Exploring the Structure of Private Foundations

Thomas Dudley, Alexandra Fetisova, Darren Hau

December 11, 2015

1 Introduction

There are nearly 90,000 private foundations in the United States that manage approximately $700 billion in assets, give nearly $50 billion in grants each year, and benefit from large tax breaks [2]. Despite their importance in the social sector and their favorable tax status, private foundations have very little external accountability. Control over private foundations is highly concentrated among the founding members, and is not influenced by shareholders, customers, or even competitors [3]. Theories in organizational behavior suggest that, given the lack of external oversight, private foundations might not be particularly effective vehicles for the funding of charitable services. For instance, the initial charitable purposes of a foundation might become secondary to administrative concerns over time, resulting in fewer grants distributed and increased overhead.

We apply unsupervised learning methods to better understand the landscape of private foundations in the United States, using the financial information of all private foundations submitted to the IRS from 1995-2011. Using dimension-reduction techniques such as PCA, we look for a compact way to summarize the universe of foundations. Using clustering methods such as k-means and Gaussian mixture model, we then find groupings of foundations that have similar financial positions. We discover that foundations can be split into three clearly separable groups in just two dimensions based on either their expenses or their asset composition.

2 Related Work

The foundation movement is a relatively new phenomenon in American philanthropy [3]. Unlike public charities, private foundations tend to contribute to charitable work solely by making grants [8], which reflects a larger shift from participatory activism to "checkbook activism" [6, 5].

The concern that the foundations might accumulate or misuse their money instead of distributing them resulted in a tighter regulation of foundations' spending in the Tax Reform Act of 1969 (B. Marx). Marx uses regression to analyze the foundations' finances before and after 1969. As Marx finds, the regulation indeed discouraged donors from giving to foundations for the purposes other than charity. However, it also increased administrative costs to run a foundation, which repelled donors from giving charitable donations.

The 1969 law also required foundations to spend at least 5% of their funds to qualify for tax breaks. Sansing and Yetman [8] examine whether this minimal requirement affects the spending of foundations using regression analyses. They discover that only about half of foundations adhere to the 5% benchmark, but the ones that spend more tend to grow faster and attract more donors. Moreover, Deep and Frumklin [1] suggest that the 5% grant policy has not provided foundations with the right incentives to distribute their wealth to non-profits.

Although foundations are of growing importance in the United States, the 1969 policy has remained mostly unchanged [4]. At the same time, to our knowledge, there have been no studies applying unsupervised learning to foundations. Our analysis seeks to add to the discussion on the regulation of foundations by grouping foundations by their financial behavior.

3 Dataset and Features

By law, private foundations are required to report key financial information to the IRS via an annual filling of form 990-PF. Required disclosures include information about sources of revenue, operating expenses, financial assets, major program activities, and charitable purposes. This information is digitized by the IRS and made available to the public. The National Center for Charitable Statistics has collected the full set of form 990-PFs for the years 1989 – 2013 [2].

In total, this dataset contains 23 financial features as well as numerous categorical variables describing the charitable purpose of the foundation. Annual counts of organizations vary over time from about 34,000 fillings in 1992 to around 99,000 fillings in 2011. We focus initially on 2011, as it is the most recent year with good coverage. After grouping foundations in 2011, we also examine the development of those groups over time.

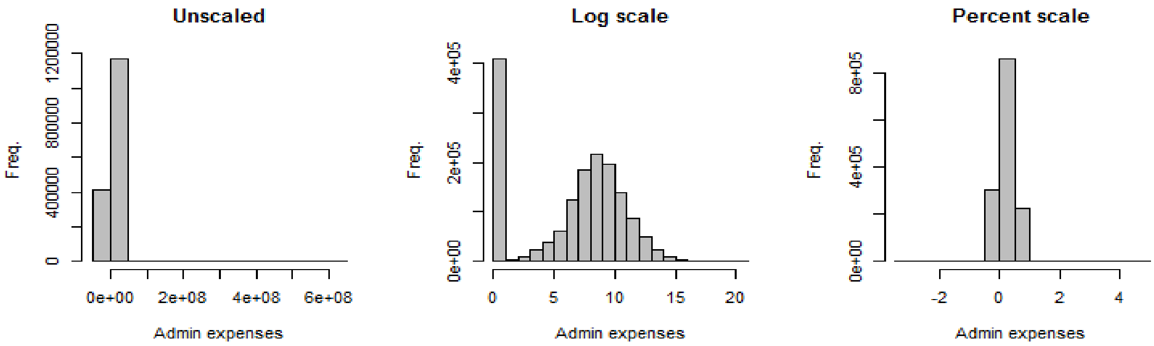

Figure 1: Distribution of Administrative Expenses (2011)

