# THE MAPINATOR CLASSIFICATION: CROSS TIER PLACEMENT RATES

MICHAEL PETERS, JAMES YU (UBC)

Abstract. The paper uses public data from the mapinator project to classify Universities based on the way they place their economics phd graduates. Universities that have similar cross tier placement rates are classified together. The classification resembles many well known academic rankings, but is based on a very different methodology which is more useful for predicting graduate outcomes. The information in the classification can be used to guide graduate school choices. It is also useful in guessing the outside options available to graduates during recruiting.

There are many rankings of economics departments. Some are based on publications (Tilberg Ranking), working papers (RePec), or surveys (US News and World Report). The results here don't have much to add to these 'rankings', in fact we don't strictly speaking have a ranking here.

What this paper is about is complementing those rankings by adding a different kind of information. Our methodolgy does not use publications, or surveys, only information about graduate placements. The only paper we have found with a similar approach is Amir and Knauff (2008). They scraped websites of top schools then traced their faculty back to the departments where they graduated. This created a network that tracked trades across a subset of important universities. They didn't classify universities the way we do here, instead they ranked them using google page rank.

Our data includes schools other than top ranked ones, and tracks placements outside academics. In this sense we know what happened to graduates *who didn't* get jobs in top schools. In other words, we have placements that Amir and Knauff (2008) never monitored.

In any case, our point is different from theirs. Most economists broadly agree on what the top economics graduate schools are. Our objective is to identify how the top schools differ from other schools, and to provide similar information about schools outside the top tier.

For example, one might ask how likely it is that a graduate from a top school will later get a job at a top school. More important for potential graduates, how do these chances differ if they go to a prominent but less prestigious school.

We provide a way to identify 'market failures'. By this, we mean placements graduates who were not successful at getting offers from top universities and firms. Since Ph'd graduates in economics are too valuable to remain unemployed, we identify graduates whose post graduation jobs probably could have been acquired without their participation in the international job market. It is by no means

*Date*: July 16, 2024.

The authors are grateful for comments and advice from Li Hao, Wei Li, Sergei Serverinov, Vitor Farinha Luz, Sam Gyetvay, Jesse Perla, Vadim Marmer, Kevin Song and Hiro Kasahara.

obvious how to do this, but we suggest an approach here that is at least a workable start.

More generally, our classification provides a relatively intutitive way to think about complex network trading data. In a related paper (Mapinator Classification: Theory) we provide a theoretical formulation that can be used to estimate the values of the graduates of the different tiers and make inferences about market efficiency and the benefits of providing inforormation to the market (like the information in this paper).

The data we use comes from a public database of outcomes created by the Mapinator Project at the University of British Columbia. This project is sponsored by https://econjobmarket.org. Much of the work on the project has been completed by honours students in economics at the Vancouver School of economics. One example is the following interactive program - interactive program - that can be used to explore the database. This was written by Amedeus D'Souza (at the time of writing, a phd student at Chicago) who was an honours student at the Vancouver School of Economics at the time.

**The Adjacency matrix.** The easiest way to understand the data is to look at the visualization at https://sage.microeconomics.ca. It shows a directed graph that illustrates the transaction flows of economics ph'd graduates between universities from 2003 to the present - essentially a visualization of the worldwide market. As with all network data, it is hard to interpret visually.

To see how this paper will attempt to understand the data, it is useful to try to visualize it using an *adjacency matrix.* The data from the directed graph can be converted into a *raw* adjacency matrix by creating a matrix which has one row for every market participant on the demand side of the market (every university, every private sector firm and every government adjacency who hires phd graduates in economics) along with one column for every university that produces graduates. Any cell $(i, j)$ of this matrix contains the number of graduates from the university represented by cell $j$ who were hired by the institution represented by row $i$.

In the data used here, there are 1114 universities that both produce graduates themselves, and also hire graduates from other universities. There are 642 universities that hire graduates at the assistant professor level, but don't themselves produce any graduates. We'll refer to them here as *teaching universities.* There are 227 private sector institutions and 152 government institutions as well. The other three rows in the matrix track placements at universities which are not at the assistant professor level.

The number of columns in the raw adjacency matrix would be just 1114, the number of universities that graduate students who participate on this international job market. There are 2777 institutions who try to hire phd graduates.

Let $A$ represent this raw adjacency matrix, with representative entry $a_{ij}$ interpreted as above, the number of hires by $i$ from institution $j$. The basic objective in the paper is to sort the various institutions into communities based on the idea that members of the same community will have similar placements essentially because the products they produce are close substitutes.

