a wider set of programs; this would enable comparison across a wider range of philosophic schools and different typical career paths of students.

For each school in our sample we identified the core microeconomics courses required of students in the program for academic year 2009-10. In general these courses are listed in program descriptions, we also had access to course listings. In identifying courses we did not include courses that teach mathematics for economists. For all programs one full academic year of core microeconomics is taught.[12]

Our primary source of information about required readings was course syllabi that we obtained either from websites or instructors. Most courses have assigned required texts; some also assign required articles. For most courses the syllabus lists required readings for each class session; these required readings include chapters and sections from required texts as well as articles. Whenever a syllabus was not clear about what was required we contacted the instructor or teaching assistant and invariably they were able to clarify the situation for us.[5] In a few cases the primary reading source for a course was notes prepared by the instructor and in those cases we obtained the notes directly from the instructor. There were a number of courses for which a few class sessions were based exclusively on class notes; in those cases we obtained the lecture notes for those class sessions from the instructor.[6]

---

[3] For an interesting discussion comparing what was taught at two of these schools, Chicago and Yale, 50 years ago, see McDonald (2009).

[4] http://grad-schools.usnews.rankingsandreviews.com/best-graduate-schools/top-humanities-schools/economics-rankings. We chose not to include Princeton mainly to have somewhat more geographic variation.

[12] For all schools but Penn the courses were taught in sequence in the first year of the program - at Penn the required courses were taught concurrently in the fall of the first year for 2009-10.

[5] There were three courses for 2009-10 for which we relied on a syllabus from an adjacent year. For Chicago we used a syllabus for course 301 and the first half of 302 from the year before. This course has been taught for many years by Gary Becker and Kevin Murphy and we were assured by the teaching assistants for this course for 2009-10 as well as other years that the syllabus was identical. For Penn, for course 203 part (i) we used a syllabus for 2010, the following year. The syllabus for 2009 lists the same topics as the syllabus for 2010, but does not provide a detailed list of assigned readings, whereas the syllabus for 2010 does. The instructor for 2010, Qingmin Liu, told us that he taught the same material and assigned the same or very similar readings. For Stanford, for course 203, taught in winter of 2010, we relied on the syllabus from the preceding year. The course was taught by Doug Bernheim who had taught the course in the preceding year and we are confident material was very similar for 2010.

[6] In a few cases lecture notes were available online; but in most cases we contacted the instructor and they provided the necessary materials.

In making judgements about whether or not a given readings should be viewed as required, we examined the language used in asigning the reading. In many cases instructors rely on a single primary text and clearly state that all other readings are supplemental and not required. In a few cases they have two or three primary texts, often assign readings alternately from these texts for different parts of the course. The most difficult cases are those in which one reading is not as clearly identified as the primary required text. In these cases we generally coded as required readings that were clearly assigned in relation to given class sessions or units, but did not code as required readings that were assigned more generally for the class as a whole - for example books that are mentioned as useful references.[7]

Appendix Table A1 lists the main textbooks required across the set of core courses in our sample. The second column refers to number of concepts that are taught linked to these texts - we define this measure below. The dominant text is *Microeconomic Theory* by Mas-Colell, Whinston and Green published in 1995. There are several other important texts, especially for game theory.

The instructors in the courses in our sample vary widely in their tenure. For 2009-10 we identified 29 instructors teaching in the core microeconomic doctoral courses in our sample. Of this number, eight earned their degree prior to 1980, eight earned their degree between 1980 and 1989, seven earned their degree between 1990 and 1999, and the remaining six earned their degree in year 2000 or later.

Job Market Placement

We collected data on students on the "job market" either in year 2012-13 or 2013-14. We used program websites listing job candidates during the year in question to identify students. In total we have 330 students in our sample. This is close to the universe of students on the market from these schools; we dropped only a handful of students for whom we could not find placement information. In nearly every case we were able to obtain a curriculum vitae for the student; in the few cases when this was not possible we used alternate sources to gather information about the student. In most cases the program lists job placement following the job market completion. In the cases where a student's outcome was not listed we reached out to their advisor or contacted them directly to obtain their placement information. 57 of the students placed at a full-time faculty position at one of the other universities hosting one of the other 7 programs and they are the sample for our job placement study.[14]

The students in our sample typically would have earned their degree in the year of job placement, thus approximately 3 or 4 years subsequent to 2009-10, our core micro data year. Thus we believe most would have taken the core at their program in this year or at worst in a neighboring year. To ensure this, we did not include in the job market sample any students who had completed their program more than one year prior to entering the job market (for example having taken a temporary position for more than one year).

