The Size Distribution of US Banks and Credit Unions[+]

John Goddard[a], Hong Liu[b], Donal McKillop[c], John O.S. Wilson[d,x]

Abstract

This study examines the firm size distribution of US banks and credit unions. A truncated lognormal distribution describes the size distribution, measured using assets data, of a large population of small, community-based commercial banks. The size distribution of a smaller but increasingly dominant cohort of large banks, which operate a high-volume low-cost retail banking model, exhibits power-law behaviour. There is a progressive increase in skewness over time, and Zipf's Law is rejected as a descriptor of the size distribution in the upper tail. By contrast, the asset size distribution of the population of credit unions conforms closely to the lognormal distribution.

JEL Classification: G21; L10; L16

Keywords: Firm Size distribution, Zipf's Law, Gibrat's Law, Banks, Credit Unions

---

+ The authors gratefully acknowledge the helpful comments of two anonymous referees on an earlier draft of this paper. The usual disclaimer applies.

a Bangor Business School, Bangor University, Bangor, Gwynedd, LL57 2DG, UK. Tel: +44 1248 383221. Email: j.goddard@bangor.ac.uk

b Accounting and Finance, Business School, University of Glasgow, Glasgow, G12 8QQ, UK. Tel: +44 141 330 6124. Email: hong.liu@glasgow.ac.uk

c Queen's University Management School, Queen's University of Belfast, 185 Stranmillis Road, Belfast, BT9 5EE UK.. Tel: +44 28 9097 4852. Email: dg.mckillop@qub.ac.uk

d School of Management, University of St Andrews, The Gateway, North Haugh, St Andrews, Fife, KY16 9SS, UK. Tel: +44 1334 462803. Email: jsw7@st-andrews.ac.uk

x Corresponding author.

The Size Distribution of US Banks and Credit Unions

1. Introduction

This study examines the empirical size distribution of US banks and credit unions. It is well known that most empirical firm-size distributions are highly skewed. If firm sizes are subject to proportional random growth, consistent with Gibrat's Law so that log sizes follow random walks, a lognormal cross-sectional firm-size distribution emerges over time (Gibrat, 1931; Sutton, 1997). There is, however, extensive evidence that the lognormal provides a poor approximation to empirical firm-size distributions in the upper tail, which typically exhibit greater skewness than is consistent with lognormality. Certain modifications to Gibrat's Law, however, are capable of producing a cross-sectional size distribution that exhibits power-law behaviour.[i] A strand in the empirical literature examines the application of lognormal and power-law distributions to cross-sectional firm-size data (Simon and Bonnini, 1958; Quandt, 1966; Lucas, 1978; Cabral and Mata, 2003).

Pareto (1897) describes the distribution of a collection of N subjects ranked by size, where the density function, denoted f( ), obeys a power law in the upper tail:

$$f(X) = \frac{\alpha}{\Theta}\left(\frac{X}{\Theta}\right)^{-(\alpha+1)} \qquad \text{for } X \geq \Theta \qquad [1]$$

where X is size, Θ is the size threshold above which the Pareto distribution applies, and α is constant. Zipf's Law describes the special case α=1 (Zipf, 1949). Axtell (2001) reports that Zipf's Law provides a close approximation to the entire size distribution of US manufacturing firms.[ii]

The recent financial crisis has heightened interest in the role of large banks in the financial system. Some recent bank efficiency studies, based on data from the 1990s and 2000s, have challenged the earlier received wisdom, derived mainly from the analysis of

1980s data, that scale economies were rapidly exhausted beyond relatively modest bank sizes. For example, Hughes and Mester (2011) and Wheelock and Wilson (2012) report evidence of scale economies at all sizes, which cannot be attributed solely to an implicit "too big to fail" TBTF subsidy. Across the empirical banking literature as a whole, however, the evidence as to whether large banks operate at lower average costs than their smaller counterparts is rather weak and contradictory (Davies and Tracey, 2012).

Awareness that the failure of one large bank could tip many other banks into default is widely perceived to have contributed to the emergence among depositors and investors of a TBTF mentality: a presumption that the regulator would intervene and bail out a TBTF bank, rather than permit failure. Steps taken by governments in the US, Western Europe and elsewhere at the height of the crisis in 2008 to bail out numerous large banks that would otherwise have become insolvent have heightened concerns over TBTF. Among the consequences is a weakening of market discipline, with little incentive for depositors and investors to monitor risk-taking. Competition is distorted, because the TBTF bank benefits from an implicit safety-net insurance public subsidy (O'Hara and Shaw, 1990; Morgan and Stiroh, 2005; Mishkin, 2006; Schmid and Walter, 2009; Brewer and Jagtiani, 2012).