A *model* is a finite collection of communities $\mathcal{C}$ (including a specification of the members of each community). The paper attempts to identify the best model using maximum likelihood.

Once we have chosen a model with a small number of communities, we can create something we'll just refer to as the *adjacency matrix* without the 'raw' qualifier. Each cell $a_{ij}$ in this matrix records the number of graduates from community $j$ that were hired by some member of community $i$. Of course, this matrix has the advantage that it makes the raw data easier to read. However, it also makes it possible to look for patterns in the data that can reveal some of the economics that underly it. For example, we might ask what a tier based classification would look like, and whether we could distinguish it from a geographic classification. For the most part, we'll defer these questions to another paper and stick with the methodology that was used to create the classification.

This is the adjacency matrix associated with our best model. Looking down the column for Tier 4, for example, you can see that there were 43 graduates of Tier 4 universities from this sample of the data who were hired by Tier 1 universities. Tier 1 universities hired mostly from other Tier 1 universities.

| | **Tier 1** | **Tier 2** | **Tier 3** | **Tier 4** | **Tier 5** | **Row Totals** |
|---:|---:|---:|---:|---:|---:|---:|
| TYPE 1 (20 insts) | 1194 | 338 | 172 | 43 | 4 | 1751 |
| TYPE 2 (58 insts) | 950 | 853 | 314 | 91 | 10 | 2218 |
| TYPE 3 (180 insts) | 1073 | 1381 | 986 | 146 | 34 | 3620 |
| TYPE 4 (334 insts) | 270 | 506 | 451 | 409 | 33 | 1669 |
| TYPE 5 (522 insts) | 0 | 42 | 53 | 28 | 163 | 286 |
| Public Sector (152 insts) | 568 | 545 | 285 | 57 | 23 | 1478 |
| Private Sector (227 insts) | 777 | 441 | 230 | 55 | 26 | 1529 |
| Postdocs (598 insts) | 86 | 130 | 135 | 56 | 27 | 434 |
| Lecturers (413 insts) | 34 | 59 | 69 | 41 | 48 | 251 |
| Unmatched (1 insts) | 367 | 599 | 752 | 413 | 341 | 2472 |
| Other Groups (38 insts) | 210 | 165 | 83 | 32 | 9 | 499 |
| Teaching Universities (642 insts) | 240 | 377 | 420 | 159 | 69 | 1265 |
| Column Totals | 5769 | 5436 | 3950 | 1530 | 787 | 17472 |

There is a row in the matrix called 'Unmatched'. This consists of graduates who either disappeared in the sense that no record of their employment could be found by web search, or graduates who got jobs in institutions which have never been recorded at econjobmarket.org. The institutions which haven't been identified by econjobmarket have never advertised on econjobmarket and have never graduated a student who registered on econjobmarket.

Of course, this is an arbitrary definition of the international job market. Examples suggest that many of the graduates in this category end up in jobs that don't necessarily need a ph'd in economics. For example, a credit analyst in a small bank, an administrator, a lecturer in a local college. Others disappear into countries that don't use English, making it hard to track them. Such countries often have local markets that work well enough on their own.

Since all the grad students involved in these placements did take the time to register on econjobmarket, they probably wanted offers from the international job market, but didn't get them.

Tiers 1 and 2 consist of universities most would be able to predict. Those two tiers along with the private and public sector hire most of their graduates from tieir 1. The other tiers hire most of their applicants from the higher tiers. The reverse seems to be true for placements. This makes the adjacency matrix look heavy below the diagonal.

This suggests a value based tier system. This also appears to differ from a geographic clustering which would tend to have nearest neighbour clusters around the diagonal.

Secondly notice that only 20% of graduates from tier 1 get jobs in a tier 1 school. more than 40% of graduates from tier 1 schools don't get academic jobs at all. At the other extreme, only 4 tier 5 graduates get jobs at one of the top tier academic schools. This pattern can be expained using a value based tier system along with trading frictions. However, this data does suggest that even though the school where you graduate will have an important impact on where you end up, it is not by itself a very good predictor.

Finally, at first glance, it might appear that the tier 1 produces almost as many graduates as all the others combined. This is not true because there is a large bias in sampling toward higher ranked schools that tend to produce web based lists of their students placements.