[7] Berkeley course 201A had 3 texts in addition to class notes. The instructor stated that the class notes were the most primary source but also listed 3 texts as "recommended" and listed specific readings from these texts for each class session. The class notes are fairly technical and in our judgement it was important to include the readings as required as well: We belive a conscientious student would consult these readings.

[14] In the vast majority of cases a student placed at either the economics department or business school. In a few cases a student placed at a medical school or some other professional school. We included these as placements at the relevant university as long as it was a full-time faculty position. We did not include in the sample students who listed just a temporary position at another university in our sample of 8, and as noted at an earlier point did not include students who placed at their home university, since their may be many reasons why this occurs.

6. Results

In this section we present results from our empirical investigation of concepts taught at the top 8 graduate econ programs in their micro core courses, and the impact on job placement flows between these programs. In general, the results show that there is a lot of diversity in concept coverage, a natural clustering of schools, and that conceptual closeness does seem to impact job placement.

Appendix Table A2 provides a list of the higher levels of our concept map. It lists 17 general topics and under each general topic a set of topics, listing 108 topics in total. The bottom level of concepts is not shown and fits under the sub-topics, listing the individual concepts. Overall we coded 1370 concepts over all 8 programs in our study.

Table 2 presents a tabulation of the concepts covered in each program in our study. Column 2 lists the total number of concepts covered by each program and column 3 displays the fraction of concepts covered by the program relative to the universe of concepts covered over all programs.

All but three of the programs covers between 40 and 50 percent of the universe of concepts. The three exceptions are Harvard and Stanford, which cover somewhat over 50 percent of concepts, and Yale, which covers just under 40 percent.[12] We emphasize that we do not support a normative interpretation concerning the total number of concepts taught in a program. Not all concepts are equally difficult to teach, and teaching fewer concepts may enable a program to teach its set of concepts in more depth.

The lower panel of Table 2 shows analogous coverage statistics at the topic level. For this part of the table we code a program as "covering" a topic if any part of that topic is covered in assigned readings, thus erring on the side of coding programs as covering many topics. Not surprisingly, programs cover a higher percentage of broad topics but still significantly less than the universe of topics covered over all 8 programs. All but three cover between 70 and 80 percent of topics. Harvard and Stanford again have slightly higher coverage, and UC-Berkeley has slightly lower coverage.

An important takeaway from the numbers in Table 2 is the recognition that no one program covers close to the totality of concepts or even broad topics covered across all programs. Thus there is substantial diversity across programs, at the level of concept detail at which we are measuring content, as well as at the level of broader topics. This result can be viewed as presenting some evidence about product diversity in graduate education. In general in markets we expect product diversity and this market is no exception. Why this is the level of diversity we see for the year in question, its implications, and whether it is in any sense an "optimal" degree of diversity are important questions raised by our empirical findings.

Given that every program is one full year of microeconomics, and calculating based on 24 class sessions per semester, or 48 total, we estimate that a typical program covers approximately 11 to 14 concepts per class session. Of course these figures are not based directly on what is covered in class but on readings. Most likely not all concepts covered in assigned readings are actually covered in class so that the number of concepts covered in a typical class session is most likely below this.

[12] Stanford is on a quarter system. The first two courses, taught in the fall and winter, are responsbile for the larger number of concepts. The fall course in particular covers many topics. The syllabus for the course states that the lecture notes are the primary material, and states that the Mas-Colell, Whinston and Green has been ordered at the bookstore for the course. We do not have the lecture notes and code as required readings readings assigned from the Mas-Colell, Whinston and Green text. The volume of required reading is large. For example approximately two weeks are devoted to general equilibrium and chapters 15-17 of the Mas-Colell, Whinston and Green book are assigned for this part of the course. These chapters are quite heavy with concepts, especially Chapter 17.

Table 3, viewing the data from the perspective of concepts rather than programs, shows the number of concepts covered by one program, two programs, and so on. Interestingly, a signficant number of concepts - 390, or nearly one-third of the total, are covered by only a single program. In contrast, 153 concepts are covered by all programs, just 12 percent, and 150 concepts are covered by all but one program. At the level of broad topics, 37 percent of topics are covered by all programs; while this is substantially larger than the percentage of concepts covered by all programs, it is still fairly low. Appendix Table A3 breaks the information in table 3 down by program: It shows for each program the number of concepts covered by that program that are covered by no other program, by 1 other program, and so on. Every program covers at least some concepts not covered in any other program, with Chicago and Stanford covering the most concepts not covered by others. These two tables reinforce the finding of substantial diversity across programs in concept coverage.