Historically, geographic and product market regulations have constrained the growth of US banks, and the largest US banks were small by comparison with the largest European and Japanese banks (DeYoung, 2010). Since the 1980s, however, financial deregulation has eased many of the earlier constraints on growth.[iii] Consolidation through M&A (merger and acquisition) reduced the number of commercial banks, from around 14,400 in 1980 to around 6,400 in 2010.[iv] Industry concentration measured, for example, using the ten-firm concentration ratio for the assets of all separately constituted commercial banks (irrespective of ownership), has increased from 0.300 in 1980 to 0.555 in 2010. It seems likely that many of the banks that disappeared were simply too small to compete effectively in a less heavily regulated and more highly competitive banking market.

Despite the increasing dominance of large banks, and their importance for financial stability, to our knowledge only one previous study has examined the statistical characteristics of the firm-size distribution for banks. Janicki and Prescott (2006), henceforth JP, report estimation results obtained by fitting a Pareto distribution to the upper tail of assets, deposits, loans and employees data for US commercial banks and Bank Holding Companies (BHC) at ten- or fifteen-year intervals between 1960 and 2005. Zipf's Law is found to provide a close approximation to the data in all of the years examined except 2005, for which the upper tail is heavier than is consistent with Zipf's Law.

An important technical limitation of the methodology used by JP is the imposition of an arbitrarily chosen value for the threshold parameter Θ, corresponding to the asset size of the bank ranked 3,000th in the firm-size distribution. A consequence of this limitation is that JP are unable to provide an accurate description of changes in the shape of the firm-size distribution, especially changes in the location of Θ, over time. The present study overcomes this limitation, by applying an estimation method that permits a choice between fitting the lognormal distribution over the entire range of firm sizes, and fitting a combination of a truncated lognormal distribution for smaller firms and a Pareto distribution for firms above the size threshold Θ. In the latter case, the choice of Θ is determined by the data, and Θ is permitted to vary over time. It is shown that the true location of Θ for commercial banks is, in fact, considerably nearer the upper end of the firm size distribution than the arbitrary location assumed by JP; and that the Pareto upper tail describes the size distribution of a correspondingly smaller proportion of the population. Between 1995 and 2010, for example, the proportion of the population of commercial banks represented by the Pareto upper tail fell from 6.2% (611 banks) to 2.8% (180 banks). Owing to the highly skewed nature of the firm-size distribution, however, these 180 banks accounted for more than 85% of banking industry assets in 2010.

This study also compares the size distribution of commercial banks with that of (not-for-profit) credit unions. The growth of credit unions is constrained by common-bond restrictions on the groups with which they are permitted to transact, despite moves in recent years towards the easing of several of the restrictions. It is of interest to examine whether the combination of a non-profit orientation and a restrictive regulatory environment generates a similar or differently-shaped firm-size distribution over the long term. The empirical results suggest that the shape of the size distribution does, indeed, differ markedly between commercial banks and credit unions. The size distribution of commercial banks is described by a truncated lognormal distribution and a Pareto distribution in the upper-tail. There is a pronounced trend in the upper-tail shape parameter that reflects an increase in skewness over time. By contrast, the lognormal distribution provides an accurate description of the entire size distribution of credit unions throughout the observation period.

The rest of this study is structured as follows. Section 2 describes the estimation method. Section 3 reports estimation results for the population of US commercial banks for the period 1976-2010. Section 4 reports an application of the same methods to the population of US credit unions over the period 1995-2010. Finally, Section 5 summarizes and concludes.

2. Estimation Method

Let X_i denote the assets of firm i in a particular year, and let $x_i=\ln(X_i)$. Let $X_{[i]}$ denote the value of the i'th observation when the firms are ranked in descending order of asset size, so $X_{[1]} \geq X_{[2]} \geq ... \geq X_{[n]}$, and let $x_{[i]}=\ln(X_{[i]})$. Let $\Theta_k = X_{[k]}$ denote some threshold value of k that is, initially, assumed to be pre-selected, and let $\theta_k = x_{[k]}$.

We examine two candidate distributions for X_i:

(i) $x_i \sim N(\mu_0, \sigma_0^2)$ for all i

(ii) $x_i \sim TN(\mu_k,\ \sigma_k^2, -\infty, \theta_k)$ for $x_i < \theta_k$, and $X_i \sim Par(\Theta_k, \alpha_k)$ for $X_i \geq \Theta_k$

In (i), the distribution of X_i is lognormal with mean and variance parameters μ_0 and σ_0^2, for all i. In (ii), the distribution of X_i is truncated lognormal with mean and variance parameters μ_k and σ_k^2, and an upper truncation point θ_k, for $X_i < \Theta_k$; and Pareto with location and shape parameters Θ_k and α_k, for $X_i \geq \Theta_k$.