The following table illustrates the biased sampling:

| | **Type 1** | **Type 2** | **Type 3** | **Type 4** | **Type 5** |
|---|---|---|---|---|---|
| Registrations - 2018-23 | 2664 | 3632 | 4472 | 2282 | 1339 |
| Registrants without placements | 431 | 1016 | 2113 | 1388 | 863 |
| Adjustment | 1.19301 | 1.38838 | 1.89572 | 2.55257 | 2.81303 |

The first row gives the number of students who register with econjobmarket in each of the last 5 years sorted by tier. The second row indicates registrants for which placements have not yet been found (in the sense that no one has yet searched for them). There are many reasons for this sampling bias. For example, top tier universities mostly like to advertise their placements (oddly there are a number of them who insist on claiming a large number of placements without revealing any information about them). The top 100 universities have pages on the search site showing which of their graduates have no recorded placements, so it is just faster to search from the top tiers.

In any case, some of the appearance in the adjacency table comes from this skewed search method. Assuming the applicants who haven't been found have the same kinds of outcomes as the ones who have makes it possible to guess what the actual adjacency matrix looks like. This is it:

| | **Tier 1** | **Tier 2** | **Tier 3** | **Tier 4** | **Tier 5** | **Row Totals** |
|---|---|---|---|---|---|---|
| TYPE 1 (20 insts) | 1424 | 469 | 326 | 110 | 11 | 2340 |
| TYPE 2 (58 insts) | 1133 | 1184 | 595 | 232 | 28 | 3172 |
| TYPE 3 (180 insts) | 1280 | 1917 | 1869 | 373 | 96 | 5535 |
| TYPE 4 (334 insts) | 322 | 703 | 855 | 1044 | 93 | 3017 |
| TYPE 5 (522 insts) | 0 | 58 | 100 | 71 | 459 | 688 |
| Public Sector (152 insts) | 678 | 757 | 540 | 145 | 65 | 2185 |
| Private Sector (227 insts) | 927 | 612 | 436 | 140 | 73 | 2188 |
| Postdocs (598 insts) | 103 | 180 | 256 | 143 | 76 | 758 |
| Lecturers (413 insts) | 41 | 82 | 131 | 105 | 135 | 494 |
| Unmatched (1 insts) | 438 | 832 | 1426 | 1054 | 959 | 4709 |
| Other Groups (38 insts) | 251 | 229 | 157 | 82 | 25 | 744 |
| Teaching Universities (642 insts) | 286 | 523 | 796 | 406 | 194 | 2205 |
| Column Totals | 6883 | 7546 | 7487 | 3905 | 2214 | 28035 |

This table suggests that the largest producers of graduates are tiers 2 and 3. However, using the number of institutions in the classification, tier 1 produce by far the largest average number of placements per institution.

**Estimation.** The next step in the process is to use a method of find the appropriate community model. We use a variants of the *stochastic block model* to do this, following Karrer and Newman (2011); Peixoto (2014); Wang and Bickel (2017).

In most of the literature the approach is to take the entire network graph and create from it a square adjacency matrix which has as many rows and columns as there are nodes in the graph. In other words there is a distinct row and column for every institution that either hires or graduates phd students in economics. This isn't what we want here because we already know the indentities of the institutions that graduate phd students (and the ones who don't). As some institutions only hire, large blocks of this adjacency matrix would be full of zeros.

Perhaps more important, we already know pretty well how to allocate the non-academic institutions in the placement data. We expect all the government institutions and teaching universities, for example, to hire in pretty much the same way as other institutions in their communities. For this reason, we'll work with an adjacency matrix which has one row for every hiring institution, but only has a column for every academic university that graduates students.

We'll refer to this adjacency matrix as the *raw* adjacency matrix and continue to refer to it as $A$, with enteries $a_{ij}$ giving the number of graduates from *academic* institution $j$ that were hired by institution $i$.

Since we will assign the institutions who don't produce graduates to communities manually, finding the best community model involves choosing the number of different academic communities, then assigning each university to one of the communities.

Once we have an assignment of institutions to communities, then we can create a new matrix which we just refer to as the adjacency matrix (that is, without the 'raw' qualifier. Each element $i, j$ of this matrix will contain the number of graduates from academic community $j$ that were hired by community $i$. The dimension of this matrix is $K \times k$ where $k$ is the number of academic communities and $K$ is just the sum of $k$ and the number of communities we have predefined to contain the institutions who don't graduate students. Table in the introduction is a compressed adjacency matrix with 5 academic communities. Since the number of non-academic communities is fixed, we can refer to a community $\mathcal{C}^k$ consisting of $k$ academic tiers. We don't know which institutions belong to each tier, we have to estimate that.