Table 4 presents the correlation measure of overlap between pairs of programs. The correlations are positive but relatively modest in magnitude, running from below 20% to just over 50 %, revealing modest overlap in coverage of concepts between pairs of programs. The correlations reveal some clustering, the topic we explore next. For example, Harvard shares a relatively high correlation with MIT, Northwestern, and Stanford, while Northwestern shares a lower overlap with MIT, but high with Stanford.

Clustering Results

Figure 1 shows the clustering results for the method using Euclidean distance and the average distance method of computing cluster distances. The results show two main clusters: Harvard, MIT and Stanford in one cluster, and the remaining 5 programs in the other. The two tightest pairings are Northwestern and Yale, and Harvard and Stanford. In turn, Berkeley links with the Northwestern-Yale cluster and MIT links with the Harvard-Stanford cluster. Lastly, Chicago and then Penn are added to the Northwestern-Yale-Berkeley cluster. When we explore alternative methods for clustering, such as city-block, clusterings are mainly the same but not identical; the most variable program is Penn.

The two clusters suggest a natural interpretation of different schools of thought or philosophies in terms of teaching core micro. Both Northwestern and Yale have reputations as emphasizing technical rigor, with Penn also typically grouped in that category, though in our data Penn does teach somewhat different concepts. Harvard is known to emphasize learning microeconomic tools for practical policy application. Both MIT and Stanford might be thought of as “middle-of-the-road” programs emphasizing a wide-ranging approach to micro. They end up paired with Harvard as having an orientatin stressing application. This leaves Chicago, with its own distinctive heritage, and Berkeley. In the year we study part of the Chicago core micro curriculum was more applied, as taught by Gary Becker and Kevin Murphy, but focused on somewhat distinct concepts different than say Harvard. The other part was taught by more traditional economic theorists and was more technical. The combination places Chicago in the second cluster. Berkeley had a fairly technically oriented core micro program in the year we study, including a strong emphasis on general equilibrium analysis; for these reasons it ends up in the second cluster, closer to Penn, Northwestern and Yale.

Job Placement Results

The final part of our empirical analysis investigates job placements among the 8 programs.

Table 5 displays results for the two-stage model in which the first stage is graduate program placement and the second stage is job placement. We estimate this model for the two large clusters, with Harvard, MIT and Stanford forming cluster a and the remaining 5 programs cluster b. The first stage is estimated over all students in our sample, which contains all students who attended graduate school at one of the 8 programs and were listed on the job market in either years 2013-14 or 2014-15. The second stage is restricted to the subset of these students who accepted faculty positions at one of the other 7 programs. A key parameter is the correlation coefficient ρ, which links the two stages statistically. The results in the first stage are for a model in which ρ is estimated, while the results in the second stage are for a model in which ρ is constrained to be zero. Comparison of the log likelihoods between the two models allows construction of a likelihood test for testing the hypothesis that ρ is in fact zero. As shown, the log likelihoods have only a tiny difference, so that based on the log likelihood test we cannot reject the hypothesis that rho is zero. Thus we conclude that there is no evidence in support of a link between the first stage program placement and second stage job placement parts of the model in terms of the unobservable components u and v. This enables us to present additional, simpler results just for the job placement part of the model, which we discuss below. First however we discuss the results for the remaining parameters for the two-stage model.

For stage 1, program placement, we explored a range of variables, gleaned from students' curriculum vitae. The results we display include 3 explanatory variables. Foreign is a dummy variable coded as 1 if the student is not from the U.S. As shown, this coefficient is statistically significant and the negative sign indicates that being from a foreign country means the students is more likely to place in a program in cluster b, not cluster a. In additional work not reported we examined specific countries and regions of origin and in general found only small effects for individual countries and regions. The second variables is Grad which is a dummy variable coded as 1 if the student already has a graduate degree, typically a master's degree, upon applying to doctoral economics programs. This variable is not statistically significant. Finally, the third variable is a dummy variable Math, coded as 1 if a student has a math or math-related major in college such as engineering. This variable has a relatively large positive magnitude, suggesting that students with a more math intensive background are more likely to place at a program in cluster a. However the variable is not statistically significant.