For (i), the maximum likelihood estimators of μ_0 and σ_0^2 are $\hat{\mu}_0 = n^{-1}\sum_{i=1}^{n} x_i$ and $\hat{\sigma}_0^2 = n^{-1}\sum_{i=1}^{n}(x_i - \hat{\mu}_0)^2$. The maximized log-likelihood function is $\ln(L_0)$ = $\sum_{i=1}^{n} \ln \phi[(x_i - \hat{\mu}_0)/\hat{\sigma}_0]$, where ϕ is the standard normal density function.

For (ii), the likelihood function that would be formed over the two segments of the distribution is discontinuous at the truncation point θ_k. Accordingly, we estimate the parameters μ_k, σ_k^2, k, Θ_k and α_k in two stages.

Stage 1

Let $\bar{x}_k = (n-k)^{-1}\sum_{i=k+1}^{n} x_{[i]}$ and $s_k^2 = (n-k-1)^{-1}\sum_{i=k+1}^{n}(x_{[i]} - \bar{x}_k)^2$ denote the sample mean and sample variance of all $x_{[i]} < \theta_k$. For $x_i \sim TN(\mu_k,\ \sigma_k^2, -\infty, \theta_k)$:

$$E(x_i) = \mu_k - \phi[(\theta_k - \mu_k)/\sigma_k]/\Phi[(\theta_k - \mu_k)/\sigma_k] \qquad [2]$$

$$V(x_i) = \sigma_k^2\{1 - [(\theta_k - \mu_k)/\sigma_k]\phi[(\theta_k - \mu_k)/\sigma_k]/\Phi[(\theta_k - \mu_k)/\sigma_k]\} + \sigma_k^2\{\phi[(\theta_k - \mu_k)/\sigma_k]/\Phi[(\theta_k - \mu_k)/\sigma_k]\}^2 \qquad [3]$$

where Φ is the standard normal distribution function. The estimates $\hat{\mu}_k$ and $\hat{\sigma}_k^2$ are obtained by selecting the pair of values for μ_k and σ_k^2 that minimise the distance function $D = [E(x_i) - \bar{y}_k]^2 + [V(x_i) - s_k^2]^2$.

Stage 2

Estimation of α_k proceeds by applying maximum likelihood to the observations $X_{[1]},...,X_{[k]}$, which are assumed to follow the Pareto distribution. The log-likelihood function takes the form $\sum_{i=1}^{k} \ln[(\alpha_k / \Theta_k)(X_{[i]} / \Theta_k)^{-(\alpha_k+1)}]$. Zipf's Law describes the special case α_k=1, for which Stage 2 (the estimation of α_k) is not required. A grid-search is used to identify the value of k that maximises the pseudo-log-likelihood function, constructed as follows:

$$\ln(L_k) = \sum_{i=k+1}^{n} \ln\{\phi[(x_{[i]} - \hat{\mu}_k)/\hat{\sigma}_k]\}$$

$$+ \sum_{i=1}^{k} \ln\{(\hat{\alpha}_k / \Theta_k)(X_{[i]} / \Theta_k)^{-(\hat{\alpha}_k+1)} X_{[i]} \{1 - \Phi[(x_{[i]} - \hat{\mu}_k)/\hat{\sigma}_k]\}\} \quad [4]$$

In the special case described by Zipf's Law, $\hat{\alpha}_k$ is replaced by one in [4].

By convention, the rank-size relationship is commonly depicted using a plot of log rank against log size. If asset sizes follow a Pareto distribution in the upper tail, the log rank-log size relationship is linear over the relevant range of asset-size values, with a slope of $-\alpha_k$. This follows from

$$j = k\int_{X_{[j]}}^{\infty} f(z)dz = k\Theta_k^{\alpha_k} X_{[j]}^{-\alpha_k}$$

where $j \le k$ is the continuous analogue of the (discrete) rank, and f is the density function of the Pareto distribution. Applying a log transformation yields

$$\ln(j) = \ln(k) + \alpha_k\theta_k - \alpha_k x_{[j]}$$

In the special case described by Zipf's Law, the slope of the log rank-log size plot is –1.

3. The Size Distribution of US Banks

The data source for the empirical analysis of the assets size distribution of US commercial banks is Reports on Condition and Income (Call Reports) provided by the Federal Financial Institution Examination Council. Data on savings banks, savings and loan associations, investment banks, mutual banks and credit-card banks are excluded. Table 1 reports descriptive statistics based on the fourth-quarter reports at five-year intervals from 1980 to 2010. The upper panel reports an analysis in which commercial banks are the units of observation, disregarding ownership. Accordingly, in the 'all commercial banks' population, commercial banks that are constituents of the same BHC are treated as separate observations. Most previous studies of the evolution of the US banking industry focus on banks, rather than BHC (Berger et al., 1995; Jones and Critchfield, 2005). In part this is because many price and non-price decisions are taken at bank level. Regulation concerned with competition and anti-trust issues also tends to focus on the bank, and not the BHC.