One important assumption in what we do is to assume that institutions who are in the same academic community are also in the same hiring community. From some processing we'll do below, this appears to be largely true though there are some exceptions.

Define as estimate $\hat{\mathcal{C}}^k$ as a particular assignment of universities into the $k$ communities. The cell $c, c'$ of this compressed matrix $A[\hat{\mathcal{C}}^k]$ contains the value

$$\sum_{i \in c} \sum_{j \in c'} a_{ij}.$$

In other words, the cell $(c, c')$ records the number of graduates from academic community $c'$ who were hired by hiring community $c$.

We start with the assumption that the random variable $a_{ic'} = \sum_{j \in c'} a_{ij}$ has a poisson distribution with parameter $\lambda_{ic'}$ which is common to each institution $i$ in community $c$. Write for the moment, $\lambda_{ic'} \equiv \overline{\lambda}$. Since the sum of $n_c$ independent poisson variables with common parameter $\overline{\lambda}$ is a poisson random variable with parameter $n_c\overline{\lambda}$, the random variable

$$\sum_{i \in c} \hat{a}_{ic'} = \sum_{i \in c} \sum_{j \in c'} a_{ij}$$

can be written as a poisson random variable with parameter $n_c\overline{\lambda} \equiv \lambda_{cc'}$.

Then the distribution of $\sum_{i \in c} \hat{a}_{ic'} = \hat{a}_{cc'}$ is poisson with parameter $\lambda_{cc'}$.

Assuming the random variables in the cells of the adjacency matrix are independent, then we can think of the adjacency matrix itself as a random variable. The probability of the adjacency matrix $\hat{A} = \{\hat{a}_{cc'}\}$ is

$$\prod_{c} \prod_{c'} \frac{e^{\lambda_{cc'}} \lambda_{cc'}^{\hat{a}_{cc'}}}{\hat{a}_{cc'}!} \tag{0.1}$$

Once we assemble a potential adjacency matrix $\hat{C}$, then we can find the parameters $\left\{\hat{\lambda}_{c.c'}\right\} \equiv \hat{\lambda}(C)$ that maximize (0.1) and calculate the likelihood of $C$ given those parameters.

Hypothetically we want a community structure that maximizes

$$G(C) = \prod_{c \in \mathcal{C}} \prod_{c' \in \mathcal{C}} \frac{e^{\hat{\lambda}_{cc'}} \hat{\lambda}^{\hat{a}_{cc'}}}{\hat{a}_{cc'}!}, \tag{0.2}$$

where the parameters $\hat{\lambda}_{cc'}$ are given by $\hat{\lambda}(C)$. Our problem then becomes to choose $C$ to maximize $G(C)$.

In practice there are two problems. The first is that starting with a guess about the number of academic communities, say 5, there is a very large number of ways to partition our thousand or so academic institutions into 5 tiers. To do so properly we would have to check each of those partitions.

The second problem is that the more tiers there are, the more parameters there are that can be used to maximize likelihood.

**Assigning universities to communities given a fixed number of tiers.** The stochastic block model approach handles these problems in two ways. First, it uses a *greedy* algorithm that makes incremental improvements in likelihood by reallocating community members until improvements can't be found. Then it uses a penalty based approach to pick the number of communities.

We should mention that the usual approach uses network connections rather than trade flows to create communities. Secondly, we specify a number of communities exogenously. So some of the econometric theory in that literature does not apply.

To see our approach, start with a fixed number of $k$ tiers. To find the best community structure with $k$ tiers, the algorithm starts by assigning all universities to the same tier. At this stage, each university places different numbers of applicants to other universities, government, the private sector, and teaching universities and whatever other categories exist. Then the likelihood of the data is computed using the method described above. This initial assignment of all universites to the same tier, and the resulting likelihood are then temporarily saved as the *status quo* model.

Then, on each iteration of the algorithm, one university is randomly selected to be reassigned to a randomly chosen tier and the likelihood and poisson placement rates recalculated. If the likelihood increases after the reassignment, the new assignment along with its corresponding likelihood and placement rates are adopted as the new status quo and the process is repeated. If the likelihood doesn't increase on this step, the status quo remains unchanged, and a new random reassignment is made. This process continues until the likelihood stops rising.