Our main focus is stage 2, job placement. Recall this model is estimated only over a relatively small sample of 57 students who placed at one of the other 7 programs in our sample. Critically, in the model with ρ constrained to zero we find a positive, statistically significant coefficient for our variable Close, which is a dummy variable that takes the value 1 if the program under consideration for job placement is in the same cluster, hence conceptually close. Thus we do find evidence that the conceptual closeness of programs has an impact on job placement. The coefficient on Close is still positive, but of smaller magnitude and not statistically significant when we allow ρ to be freely estimated. We also include two important control variables in this model. Rank is a dummy variable that takes the value 1 if the program under consideration has a similar ranking to the student's home program, in the specification shown here meaning the two programs are either both above the median rank, or both below. We might have expected that this might be an important drive of job placement, but in fact the coefficient on this variable is relatively small and not statistically significant. Finally, the variable Metro is a dummy variable that takes the value 1 if the program under consideration is in the same metro region as the student's home program. Again, we might have expected this to be an important driver of job placement. In fact the coefficient on this variable is positive, as expected, but is not statistically significant.

Table 6 shows a set of results for the simpler models in which we estimate only the job placement outcomes and omit the first stage, which from the previous analysis does not have a statistically significant interaction with job placement. The top panel shows results for which conceptual closeness between programs is based on the clustering analysis, while the bottom panel shows results for which conceptual closeness is based on the correlation coefficient of concept coverage between pairs of programs. [Note that this model is not identical to the model in column 2 of Table 5, because even though ρ is constrained to zero in that model there is still a random effect for the second stage, whereas these models are more straightforward multinomial logit specications.] We include Rank and Metro as control variables in both specifications.

In both specifications we find a positive effect of conceptual closeness of a target program to the candidate's home program. The effect is statistically significant for the cluster-based analysis, and not statistically significant, though of relatively large magnitude, for the correlation-based analysis. Overall, the results provide support for the hypothesis that conceptual closeness of programs is linked with job placement flows. Of course further work with more data would be very valuable to corroborate these findings, but they certainly suggest that individuals trained in more similar core micro concepts are a more natural fit as faculty with a program.

6. Conclusion: Interpretation and Implications

In this paper we have presented empirical findings for a study of the concepts taught in core microeconomics doctoral courses at a set of top US economics programs in 2009-10. Our main finding is that there is substantial diversity across programs in the concepts they cover based on required readings. In addition we find that programs are statistically more likely to hire as new assistant professors individuals coming from programs more similar to their own, as identified by a cluster analysis, even after controlling for other factors.

Our results point in several directions of additional analysis. One is developing a better understanding of what drives the choice of topics to cover in core courses. This topic can be explored empirically and theoretically. Empirically, it would be interesting to understand to what degree topic and concept coverage are linked to individual faculty who teach in the core. A faculty may bring in new material when he or she begins teaching in the core, or may evolve distinct topics over the years teaching a core course. A second, related issue is the extent to which what is taught transcends individual faculty and is better understood as determined at the program level. It may be possible to explore department level changes, for example in chairmen, retirements, and new senior hires, and explore whether this has an impact on core coverage.

A second open topic for further research is constructing theoretical models of educational curriculum. We are not aware of prior work specifically studying, from an economics and knowledge-based approach, educational programmatic decision-making about what concepts to teach in a curriculum. This type of modeling can not only inform our understanding of individual program choices but also, at the level of fields such as economics, produce predictions about equilibrium among programs in the market for degrees in particular fields. To model the decisions programs make about their curriculum requires specifying the objective function that guides program decision-making. Presumably programs are guided by the desire to attract students, place students (which in turn attracts further students), and attain high rankings. Attaining a high ranking is itself a complex endeavor, involving research productivity, student placement both in academic and non-academic jobs, student satistfaction, degree completion rates, as well as perhaps a more intangible sense of program identity which may in fact have some persistence over time. In turn, these objectives drive choices about faculty hires and what material to teach in core programs, as well as a host of other decisions around the student experience and faculty management. Finally, the model must specify the nature of competition in fields of education. Our results indicate that a valid model will likely generate predictions of significant diversity across programs as one implication.

A separate area of analysis is exploring further empirical linkages, especially over time. How stable are clusterings over time? What about placements outside academic, how do these relate to curriculum design?

This paper is just a first step in what may be a much larger endeavor to understand education curriculum choices in a more systematic fashion.

Appendix I: Additional Formulas

The total number of concepts covered in program j is then $n_j = \sum_{i \in \mathcal{A}} X_{ij}$. The fraction of concepts covered in program j is

For each concept i we define $a_i = \sum_{j \ Programs} X_{ij}$ to be the number of programs that cover i.