For purposes of comparison with the JP study, the lower panel of Table 1 reports an analysis in which 'ownership groups' are the units of observation. The data for any banks that are constituents of the same BHC are aggregated to form a single observation, while the data for independently-owned commercial banks are retained as separate observations.[v] BHC accounts report data that reflects the summation of both banking and non-banking business. Since the 1990s, and especially since the passing of the Gramm Leach Bliley Act in 1999, the importance of non-banking business to BHC has increased, especially for larger entities (Avraham et al. 2012; Copeland, 2012). In 2008, investment banks such as Goldman Sachs and Morgan Stanley (which hold the bulk of their assets outside banking subsidiaries) converted to BHC status. In order to avoid the inclusion of non-banking assets, bank subsidiaries only are included in the construction of aggregated BHC data.[vi] Off balance sheet business is also excluded from the analysis.[vii]

According to Table 1 the number of commercial banks has fallen steadily throughout this period, owing primarily to M&A. Industry concentration fell slightly during the 1980s, but has increased throughout the 1990s and 2000s. These trends are apparent in the five-, ten- and twenty-firm concentration ratios, and the Herfindahl index.

Table 2 reports the empirical analysis of patterns in the assets size distribution of US banks, for the 'all commercial banks' and 'ownership groups' definitions. The data are yearly for the period 1976-2010, but to conserve space Table 3 reports results only at five-year intervals from 1980 to 2010. Figures 1 and 2 present summary yearly results for the entire observation period, in the form of plots of the estimated k and α_k. The comparison between the two candidate distributions, (i) lognormal and (ii) truncated lognormal with Pareto upper tail, favours (ii) in every year reported in Table 3, and in all of the intermediate years that are not reported. A Kolmogorov-Smirnov test rejects the null hypothesis of lognormality for the entire size distribution in every case.[viii]

For the 'all commercial banks' population definition, the percentile of the asset size distribution at which the estimated threshold parameter Θ_k is located varies between a minimum value of 95.5 and a maximum of 96.8 between 1976 and 1990, when there is no discernible trend. Between 1991 and 1995 this percentile drops to 93.8; but from 1996 onwards there is a steady trend in the reverse direction. The percentile of the asset size distribution at which Θ_k is located attains its highest value of 97.2 in 2010 (see Figure 1). Despite the decrease in the number of banks located in the upper tail, their share of total assets, measured by the k-firm concentration ratio CR_k, increases steadily over the observation period, from 0.719 in 1980 to 0.856 in 2010. While the Pareto upper tail comprises only 2.8% of all commercial banks operating in 2010, these banks account for more than 85% of total banking-sector assets.

The estimated upper-tail shape parameter $\hat{\alpha}_k$ is below one, but not significantly different from one, for 1976 and 1977, the first two years in the data set. For every subsequent year, however, the estimated α_k is significantly below one; and there is a downward trend that is most pronounced during the 1980s and early 1990s. Accordingly, Zipf's Law is rejected as a descriptor of the size distribution in the upper tail, in favour of a Pareto power-law distribution with a heavier upper tail than would be expected under Zipf's Law. Relative to the rest of the distribution, the weight of the upper tail has increased since the mid-1990s, reflected in the downward trend in the estimated α_k.

For the 'ownership groups' population definition, a trend in the rank at which the estimated threshold Θ_k is located is apparent from the mid-1980s onwards. Over the entire period 1976-2010, the percentile of the asset size distribution at which the threshold value of k is located attains a minimum value of 94.4 in 1983 and a minimum of 97.6 in 2008. The estimated upper-tail shape parameter α_k is significantly below one in every year. Again Zipf's Law is rejected as a descriptor of the size distribution in the upper tail of the 'ownership groups' population, in favour of a Pareto power-law distribution with a heavier tail than is required for conformity with Zipf's Law.

Figure 3 reports scatter plots, on a logarithmic scale, of the relationship between rank and asset size, for the 'all commercial banks' population in 1980, 1990, 2000 and 2010. In each case, the theoretical rank-size plot for a lognormal distribution with mean and variance corresponding to the sample estimates over the entire sample is located (thin dotted line). The equivalent plots for the upper tail, assuming banks larger than the estimated size threshold k follow the Pareto distribution in accordance with either Zipf's Law, or the fitted shape parameter α_k, are also located (thick dotted line, and continuous line, respectively). Accordingly, the continuous line represents the maximum likelihood estimate of the Pareto upper tail, while the thick dotted line represents an upper tail with a shape parameter in accordance with Zipf's Law.[ix] The maximum likelihood estimates provide a close, if less than perfect, representation of the upper tail, with a slight tendency to overstate the number of firms at the top end of the distribution (more so in 1980, 1990 and 2000 than in 2010).