To illustrate, the following table show a monte carlo exercise which ran the algorithm 165 times using bootstrapped (with replacement) samples of 10000 placements from the mapinator data. Each university was placed into the tier to which the algorithm most often put them. Cell $(i, j)$ in the table gives the proportion of times that a university who was most often placed in tier $i$ was allocated to tier $j$.

For the top tier, the algorithm was remarkably consistent. Universities in tier 1 were placed in tier 1 by the algorithm over 98% of the time. This is less true for lower tier universities. For example, Tier 4 universities where classified as such just 63% of the time.

| | **Tier 1** | **Tier 2** | **Tier 3** | **Tier 4** | **Tier 5** |
|---|---|---|---|---|---|
| Tier 1 | 0.982054 | 0.00548355 | 0.00398804 | 0.00847458 | 0.0 |
| Tier 2 | 0.0433359 | 0.893875 | 0.057396 | 0.00539291 | 0.0 |
| Tier 3 | 0.00131279 | 0.0534789 | 0.755407 | 0.163822 | 0.0259794 |
| Tier 4 | 0.0 | 0.00381175 | 0.134479 | 0.633615 | 0.228095 |
| Tier 5 | 0.0 | 3.31686e-5 | 0.00202328 | 0.202295 | 0.795648 |

The corresponding table when there were assumed to be just 4 types is given below. Generally with 4 possible tiers the algorithm was more consistent

| | **Tier 1** | **Tier 2** | **Tier 3** | **Tier 4** |
|---|---|---|---|---|
| Tier 1 | 0.9509 | 0.0490998 | 0.0 | 0.0 |
| Tier 2 | 0.0385484 | 0.899714 | 0.0617377 | 0.0 |
| Tier 3 | 0.000514249 | 0.0466252 | 0.873109 | 0.0797514 |
| Tier 4 | 0.0 | 0.0 | 0.0711762 | 0.928824 |

In the discussion below, we use a different approach to evaluate the tier estimate. We use various definitions of closeness just to check how close each individual university is to the tier to which it is assigned.

**Number of Tiers.** As mentioned above, if we evaluate maximum likelihood by itself, the maximum likelihood estimate will always have as many tiers as there are institutions. Instead estimate the number of tiers by using a penalized likelihood approach (Wang and Bickel (2017)). The estimation procedure is simple enough. We find $k$ to solve

$$\min_{k} \left\{ -\ln\left(G(\hat{\mathcal{C}}_*^k)\right) + \delta \frac{k(k+1)}{2} n \ln(n) \right\}$$

where $n$ is the total number of placements in the sample and $G$ is defined by (0.2).

The variable $\delta$ is a tuning parameter. Obviously this is chosen to get an interior solution. However to find it, we used the heuristic procedure defined in Wang and Bickel (2017). Define

$$\beta_k(\delta) = -\ln\left(G(\hat{\mathcal{C}}_*^k)\right) + \delta \frac{k(k+1)}{2} n \ln(n).$$

We calculated function for each $k$ between 2 and 10. Then we calculated

$$w_k(\delta) = \frac{\beta_k(\delta)}{\sum_{k'=2}^{10} \beta_{k'}(\delta)}.$$

We then picked $\delta$ to maximize the entropy measure

$$-\sum_{k=2,10} w_k(\delta) \ln(w_k(\delta)).$$

The resulting value for the tuning parameter was .001. The calculation yields an 'optimal' of 5.

This approach is motivated by a theorem in Wang and Bickel (2017)which says that conditional on there being a true community structure, this approach will find it with probability 1 as the data set gets larger. As mentioned we do not know whether the econometric theory described in their paper applies here. However it is consistent with our monte carlo simulations which suggests that tier assignment becomes less stable as the number of tiers increases.[1]

**Other Measures.** To help evaluate the classification, we created a number of measures that can be used to evaluate the classification. We began by doing a basic structural estimation of the model described in the working paper. Mapinator Classification: Theory

This estimation creates two bits of information. The first is a vector of values for graduates from each tier. This is given by

(0.3)

$$(1.0, 0.5647769732486636, 0.3980183618514731, 0.24138921267998178, 0.1407247181473242)$$

with values listed in order for each tier starting with the top tier which is normalized to 1. We discuss estimation in another paper. At this point we are only trying to find reasonable ways of weighting hires from the different tiers. This method seems as good as any.