We faced a few challenges in getting this data in a form that could be easily analyzed. The size distribution of private foundations is highly skewed. A few prominent organizations, most notably the Bill and Melinda Gates Foundation, are much larger than the typical private foundation. To correct for this, we identified two possible scaling methods. The first strategy is to use the log function to scale variables. This resulted in data which look to have a Gaussian component, as well as a large point mass at zero. The second strategy was to transform totals with log, and turn component variables into percentages. For example, total revenue would become log(total revenue), while revenue from contributions would become (revenue from contributions)/(total revenue). The advantage here is that we think these transformations might meaningfully distinguish foundations from one another. The downside is we had to drop around 15% of the data because of negative values, unrealistic percentages, etc. An example of the unscaled, log, and percent distributions for administrative expenses is shown in Figure 1. Since the log scaling resulted in higher granularity, we mostly used log-scaled data.

4 Methodology and Results

4.1 PCA

Some of the features in the data are linearly correlated (e.g. different types of revenue and total revenue), which suggests that the data has fewer dimensions than the number of features. Therefore, our first step was to perform Principal Component Analysis in order to reduce the dimensions of the data for better data visualization. Using R, we applied PCA to the scaled 2011 data. We summarize our output by looking at the percent of variance explained by each PC in Figure 2.

In the log-transformed data we see a dramatic amount of variance explained by the first PC. Upon examining the factor loadings, we observe that they are all of the same sign and of similar magnitude. We interpreted this to mean that the first PC was picking up a general foundation size. The plot suggests that at most five PCs are needed to summarize the data, and perhaps as few as two would be sufficient. We find a similar, less dramatic result for the percent transformation. Again the first PC explains a lot of the variance. Additionally, we see elbows after the first three and six PCs.

4.2 K-Means

Our first approach to discover structure in the data was to apply k-means clustering. We fit the model with multiple choices of k, using twenty starts to avoid getting stuck in a local optimum. The ratio of between cluster sum-of-squares and total sum-of-squares provides a good approximation of the variance explained by the model. We present output for both the log and percent scaling in Figure 3. In the log

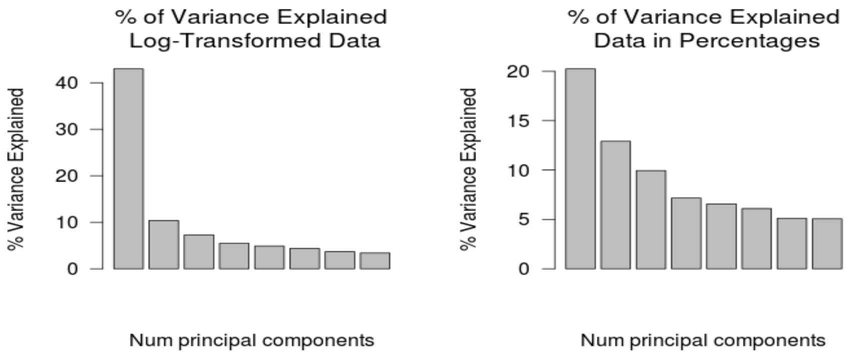

Figure 2: Percentage of Variance Explained vs. Number of Principal Components (2011)

scaling, we see an elbow around 10 clusters. In the percent scaling, there is a much less noticeable elbow around 14 clusters.

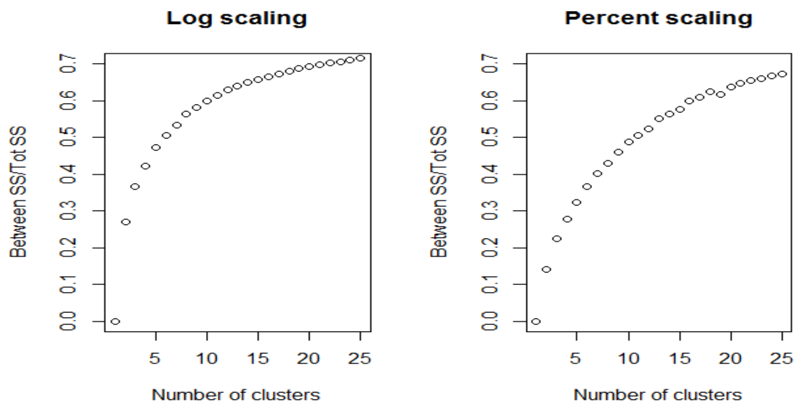

Figure 3: K-Means Output: Cluster Sum-of-Squares / Total Sum-of-Squares vs. Number of Clusters