Table 3 reports summary results for three alternative firm size measures: loans, deposits and employees. As before, in every estimation the choice between the two candidate distributions favours (ii) truncated lognormal with Pareto upper tail. For each size measure and for each of the two population definitions, the threshold value of k and the estimated upper-tail shape parameter α_k are reported. The results are similar to those for the assets size measure. In all cases k diminishes over the entire observation period, and the estimated α_k shows a tendency to decrease in magnitude.[x]

The notion that the commercial banks population divides into two discrete categories defined by scale is consistent with the characterization of the evolution of the US commercial banking industry described by DeYoung, Hunter and Udell (2004), henceforth DHU. Prior to the 1980s, the industry was dominated by a large number of small, community-based banks, offering differentiated or customized loan products and a highly personalized service, and operating at relatively high unit cost. Subsequently, deregulation and technological change created new strategic opportunities for growth that were realized initially by within-market M&A, and later by larger-scale M&A that greatly extended the market reach of the merged entities. Those banks that grew most aggressively came to bear less resemblance to community banks, by adjusting to a high-volume low-cost retail banking model reliant on scale economies, in which automated production and distribution processes deliver standardized products and services at low unit cost. A consequence has been the emergence of "... a strategic wedge between the large and growing banks on the one hand and the smaller community banks on the other" (*op cit*, p110). Although many community banks have also grown, they continue to operate under a traditional high unit-cost and high value-added retail banking model.

A rule-of-thumb used by DHU defines community banks as those with assets below around USD 1 billion (2001 prices). This figure corresponds quite closely to the threshold asset size value of USD 1.2 billion for the year 2000 that is quoted in Table 2, below which the fitted size distribution is truncated lognormal and above which Pareto. According to Table 2, the cut-off asset size value that separates the smaller community banks from the larger ones that have outgrown community status increased to around USD 2 billion in 2005, and USD 3 billion in 2010 (all values in current prices). The number of banks in the upper tail fell from 325 (96.1 percentile of the assets-size distribution) in 2000 to 267 (96.4 percentile) in 2005, and 180 (97.2 percentile) in 2010.

4. The Size Distribution of US Credit Unions

Credit unions are cooperative not-for-profit financial organizations under mutual ownership that provide basic banking services to their members. A credit union's worth is based on book value rather than market value. Accordingly, credit unions are not subject to market-driven expectations for growth and earnings performance. Credit unions are subject to a common bond, which defines the groups with which each credit union is permitted to transact. The common bond might be defined by residence in a particular geographic area, employment in a particular company or industry, or religious or some other affiliation. Deregulation has, to a limited extent, eased the constraint on the growth of individual credit unions imposed by the common bond. The Credit Union Membership Access Act of 1998 permitted federally-chartered credit unions to operate with a multiple common bond, or to transact with any resident of a geographical area defined as a 'community'. Similarly the Gramm-Leach-Bliley Act in 1999 removed many restrictions previously imposed on the activities of banks. Neither banks nor credit unions face restrictions on the prices they charge for specific products. Credit unions, however, face some restrictions on their volumes of small-business lending, while banks are subject to limits on lending concentration. In common with commercial banking, the US credit union industry has experienced significant consolidation through M&A in recent decades.[xi]

Comment [DM1]: Should consider deleting this sentence – mention the Act earlier

The data source for the empirical analysis of the assets size distribution of US credit unions is the '5300 Call Reports', published by the National Credit Union Association (NCUA). Table 4 reports an analysis of patterns in the assets size distribution of US credit unions, based on the December reports at five-year intervals within the period 1995-2010. A steady decline in the number of credit unions, which has been underway since the 1970s, has seen numbers fall from around 11,700 in 1995 to around 7,300 in 2010. Consistent with the pattern for US banks, industry concentration, measured by the five-, ten- and twenty-firm concentration ratios and the Herfindahl index, has increased steadily throughout this period; but concentration remains considerably lower than it is for commercial banks.