We do the same thing with placements by using structural estimates of the average offer values (scaled between 0 and 1) for each hiring community. The table is as follows:

| | **Values** |
|---:|---:|
| TYPE 1 (20 insts) | 0.607084 |
| TYPE 2 (58 insts) | 0.486908 |
| TYPE 3 (180 insts) | 0.407543 |
| TYPE 4 (334 insts) | 0.163591 |
| TYPE 5 (522 insts) | 0.00646005 |
| Public Sector (152 insts) | 0.435835 |
| Private Sector (227 insts) | 0.477125 |
| Postdocs (598 insts) | 0.242724 |
| Lecturers (413 insts) | 0.127996 |
| Unmatched (1 insts) | 0.0 |
| Other Groups (38 insts) | 0.127996 |
| Teaching Universities (642 insts) | 0.419238 |

Again, since we are at this point just looking for a reasonable way to weight placements, this seems as reasonable as any.

[1]It is worth it to note here that assigning all universities to a single tier would be perfectly stable with respect to the simulations we ran.

We then take all the calculated placements and hires to produce a report like the following one for the University of British Columbia.

```
name => University of British Columbia
tier => 2
id => 57
hires=>Any[34, 15, 11, 5, 0]
hiring_value => 48.05680264249607
hiring_quantile => 0.9655172413793104
hiring_ratios => Any[30.422132878859863, 1.0, 2.1548503530352365e-13, 1.076614498

placements => Any[3, 17, 35, 15, 0, 3, 4, 2, 0, 22, 0, 5]
placement_value => 32.61420917644597
placement_quantile => 0.5344827586206896
placement_ratios=>Any[9.444293145204486e-47, 1.0, 2.3661427347761196e-35, 4.51319

euclidian => Any[84.93332679225512, 27.083885631125842, 49.88693636367227, 59.506

ratios => Any[2.873155410103162e-45, 1.0, 5.0986835073640814e-48, 4.858971785838
```

The vector beside 'hires' is the number of recorded hires at ubc from each tier in order. The hiring value is just the product the product of the hiring values vector we used as weights multiplied by the vector of hires. The hiring quantile is just the proportion of all hiring values for tier 2 schools that are no larger than those at UBC.

The hiring_'ratios' vector is just the ratio of the likelihood of UBC's hires given the estimated hiring rate for each tier divided by the likelihood of UBC's hires in tier 2. As you can see from the rates, UBC's hiring performance looks more like that of a tier 1 school.

We do the same thing for placements using the estimated average offer values in each hiring tier. UBC's placement vector is more middle of the road for tier 2.

The 'ratios' vector at the bottom of the dislay is like the hiring and placement ratios except using the product of the probabilith of the hiring and placement vectors.

The 'euclidian' measure is just the euclidian distance between UBC's hiring and placement vector with the tier estimates for each tier.

Using the hiring values for all the universities in the data provides the following distribution of hiring values across the estimated tiers. The blue line represents the distribution of hiring values for tier 1. The hiring value itself is on the horizontal axis, the proportion of universities in the tier with lower hiring values is measured on the vertical axis.