Our first measure of overlap in concepts between program pairs is the ratio of the number of concepts both schools teach to the universe of concepts both schools teach, defined as the intersection of their concept sets divided by the union of their sets. Consider two programs j and k. We denote this measure of overlap of concepts taught between programs j and k by v_{ik}. Formally:

$$v_{jk} = v_{kj} = \frac{\sum_{i \in \mathcal{A}} X_{ij} X_{ik}}{n_j + n_k - \sum_{i \in \mathcal{A}} X_{ij} X_{ik}}.$$

To give a sense for this measure, suppose for example that for each program 1/2 of its concepts are taught at the other program and 1/2 are not. The measure of overlap will then equal 1/3. If 2/3 of the concepts taught in each program are taught in the other program, the measure is 1/2. If 3/4 of the concepts taught in each program are taught in the other, the measure is 0.6. It will typically not be the case that each program has the same fraction of concepts taught in the other program. For example, suppose one program has 1/2 of its concepts taught in the other program, and the second program has 4/5 of its concepts taught in the first program. In this case the first program is teaching more concepts (in the ratio 8 to 5) and the measure is 4/9.

Our second measure of concept overlap is an asymmetric measure: for a pair of programs j and k, we define s_{jk} to be the ratio of the number of concepts that are taught in j that are also covered in program k compared to the total number of concepts taught in j. Formally:

$$s_{jk} = \frac{\sum_{i \in \mathcal{A}} X_{ij} X_{ik}}{n_j}.$$

This is a measure of the extent to which program k covers the concepts taught in program j. In the case in which each program has 1/2 of its concepts taught in the other program, $s_{jk} = s_{kj} = 1/2$. Note that this is higher than the 1/3 for the first measure of overlap. In the case in which 2/3 of the concepts taught in each program are taught in the other program, $s_{jk} = s_{kj} = 2/3$. Finally, if program j has 1/2 of its concepts taught in program k, while 4/5 of the concepts taught in program k are taught in program j, then $s_{jk} = 1/2$ and $s_{kj} = 4/5$. This measure is especially valuable in highlighting asymmetries between programs, such as in this last example.

Broad topic Clustering (if want to use): The reason clustering can shift when we move from concepts to broad topics is that two schools may each cover a broad topic to the same degree of depth, but cover different specific concepts within that topic. As a result at the concept level they will appear different, but at the more aggregage broad topic level they will appear more similar. As an example, consider the case in which there are two broad topics each containing two concepts, thus 1a, 1b, 2a and 2b, and two schools, X and Y. In one scenario X covers 1a and 1b and Y covers 2a and 2b. In this case the schools do not overlap at either the concept or topic level. In a different scenario X covers 1a and 2a while Y covers 1b and 2b. In this case the two schools do not overlap at the concept level, thus will not cluster together, but each covers one-half of topic 1 and one-half of topic 2, thus they are measured as identical at the broad topic level and will cluster together.

References

Abito, J.M., K. Borovickova, H. Golden, J. Goldin, M.A. Masten, et. al. (2011): "How Should the Graduate Economics Core be Changed?" *Journal of Economic Education*, 42, 4, pp. 414-7.

Agrawal, R., J. Gehrke, D. Gunopulos, and P. Raghavan (2005): "Automatic Subspace Clustering of High Dimensional Data," *Data Mining and Knowledge Discovery*, 11, pp. 5-33.

Barone, C.A. (1991): "Contending perspectives: Curricular reform in economics," *Journal of Economic Education*, 22, 1, pp. 15-26.

Berrett, D. (2012): "3 colleges' different approaches shape learning in econ 101," *Chronicle of Higher Education*, 58, 39, pp. A6-9.

Bowen, H.R. (1953): "Graduate education in economics," *American Economic Review*, 43, 4, Pt. 2, Supplement, pp. ii-xv, pp. 1-223.

Colander, D. (2005): "The making of an economist redux" *Journal of Economic Pespectives*, 19, 1, pp. 175-98.

Colander, D. and A. Klamer (1987): "The making of an economist," *Journal of Economic Perspectives*, 1, 2, pp. 95-111.

Feinstein, J.S. (2006): *The nature of creative development*. Stanford, CA: Stanford University Press.

Feinstein, J.S. (2011): "Optimal learning patterns for creativity generation in a field," *American Economic Review Papers and Proceedings*, 101, 3, pp. 227-32.

Krueger, A.O., K.J. Arrow, O.J. Blanchard, A.S. Blinder, C. Goldin, E.E. Leamer, R. Lucas, J. Panzar, R.J. Penner, T.P. Schultz, J.E. Stiglitz, L.J. Summers (1991): "Report of the Commission on Graduate Education in Economics," *Journal of Economic Literature*, 29, 3, pp. 1035-53.