The comparison between the two candidate distributions, (i) lognormal and (ii) truncated lognormal with Pareto upper-tail, favours (i) in each of the years reported in Table 1, and in all of the intermediate years that are not reported. For most years, the estimation procedure fails to identify a value of k for which the maximised value of the pseudo-log-likelihood function in (ii) is larger than the corresponding value in (i); and for those years for which such a value of k is identified, the difference between the maximized value of the pseudo-log-likelihood function for (i) and (ii) is marginal and not statistically significant on any conventional criterion. A Kolmogorov-Smirnov test fails to reject the null hypothesis of lognormality at the 0.01 level for all years within the observation period; although there are five rejections at the 0.05 level (for years 2003-2006 inclusive, and 2010). Figure 4 shows the scatter plot of the relationship between log rank and log asset size for 2010. Only three credit unions, each with assets of more than USD 10 billion, were larger than would be expected if the entire asset-size distribution were lognormal in 2010.[xii] Overall the lognormal distribution appears to provide a satisfactory, but not a perfect, representation of the asset size distribution of US credit unions. There is no evidence of conformity with a power law in the upper tail of this firm-size distribution.

5. Conclusion

This study examines the size distribution of US financial institutions. For commercial banks, the firm size distribution is accurately described by a truncated lognormal distribution with a Pareto or power-law distribution in the upper tail. Zipf's Law is rejected as a descriptor of the firm size distribution in the upper tail for all except the first two years of a 1976-2010 observation period. A trend in the upper-tail shape parameter reflects a progressive increase in skewness over time. In 2010, the Pareto upper tail contains only 2.8% of all commercial banks, but these banks account for 85% of total banking-sector assets.

Deregulation and financial liberalization has eliminated many of the constraints on the growth of US banks that were effective prior to the 1990s. The question as to whether the existence of very large banks is justified on efficiency grounds, with economies of scale yielding average cost savings at the upper end of the firm-size distribution, remains largely unresolved in the empirical banking literature. The suggestion that governments would always bail out "too big to fail" (TBTF) large banks for fear of a contagion effect that could precipitate systemic failure, seemingly more than amply justified by events at the height of the financial crisis at the end of the 2000s, has raised concerns that large banks are subject to inadequate competitive or market discipline. The descriptive form of evidence presented in this paper does not allow us to resolve these crucial issues. However, a descriptive analysis demonstrating that the commercial banks population divides into two distinct categories defined by scale is highly suggestive, and consistent with a characterization of the US commercial banking industry that distinguishes between a large but steadily shrinking population of small competitive community-based banks, and a smaller but increasingly dominant group of large banks operating in an environment in which the intensity of competition is moderated by the existence of an implicit TBTF public subsidy.

By contrast, the lognormal distribution describes accurately the entire firm-size distribution of credit unions throughout most of a 1995-2010 observation period, when there is no evidence of power-law behaviour in the upper tail, despite some indications of a deterioration in the goodness-of-fit for the lognormal distribution during the second half of this period. Overall the extent of departure from lognormality in the entire credit union firm-size distribution appears rather modest. We conjecture that the combination of a non-profit orientation and a more restrictive regulatory environment accounts for the evolution of a firm-size distribution for all credit unions that is similar to that of the smaller commercial banks, but markedly different to that of the larger banks at the upper end of the size distribution.

Endnotes

[i] These include: proportional growth subject to a reflecting lower barrier (a minimum size threshold beneath which no firm can fall); and proportional growth with deaths at a rate that is inversely proportional to size, and births at a constant size (Gabaix, 2009).

[ii] See also Stanley et al. (1995) and Growiec et al. (2008). For empirical evidence on the size distribution of cities, regions and countries, see Gabaix (1999a,b); Eeckhout (2004); Rose (2006);and Luttmer (2007).

[iii] For example, the McFadden Act of 1927, which prohibited interstate branch banking, was repealed by the Riegle-Neal Interstate Banking and Branching Efficiency Act of 1994; and the Glass-Steagall Act of 1933, which prohibited commercial banks from transacting other financial services including investment banking and insurance, was repealed by the Gramm-Leach-Bliley Financial Services Modernization Act of 1999. Berger et al. (1995) and DeYoung (2010) describe the evolution of the US banking industry.

[iv] According to DeYoung (2010), around 350 commercial banks were acquired each year during the1980s, around 500 each year during the 1990s, and around 300 each year during the first half of the 2000s. More than 10,000 bank charters were terminated owing to acquisitions (excluding acquisitions enforced by the Federal Deposit Insurance Corporation owing to severe financial distress) between 1980 and 2005, and acquisitions accounted for more than 80% of all bank charters that terminated during this period.

[v] Commercial banks' call reports include a variable which reports (if applicable) the parent BHC code number.

[vi] Banks that form part of a BHC are subject to bank supervision and regulation. For example, they raise insured deposits, are subject to risk-based capital regulation and prompt corrective action, and have access to lender of last resort facilities via the Federal Reserve discount window.

[vii] Although accounting rules differ between countries, the most common way to record items below the line is in notes to the accounts, in supervisory reports, within banks' internal reporting systems, or in some cases not at all. Call reports provide information on the extent of certain OBS assets (such as loans held for sale) and obligations (contingent liabilities such as letters of credit). Reporting varies widely over time. Only 6% of banks held loans for sale in 1991; in 2010 this figure was around 25%. The proportion of banks that reported involvement in securitisation was smaller. Credit unions typically do not report OBS business.