We next plot the resulting clusters against the two principal components. The plot with five clusters is shown in Figure 4. At this point, we identified two challenges in moving forward. First, the principal components are difficult to interpret, as they were derived from a multitude of variables. Second, k-means does not seem to be able to pick up on the Gaussian-looking data structure that appears in Figure 4.

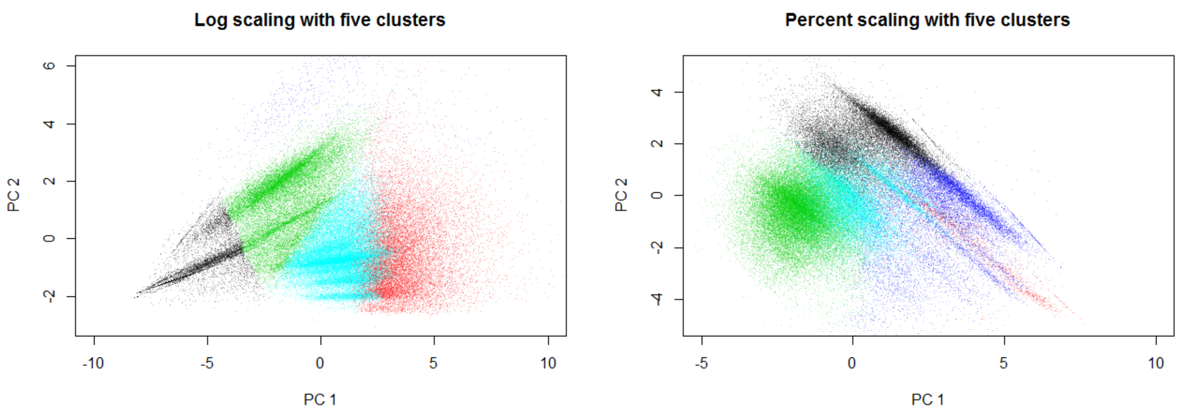

Figure 4: Clusters Plotted Against the First Two Principal Components

4.3 Mixture of Gaussians

To capture the Gaussian-looking distributions, we moved to a Gaussian mixture model, using Python's scikit-learn module [9]. To address the challenge of interpreting the PCs, we further subset the features of the log-scaled data into separate categories: expenses, revenue and assets. We also attempted to eliminate the colinearity from the data by tossing out variables which are linear combinations of other variables. As we are mostly interested in how foundations manage their endowments, we focus on only two groups of variables: expenses and assets. PCA on these subsets of variables results in the principal components shown in Table 1.

Expenses (scaled)	PC1	PC2	PC3	Assets (scaled)	PC1	PC2	PC3
Officer Compensation	-0.54	0.72	-0.43	Government Bonds	0.52	0.82	-0.23
Contributions Paid	-0.55	-0.69	-0.47	Corp. Stock	0.58	-0.54	-0.6
Total Admin Operations	-0.64	-0.02	0.77	Corp. Bonds	0.62	-0.18	0.76
% Variance Explained	57%	25%	17%	% Variance Explained	59%	24%	16%

Table 1: Principal Components in Log-Transformed Data (Expenses and Assets Subsets) for 2011

As before, the first principal component has loadings of similar magnitude and the same size for both Expenses and Assets subsets and explains more than 50% of the variance. The second principal components of the expenses subset suggests that the second dimension is the log of the ratio of a foundation's officer compensation to the contributions the foundation has distributed in grants. For assets, the second principal component is the log ratio of government bonds to corporate stock.

Applying the Gaussian mixture model (GMM) to expenses and assets gives us visually apparent and easily interpretable clusters (Figure 5). We were mostly interested in the clustering based on expenses. On the left plot in Figure 5, as the size of the organization increases, the green-colored group spends relatively more on grants and less on compensation, while the blue and red groups exhibit an opposite relationship: the larger the size, the more they spend on compensation as compared to grants. For assets, we identify three main clusters that correlate with the riskiness of their asset portfolio.