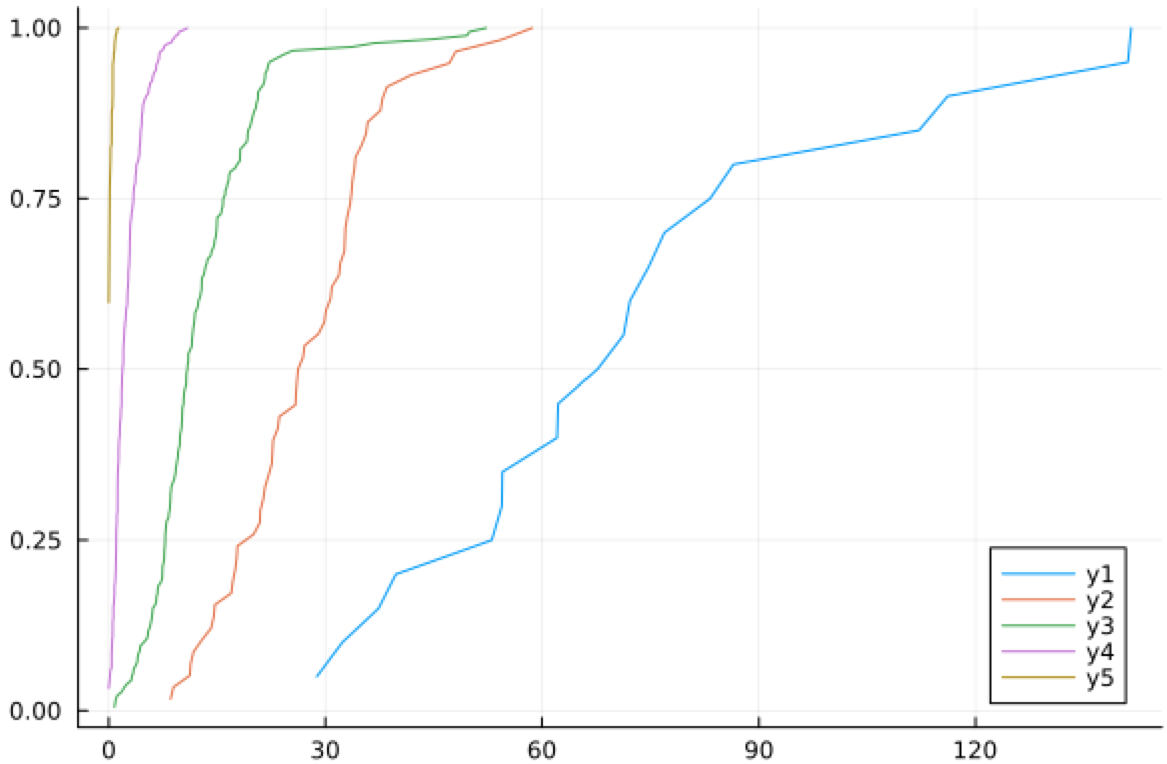


The red line provides the same information for tier 2. There is a stochastic dominance ranking of the distribution. However, note that the supports of the distritions for tier 1 and overlap. There are many tier 2 and 3 universities who are more successful at hiring than some of the tier 1 universties.

The corresponding distributions for placement values are given in this figure:

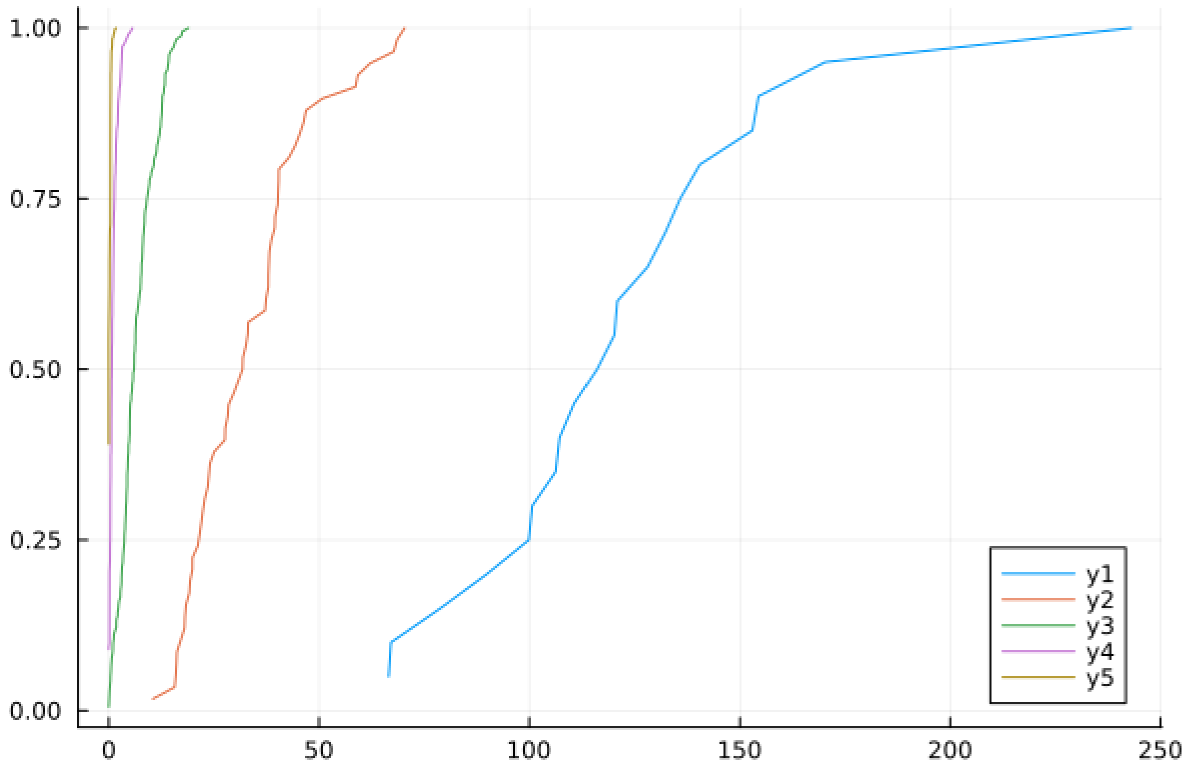


Here the difference between the tiers is very stark. The supports of the various distributions don't even overlap.