[viii] The maximised value of the pseudo-log-likelihood function (not reported in Table 1) for (ii) is substantially larger than the corresponding value for (i) in every case.

[ix] In Figure 3, both the maximum likelihood estimate of the Pareto upper tail, and the Zipf's Law upper tail, commence from the maximum likelihood estimate of the cut-off threshold.

[x] The finding that the upper tail of the log size distribution can be represented by a Pareto distribution reflects skewness and kurtosis coefficients for the log size distribution that are persistently higher than the values (of zero and three respectively) associated with the normal distribution. This approach, however, leaves open the possibility that the entire log size distribution could be represented by some other probability distribution that allows for excess kurtosis. The Student-t distribution allows flexibility in modelling excess kurtosis through variation in the degrees of freedom parameter. Experimentation with use of the Student-t distribution as an alternative to the normal distribution produced similar results for all four log size measures. As before, the formulation representing the upper tail of the log size population using the Pareto distribution and the rest of the population using a truncated Student-t distribution was preferred to the formulation representing the entire population using a Student t-distribution in every case. However, the threshold values of k that locate the switch between the truncated Student-t and Pareto distributions were persistently smaller than the corresponding threshold values obtained using the normal and Pareto distributions, as reported in Tables 2 and 3. This pattern seems plausible: the Student-t distribution describes accurately a somewhat larger proportion of the upper tail than the normal distribution; but neither of these distributions adequately represents the upper tail in its entirety.

[xi] Walter (2006) and Goddard et al. (2009) describe the evolution of the US credit union sector. Wheelock and Wilson (2011) present evidence that credit unions are, on average, too small to benefit from economies of scale.

[xii] The largest credit union, Navy Federal with assets of USD 44.2 billion in 2010, would have ranked 36th in the distribution of commercial banks by asset size in the same year.

References

Avraham, Dafna, Patricia Selvaggi, and James Vickery. (2012) “A structural view of US bank holding companies.” *Federal Reserve Bank of New York Economic Policy Review*, 18, pp. 65-81.

Axtell, Robert L. (2001) “Zipf Distribution of U.S. Firm Sizes.” *Science*, 293, pp. 1818–1820.

Berger, Allen N., Anil K. Kashyap, and Joseph M. Scalise, (1995). “The Transformation of the US Banking Industry: What a Long Strange Trip It’s Been.” *Brookings Papers on Economic Activity*, 2, pp. 54-219.

Brewer, Elijah and Julapa A. Jagtiani. (2012) How much did banks pay to become too-big-to-fail and to become systemically important? *Journal of Financial Services Research* (forthcoming).

Cabral, Luis M.B., and Mata Jose, (2003) “On the Evolution of the Firm Size Distribution: Facts and Theory.” *American Economic Review*, 93, pp. 1075-1090.

Copeland, Adam. (2012) “Evolution and heterogeneity among larger bank holding companies: 1994 to 2010.” *Federal Reserve Bank of New York Economic Policy Review*, July, 18, pp. 83-92.

Davies, Richard and Belinda Tracey (2012) “Too big to be efficient? The impact of implicit funding subsidies on scale economies in banking.” Bank of England Working Paper.

DeYoung, Robert, William C. Hunter, and Gregory F. Udell, (2004) “The Past, Present and Probable Future for Community Banks”, *Journal of Financial Services Research*, 24, pp. 85-133.

DeYoung, Robert, (2010) Banking in the United States, in Berger, Allen N., Philip Molyneux, and John O.S. Wilson (eds.) *Oxford Handbook of Banking*. Oxford: Oxford University Press.

Eeckhout, Jan, (2004) “Gibrat's Law for (All) Cities.” *American Economic Review*, 94, pp. 1429–1451.

Gabaix, Xavier, (1999a) “Zipf's Law: An Explanation.” *Quarterly Journal of Economics*, 114, pp. 739–767.

Gabaix, Xavier, (1999b) “Zipf's Law and the Growth of Cities.” *American Economic Review, Papers and Proceedings*, 114, pp. 129-132.

Gabaix, Xavier, (2009) “Power Laws in Economics and Finance.” *Annual Reviews in Economics*, 1, pp. 255-294.

Gibrat, Robert, (1931) *Les Inegalities Economiques*, Paris: Sirey.

Goddard, John, Donal G. McKillop, and John O.S. Wilson, (2009) “Which Credit Unions Are Acquired?” *Journal of Financial Services Research*, 36, pp. 231-252.

