LinkedIn Outcome Rankings Under the Hood

In October last year, LinkedIn introduced university rankings based on career outcomes of graduates who received four-year undergraduate degrees from those universities. These rankings are aimed at helping prospective students make one of the most important decisions of their lives: choosing a university for undergraduate study. Today we are excited to present university rankings at the graduate degree level, based on career outcomes of alumni who received **Master's degrees** from universities being ranked.

We believe that insights from career outcomes of alumni provide a novel and, more importantly, an extremely valuable perspective to the process of picking a school. As explained on our previous blog, our goal given the career outcome, is to publish a ranked list of universities which best set their graduates up for desirable jobs in their careers of choice.

Methodology

The sole criteria we use to score universities for a career is:

> "the likelihood of ***relevant graduates*** from that school landing ***desirable jobs*** in that career".

Desirable Jobs

We define a desirable job to be a job at a *desirable company* for the relevant profession. For example, we define *desirable finance jobs* as *finance jobs at companies desirable for finance professionals*.

We let the career choices of our members tell us how desirable it is to work at a company. To illustrate this, **imagine** there are two companies, A and B. If more finance professionals are choosing to leave company A to work at company B, the data indicates that getting a finance job at B is more desirable. The hypothesis here is that when a professional moves from one company to **another,** she gives the company she moves to a strong vote of confidence.

Similarly, the ability of a company to retain its employees is a strong indicator of that employer's attractiveness. So, hypothetically, if A and B are both attracting external employees at similar rates, but A has a much larger employee turnover than B, the data would show B to be a more desirable employer.

We define desirable companies in a profession as those which are:

“the best at **attracting and retaining talent** in that profession”.

Scoring Companies for Desirability

For each professional category, we aggregate the company transitions made by relevant professionals and lengths of their tenures at companies they worked at. We then form a company transition graph comprised of vertices corresponding to companies, and edges connecting companies to reflect transition and retention patterns amongst them. Specifically, each edge is directed and weighted by the number of professionals who transitioned from one company to the other. Retention patterns are captured by so called self-loop edges which, as their name suggests, begin and end at the same company.

Once we construct this graph, we use concepts from the graph ranking algorithm called PageRank to compute the desirability scores of the company vertices. The companies which score the highest are deemed the most desirable for this professional category.

Salient Considerations:

- To ensure that company rankings are reflective of current labor trends, we only look at transitions and tenures in the last five years.

- Company transitions and tenures of only “relevant professionals” are used. For example, say, we need top companies to rank MBA programs in the US for Finance careers. The PageRank would then only consider transitions and tenures of finance professionals who got an MBA from a school in the US. Only transitions and tenures AFTER the completion of the degree are considered.

- When computing tenures, we are careful to consolidate multiple successive positions held by members at a single company.

- Tenures are normalized by median retention/tenures for that career. So, a tenure of 2 years in career paths with shorter retention times of, say, 1.4 years, counts for more than in the case of careers with median tenures of 5 years. Medians are statistical measures robust to potential outliers in the system.

- Companies with very few relevant employees in the profession over the last five years are excluded from consideration in order to reduce possible noise in the data and produce statistically sound results.

Relevant Graduates

Since not every graduate is interested in the same profession, it is only fair to define the relevant graduates as those who end up working in that career. For example, while ranking a graduate program for the category 'Accounting Professionals', we only consider members who attended graduate programs at that school and work as Accounting professionals.

In addition, we want university rankings to reflect recent employment trends. Therefore, we only consider graduates who obtained their degrees within the past eight years.

Scoring Universities

We now have both pieces of the puzzle: the graduates who are relevant to a particular profession, and the desirable jobs for that profession. For each university and profession, we then calculate the percentage of relevant graduates who have obtained desirable jobs. These percentages allow us to rank universities based on career outcomes across different professional areas.

Things to note:

- Our algorithm is knowledgeable about university hierarchy. It knows that HBS, for example, is a sub-school of Harvard University. We leverage this knowledge to consolidate Master's degrees reported as obtained from a university and its sub-school(s). So, if a professional mentions on their profile they got an MBA from New York University, we make sure that they contribute to how NYU's Business school (Stern) performs in our rankings.

- We use confidence intervals around the percentages mentioned above. We are more confident of our estimates if we have more data points to base them on. So, for clarity, say a hypothetical school A has 10 relevant students, 5 out of whom landed desirable jobs. On the other hand, say a school B has 500 graduates who got desirable jobs, out of a total of 1,000 relevant graduates.

  Although the percentage of graduates who obtained desirable jobs is 50% for both A and B, we would trust the score for B more than score for A. Confidence intervals give us a range around the raw score, reflective of degree of statistical support for the score in the data. The higher the number of data points (higher statistical support), tighter the interval. We use the lower bounds of the interval as the score. A would score less than B. Note however that confidence intervals are designed to have a noticeable effect only when the number of data points is extremely tiny. This way, small well-performing schools aren't penalized.

- To make university rankings are less susceptible to potential noise in company rankings, we apply concepts from a popular technique called Monte Carlo Sampling. The crux is to do the following thousands of times: *Perturb the company rankings slightly and rank universities according to scores mentioned above*. These thousands of university rankings are aggregated to produce the eventual universities lists which are published.

Bias Correction:

LinkedIn has a large member base of more than 347 million professionals. However, we understand that not every professional in the world has a LinkedIn profile and there will be those who do but who haven't added recent updates. To mitigate any possible skews in the data due to underrepresentation of certain professionals on LinkedIn, we bias-correct internal statistics against external, publicly available source of truth datasets using the established technique of post-stratification weighting.

Stratification is performed on the basis of such attributes as university, degree and graduation year. A correction weight is computed for each stratum in order to match the distribution of internal data points across strata with the distribution in the source of truth dataset. The resulting correction weights are then applied to individual data points when computing the necessary statistics for ranking universities.

As a hypothetical example, suppose that an externally available government data source tells us that a school has 3,000 graduates for each of the years 2012, 2013 and 2014, while LinkedIn's member-base includes 1,800, 1,200 and 600 graduates for those years, respectively. While computing the percentages required for ranking this university, we would weight graduates from 2014 higher than those from 2012 - to be precise, 3 times as much - as a result of bias correction. Note that this weighting does not change the available sample size (3,600 graduates in the above example), but only affects the relative weights of data points from different strata.

Data Sources

As mentioned earlier, we aggregate career information that our members put on their LinkedIn profiles to draw up a list of top companies for every ranking category and subsequently the best schools whose graduates successfully obtain jobs at those companies. We leverage the investments made by LinkedIn Engineering in Standardization of several data assets on members profiles - companies, schools, degrees and titles.

To correct for potential biases in LinkedIn membership, we use public datasets made available by data collection programs like the Integrated Postsecondary Education Data System in the US.

Conclusion

LinkedIn University Rankings are the only Rankings that leverage real career outcomes information to rank universities at scale - across many professions, degrees and geos. We believe that the information that our members put on their profiles are an invaluable asset that, when aggregated, provide a powerful picture that can guide students to make truly data-driven decisions in order to maximize the likelihood of success in their careers.