Interestingly, it appears that there is no correlation between the clusters in the expenses GMM and the clusters in the assets GMM. As seen in Figure 5, in which the same labels were applied to clusters for expenses and assets, foundations from all three expense clusters appeared in each of the asset clusters. This indicates that the riskiness of a portfolio has little to do with how a foundation uses its money.

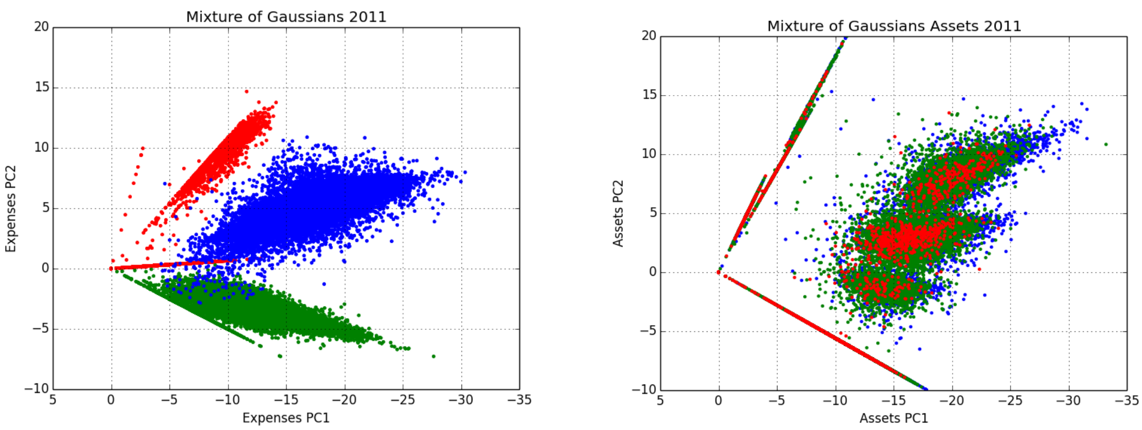

Figure 5: GMM clusters

4.4 Development of Clusters Over Time

Next, we look at the years 1995-2011 to see if there were any trends in cluster development. As Figure 6 shows, all clusters have grown in size over time. Surprisingly, we discovered that the red cluster (high compensation relative to contributions) began emerging only in 1996. Sansing and Yetman [7] identify the period from 1995 to 1998 to be the period of significant asset appreciation for foundations, which might explain the emergence of the red cluster. However, the exact reasons are not certain and require further research.