Mokyr, J. (2002): *The gifts of Athena: historical origins of the knowledge economy.* Princeton: Princeton University Press.

McDonald, J.F. (2009): "Graduate education in economics: Microeconomics at Chicago and Yale in the 1960s," *Journal of the History of Economics Thought*, 31, 2, pp. 161-80.

Report of the Commission on Graduate Education in Economics (1991): Madison, WI.

Table 1: Doctoral Economics Programs Included in the Study

University of California at Berkeley
University of Chicago
Harvard University
Massachusetts Institute of Technology
Northwestern University
Stanford University
University of Pennsylvania
Yale University

Table 2: Number of Concepts Covered in Required Readings in Each Program Included in the Study

Program	Number of Concepts Covered	Fraction of Total Concepts Covered out of Universe of All Concepts Covered Across All Programs in the Study
University of California at Berkeley	585	.43
University of Chicago	635	.47
Harvard University	698	.52
M.I.T.	572	.42
Northwestern University	602	.44
Stanford University	789	.58
University of Pennsylvania	559	.41
Yale University	532	.39

Program	Number of Broad Topics Covered	Fraction of Broad Topics covered out of All Broad Topics Covered Across All Programs in the Study
University of California at Berkeley	76	.68
University of Chicago	88	.79
Harvard University	91	.81
M.I.T.	81	.72
Northwestern University	88	.79
Stanford University	92	.82
University of Pennsylvania	83	.74
Yale University	81	.72

Table 3: Number of Programs that Cover a Concept

Number of Programs	Number of Concepts Covered by This Many Programs
1	438
2	186
3	140
4	75
5	116
6	91
7	156
8	153

Table 4: Correlations Between Pairs of Programs

	Berkeley	Chicago	Harvard	MIT	Northwestern	Stanford	Penn	Yale
Berkeley	1.00	0.33	0.38	0.13	0.41	0.18	0.27	0.41
Chicago	0.33	1.00	0.21	0.25	0.41	0.24	0.24	0.37
Harvard	0.38	0.21	1.00	0.43	0.47	0.44	0.29	0.36
MIT	0.13	0.25	0.43	1.00	0.39	0.37	0.31	0.34
Northwestern	0.41	0.41	0.47	0.39	1.00	0.52	0.47	0.5
Stanford	0.18	0.24	0.44	0.37	0.52	1.00	0.29	0.33
Penn	0.27	0.24	0.29	0.31	0.47	0.29	1.00	0.25
Yale	0.41	0.37	0.36	0.34	0.5	0.33	0.25	1.00

Table 5: Program Choice and Job Placement: Two Stage Model

	Stage 1: Program Choice	
	Coefficient estimates	
Constant	0.3761841	0.3642051
	(0.2643883)	(0.2556848)
Foreign	-1.1714958**	-1.1393117**
	(0.3344244)	(0.351953)
Grad	0.1665338	0.1494024
	(0.3266334)	(0.348846)
Math	0.3765918	0.3829121
	(0.2997247)	(0.2976788)
	Stage 2: Job Placement	
	Coefficient estimates	
Close	0.4572653	0.7869077*
	(0.3700959)	(0.3761295)
Cluster A	-0.2334381	-0.1475163
	(0.3699753)	(0.3598136)
Rank	-0.1857178	-0.1776920
	(0.3358162)	(0.4644568)
Metro	0.4533122	0.4338996
	(0.4307894)	(0.2661456)
ρ	0.95356776	0.0
	(1.0714244)	(fixed)
Log likelihood	-316.1836	-316.5732
N	330	330

**p <0.05, ** p<0.01*

Standard error for rho in column 1 estimated from likelihood ratio test for hypothesis: rho = 0.0

Table 6: Job Placement: Second Stage Alone

Panel A: Cluster model

	Coefficient estimates
Close	0.7086384*
	(0.3249231)
Cluster A	-0.0693788
	(0.3237461)
Rank	-0.2089697
	(0.3314052)
Metro	0.4705561
	(0.426726)
Log likelihood	-106.2115
N	57

Panel B: Correlation model

	Coefficient estimates
Correlation	2.6589182
	(1.714999)
Rank	-0.1184447
	(0.3300212)
Metro	0.6175498
	(0.393997)
Log likelihood	-107.4875
N	57