Rather than providing the estimates of the membership of each tier here, we refer the reader to the url Universities and Summaries page where you can browse the list of universities assigned to each tier. Clicking on the names of the universities will lead you to a page with an easily browseable list. Clicking on the name of the university will lead you to a summary page with the information described above.

For tldnr types, the critical bits of information are the tier estimate and the value and placement quantile information. The quantile esimates show you where the university you are looking at compares to others in the same tier. The higher the quantile the closer the university is to the next higher tier. The various ratios will give closeness by likelihood. These measures don't tell you how close the univeristy is to a higher or lower tier, they just show how well the university fits within the tier. The maximum value for each ratio should be the one associated with the tier to which they were assigned.

**Literature.** The only paper we have so far found that uses placement data to evaluate departments is Amir and Knauff (2008). They chose 54 departments and created faculty lists for them. They then looked up graduating institutions for the faculty members in each department and looked for links between them by linking the graduating department to the current institution. They then calculated the value of each department as a weighted average of the value of all its graduates.

Weights were calculated using a google page rank like method - all institutions start arbitrary weights. Given those weights, the value of the graduates of each institution can be calculated. The calculated values form the basis of a new set of weights which can used to repeat the calculation, until a fixed point is established.

Their objective was to create a ranking of a relatively small set of institutions. Our categorization creates a type of ranking as well since we investigate a tiered market structure. However, our objective is different since we are trying to determine whether placements can be understood as a reflection of a very coarse information partition.

Mechanically, our approach differs in that it contains a lot more data. Secondly, rather than pre-selecting departments to study, our data generates them automatically since we are just tracing placements of graduates who registered with econjobmarket. We also deal with initial placements rather than studying faculty lists.

The largest difference comes from the fact that by tracking applicants who place outside the top tier, our data gives a better overall perspective, expectially of the role of the top tier. As can be seen from the adjusted adjacency table that follows table , the top tier universities have a market share of just less than 25% (divide the number of placements by the top tier by the total number of placements). This is despite the fact that the top tier consists of only 20 institutions.

**Conclusion.** This paper provides a classification of universities into tiers. Generally graduates of the highest tier universities have the best job prospects, hardly surprising. However the classification provides some empirical insight into how first and second tier universities differ.

At very least the classification quantifies prior beliefs. Universities making an offer to some candidate which they know might be rejected, now have some probabilistic measure of how likely the probability of rejection will be. Undergraduates in economics who are considering going on to do graduate work now know what some of their prospects might be.

Though they are not strictly part of this paper, the offer values we describe in Table  suggest that the average wage offer from a tier 1 university is 25% higher than the value of an offer from a tier 2 universities. This by itself doesn't allow comparison of degrees from different tiers, since the matching probabilities are different. However, if we multiply the offer values given in Table  by the apparent placement rates given in Table  is possible to calculate the expected payoff premium associated with a degree from tier 1 relative to tier 2, it is about 20%, while the premium for a tier 2 degree over tier 3 is 16%. To the extent that choosing an investment in education is an economic decision, this kind of analysis should be valuable.

The classification may also improve the recruiting process since it can be used to auto sort applications into groups whose value is more or less similar. As the cost of processing applications has become very high because of volume, this is also a benefit.

## References

Amir, R., and M. Knauff (2008): "Ranking Economics Departments Worldwide on the Basis of PhD Placement," *The Review of Economics and Statistics*, 90, 185–190.

Karrer, B., and M. E. J. Newman (2011): "Stochastic Blockmodels and community structure in networks," *Physical Review*, 83(1).

Peixoto, T. (2014): "Efficient Monte Carlo and greedy heuristic for the inference of stochastic block models," *Physical Review*, 89(1).

Wang, Y. X. R., and P. J. Bickel (2017): "Likelihood-based model selection for stochastic block models," *Annals of Statistics*, 45(2), 500–528.