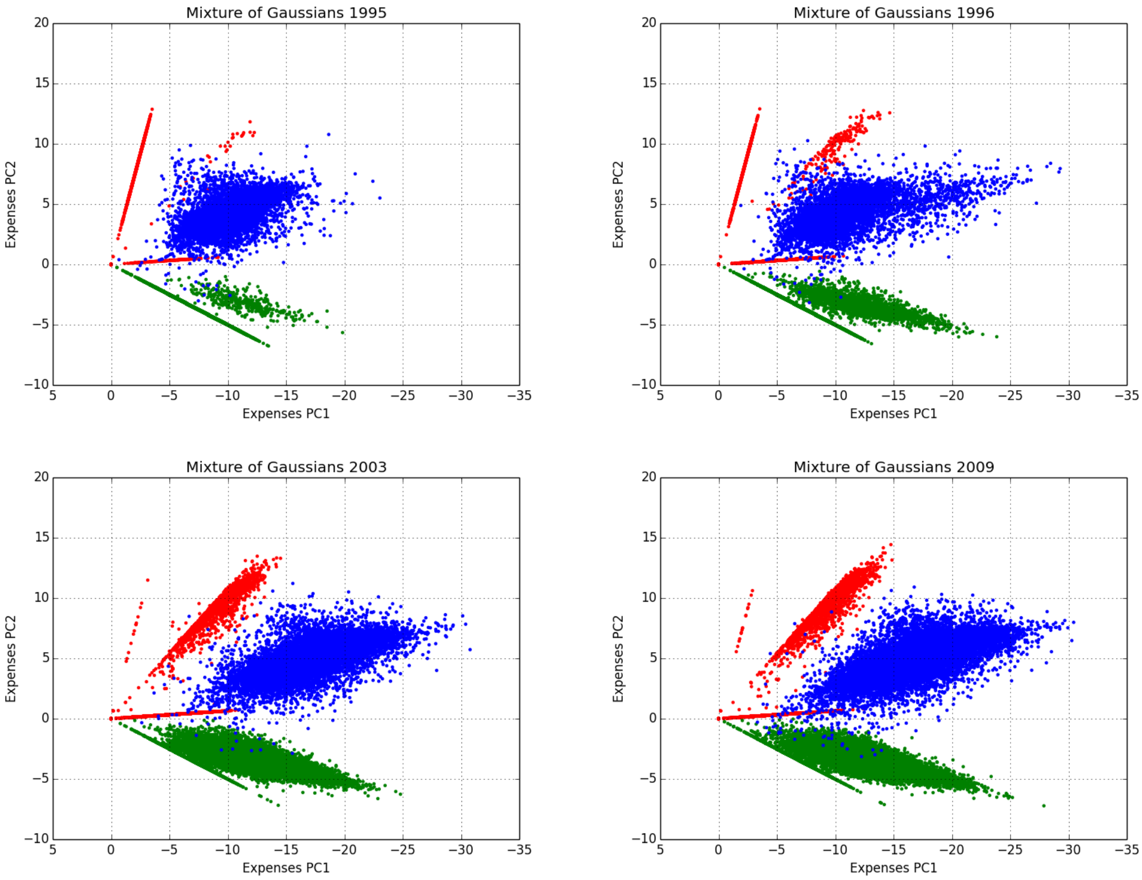

Figure 6: Cluster Emergence and Growth

5 Conclusion and Future Work

While there is a clear split between three categories of private foundations both in terms of expenses and asset allocations, we still cannot conclude with certainty which group, if any, is more efficient in managing their funds and has greater community impact. While it is tempting to assume that foundations with greater overhead (ratio of officer compensation to grant contribution) are fiscally irresponsible, the corresponding clusters could also indicate different types of foundation purposes (i.e. health vs. arts), as well as "quality-over-quantity" foundations. What we can conclude with high confidence is that the portfolio composition of a foundation is not indicative of how it spends its resources.

Further analysis would require a more micro-level approach. The first step would be to identify specific foundations in each of the aforementioned clusters, with the goal of determining what qualitative traits are shared among the foundations in a particular cluster. Further research can also focus on understanding whether foundations migrate from one cluster to another over time, what specific events cause such movement, and what factors cause the foundations to allocate their funds in different ways. By incorporating micro-level analysis, we could get closer to discovering one or more metrics for evaluating a foundation's true effectiveness and impact.

Nevertheless, our results are an important step towards understanding the foundations' spending behavior. Given that we have identified three distinct groups, policymakers may want to consider more customized regulations rather than a one-size-fits-all approach.

References

[1] Akash Deep and Peter Frumklin. *The Foundation Payout Puzzle.* Harvard University Kennedy School of Government, Working Paper No. 9.

[2] IRS Tax Stats. *Domestic Private Foundation And Charitable Trust Statistics.* https://www.irs.gov/uac/SOI-Tax-Stats-Domestic-Private-Foundation-and-Charitable-Trust-Statistics#3. Online; accessed Dec-11-2015.

[3] Tanya D. Marsh. *The Dubious Distinction: Rethinking Tax Treatment of Private Foundations and Public Charities.* Virginia Tax Review, vol. 22.1, p. 137, 2002.

[4] Benjamin Marx. *Has Regulation of Charitable Foundations Thrown the Baby out with the Bath Water?* Journal of Public Economics, vol. 129, pp. 63-76, 2015.

[5] Robert D. Putnam. *Bowling Alone: America's Declining Social Capital.* Journal of Democracy, vol. 6.1, pp. 65-78.

[6] T Skocpol. *Diminished Democracy: From Membership to Management in American Civic Life, Vol. 8.* University of Oklahoma Press, 2013.

[7] Richard Sansing, Robert Yetman. *Distribution Policies of Private Foundations.* Tuck School of Business at Dartmouth, 2002.

[8] Richard Sansing, Robert Yetman. *Governing Private Foundations Using the Tax Law.* Journal of Accounting and Economics vol. 41.3, pp. 363–384, 2015.

[9] Pedregosa et al. *Scikit-learn: Machine Learning in Python*, JMLR 12, pp. 2825-2830, 2011.