** p<0.05*

Figure 1: Cluster Map
Euclidean distance, average linkage, 2009-2010

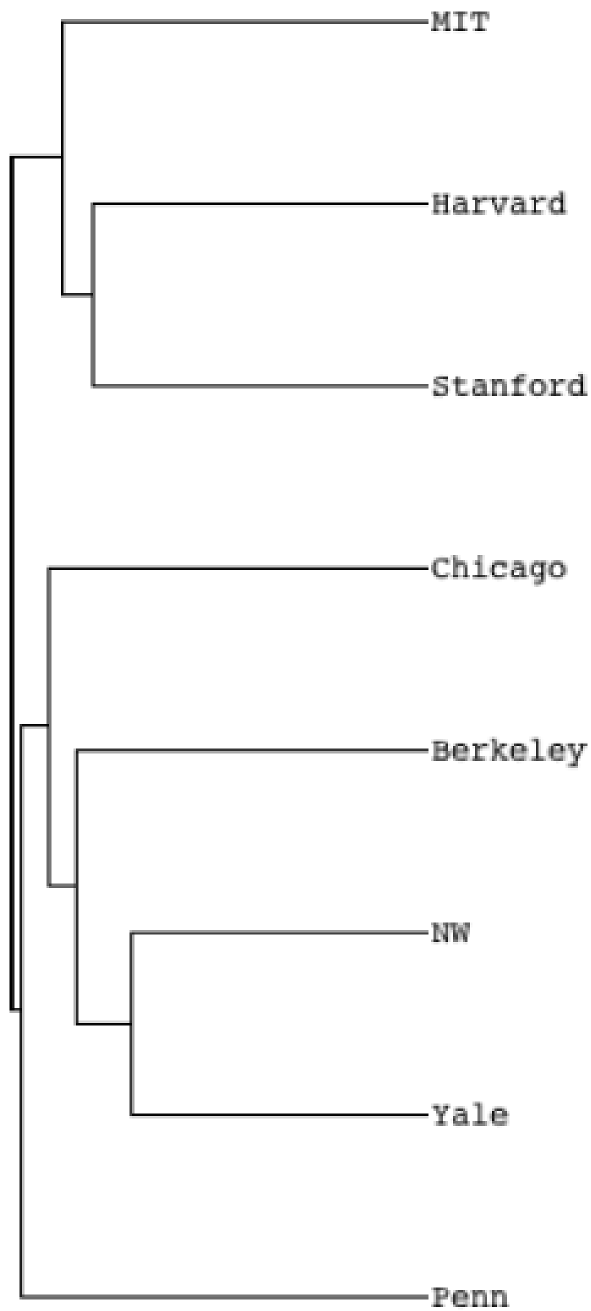

Table A1: Widely Used Textbooks in the Programs Included in the Study

Title	Author	Publication Year	Number of Assigned Concepts Linked to Text
Microeconomic Theory	Andreu Mas-Colell, Michael D. Whinston, Jerry R. Green	1995	865
Game Theory	Drew Fudenberg, Jean Tirole	1991	273
Advanced Microeconomic Theory	Geoffrey A. Jehle, Philip J. Reny	2001	255
A Course in Game Theory	Martin J. Osborne, Ariel Rubinstein	1994	171
Notes on the Theory of Choice	David M. Kreps	1988	95
General Equilibrium, Overlapping Generations Models, and Optimal Growth Theory	Truman F. Bewley	2007	95
Lecture Notes in Microeconomic Theory	Ariel Rubinstein	2006	145
A Course in Microeconomic Theory	David M. Kreps	1990	355
Microeconomic Analysis	Hal R. Varian	1992*	104
Contract Theory	Patrick Bolton, Mathias Dewatripont	2005	81

Table A2: General Topics and Sub-Topics

General Topic	Sub-Topics
Preferences	Preference orderings & properties; revealed preference.
	Representation of preferences by Utility function; properties of u functions.
	Kinds of preferences and utility functions; behavioral issues.
	Lancaster model, quality and attributes, measurement & empirics.
Demand	Budget Sets
	Demand functions/correspondences &properties
	Comparative statics; Slutsky equation, compensated price changes;
	Duality: expenditure function, Hicksian demand, duality theorems.
	Integrability; recovering preferences from data on demand.
	Empirical issues & estimation of demand.
	Strong Axiom of revealed preference; GARP
	EV and CV; welfare; Index numbers
	Aggregation - aggregate demand
Choice Under Uncertainty	Basic set-up, key assumptions, statement of EU theorem.
	Proof of EU theorem, assuming Best/Worst outcomes
	Extension to unbounded cases; mixture spaces.
	Anomalies
	State dependent U and Anscombe-Aumann model
	Subjective probability, sure thing principle, Savage model to recover preferences.
	U over money, definition of risk attitudes, certainty equivalents, risk premium.
	Comparisons of lotteries; measures of risk aversion.
	Dominance.
	Applications - insurance, mean-variance formulation, asset allocation.

Production	Production sets, production functions, properties.
	Profit functions, profit maximization, derivation of factor input demands.
	Cost functions - SR and LR, economies of scale and scope; adjustments.
	Mathematical properties of the cost function.
	Efficiency of production; nature of the firm and its objectives.
	Linear activity model.
	Aggregate Supply
Markets, Outcomes, Environment	Pareto Optimality
	Single Market equilibrium - conditions.
	Comparative statics; robust comparative statics.
	Free Entry and LR equilibrium
	Externalities & ways to deal with.
	Public goods - definition, efficiency conditions, voting models.
General Equilibrium	Exchange economy, Edgeworth Box.
	2X2 production model and results; labor supply.
	First and Second Welfare theorems; Pareto optimality.
	Existence of equilibrium: excess demand, fixed point theorems.
	More advanced topics with excess demand: Local uniqueness and Index theorem; Sonnenschein-Mantel-Debreau theorem; uniqueness.
	Tatonnement process and dynamics around equilibrium; comparative statics.
	Replication, large economies, the core, noncooperative approaches, convergence to perfect comp. equilibrium.
	GE under uncertainty: contingent commodities; rational expectations equilibrium.
	GE under uncertainty: Assets, spanning, arbitrage.
	Incomplete markets; asymmetric information.
	Equilibrium over time: paths.
	OLG models.

Static Games of Complete Information	Strategic Form Games
	Mixed Strategy
	Elimination of Dominated Strategies
	Nash Equilibrium
	Iterated Strict Dominance
	Rationalizability
	Correlated Equilibrium
	Supermodular games
	Generic Properties of Nash Equilibria
	Transforming Incomplete Information game to Imperfect Information Game
	Strategy
Static and Dynamic Games of Incomplete Information	Player type, Belief, Strategy with Incomplete Information
	Incomplete Information
	Bayesian Equilibrium
Dynamic Games of Complete Information	Extensive Form, Game Tree, Information Set
	Backward Induction
	Subgame Perfect Equilibrium
	Rubinstein Stahl Bargaining Model
	Timing Games
	Open Loop Equilibrium, Closed Loop Equilibrium, Iterated Conditional Dominance
	Repeated Games with Observable Actions
	Repeated Games with Imperfect Public Information, Perfect Public Equilibrium
Mechanism Design	General Setting
	Revelation principle and Direct Revelation Mechanism
	Groves Mechanism
	AGV Mechanism
	Revenue equivalence in auctions
	Implementation in environments with complete information
	Efficiency of Mechanisms

Dynamic Games of Incomplete Information	Setting
	Signaling Game
	Cheap Talk
	Perfect Bayesian Equilibrium
	Game trees for games with incomplete information
	Sequential Equilibrium
	Strategic Form Refinements
	Reputation Effects
	Strategic Stability
	Robustness under payoff uncertainty
Social Choice	Social Welfare Functional for 2 and general number of alternatives
	Social Choice Function
	Utility Possibility Set
	Invariance properties of social welfare functions
Bargaining	Axiomatic Bargaining
Monopoly	Monopoly problem and Natural Monopoly
	Price Discrimination
Oligopoly	Cournot and Stackelberg Equilibrium
	Collusion
	Bertrand Competition
	Contestable Markets
	Entry Deterrence and Accommodation
	Differentiated Product Competition
Coalition Games	Coalition game with transferable utility
	Coalition game without transferable payoffs
Information economics/Dynamic games of incomplete information	Principal-Agent Moral Hazard
	Principal-Agent Hidden Information
	Moral Hazard in teams
	Incomplete contract - Hold up problem

	Theory of the firm
	Adverse Selection

Table A3: Number of Concepts Taught by How Many Other Programs

	None	1	2	3	4	5	6	7
University of California at Berkeley	58	67	38	23	57	54	135	153
University of Chicago	102	58	28	21	49	69	155	153
Harvard University	67	41	77	37	84	83	156	153
M.I.T.	20	50	78	44	65	46	116	153
Northwestern University	18	18	43	46	93	75	156	153
Stanford University	106	58	73	57	96	90	156	153
University of Pennsylvania	37	38	57	37	87	59	91	153
Yale University	30	42	26	35	49	70	127	153