Growiec, Jakub, Fabio Pammolli, Massimo Riccaboni, and Eugene, H. Stanley, (2008) “On the Size Distribution of Business Firms.” *Economics Letters,* 98, pp. 207-212.

Hughes, Joseph, and Loretta J. Mester, (2011) “Who Said Large Banks Don’t Experience Scale Economies? Evidence from a Risk-Return-Driven Cost Function.” Federal Reserve Bank of Philadelphia Working Paper No. 11-27

Janicki, Hubert P., and Edward C. Prescott, (2006) “Changes in the size distribution of U.S. banks: 1960-2005.” *Federal Reserve Bank of Richmond Economic Quarterly,* 92, pp. 291-316.

Jones, K.D., and T.S. Critchfield, (2005) “Consolidation in the US banking industry: is the long, strange trip about to end?” *Federal Deposit Insurance Corporation Banking Review*, 17, pp. 30–61.

Lucas Jr, Robert E., (1978) “On the Size Distribution of Business Firms.” *Bell Journal of Economics,* 9, pp. 508-523.

Luttmer, Erzo G.J. (2007) “Selection, Growth, and the Size Distribution of Firms.” *Quarterly Journal of Economics*, 122, pp. 1103-1144.

Mishkin, Frederic.S., (2006). “How Big a Problem is Too Big To Fail?” *Journal of Economic Literature* 44, pp. 988-1004.

Morgan, Donald P., and Kevin J. Stiroh, (2005) “Too Big to Fail after all these Years.” *Federal Reserve Bank of New York,* Staff Report No. 220, September.

O’Hara, Maureen, and Wayne Shaw, (1990) “Deposit Insurance and Wealth Effects: The Value of Being ‘Too Big to Fail’.” *Journal of Finance,* 45, pp.1587-1600.

Pareto, Vilfredo. (1897) *Cours d'Economie Politique* 2, F. Pichou, Paris.

Quandt, Richard E. (1966) “On the Size Distribution of Firms.” *American Economic Review*, 56, pp.416–32.

Rose, Andrew K. (2006) “Cities and Countries.” *Journal of Money, Credit, and Banking,* 38, pp.2225–2245.

Schmid, Markus. M. and Ingo Walter. (2009). “Do Financial Conglomerates Create or Destroy Economic Value?” *Journal of Financial Intermediation*, 18, pp.193-216.

Simon, Herbert A., and Charles P. Bonini, (1958) “The Size Distribution of Business Firms.” *American Economic Review,* 48, pp.607–17.

Stanley, Michael H.R., Sergey V. Buldyrev, Shlomo Havlin, Rosairo N. Mantegna, Michael A. Salinger, and Harry E. Stanley, (1995) “Zipf Plots and the Size Distribution of Firms.” *Economics Letters*, 49, pp.453-457.

Sutton, John, (1997) “Gibrat’s Legacy”, *Journal of Economic Literature,* 35, pp.40-59.

Walter, John R. (2006) "Not Your Father's Credit Union." *Federal Reserve Bank of Richmond, Economic Quarterly*, 92, pp.353-377.

Wheelock, David C. and Paul W. Wilson, (2011) "Are Credit Unions Too Small?" *Review of Economics and Statistics*, 93, pp.1343-1359.

Wheelock, David C. and Paul W. Wilson, (2012) "Do Large Banks have Lower Costs? New Estimates of Returns to Scale for US Banks." *Journal of Money, Credit and Banking*, 44, pp.174-199.

Zipf, George K. (1949) *The Principle of Least Effort*. New York: Hafner.

Table 1 Descriptive statistics: Commercial banks, assets size measure

			Asset size distribution						Log asset size distribution			
Year	No.	Mean	CR_5	CR_{10}	CR_{20}	HHI	Top 2.5%	Top 5%	Mean	S.D.	Skew.	Kurt.
All commercial banks												
1980	14411	301527	0.210	0.300	0.382	125.6	0.672	0.734	11.085	1.190	1.057	6.025
1985	14332	345588	0.166	0.245	0.318	85.7	0.667	0.735	11.211	1.219	1.115	6.038
1990	12281	427828	0.135	0.210	0.282	67.8	0.678	0.751	11.345	1.261	1.213	6.114
1995	9907	598096	0.180	0.250	0.334	96.4	0.712	0.785	11.525	1.295	1.279	6.350
2000	8251	957021	0.282	0.373	0.477	223.9	0.794	0.842	11.678	1.317	1.302	6.919
2005	7469	1354382	0.407	0.486	0.569	406.6	0.817	0.857	11.891	1.340	1.183	6.531
2010	6480	1851634	0.473	0.555	0.664	552.5	0.851	0.882	12.034	1.313	1.275	7.328
Ownership groups: BHC and independent commercial banks												
1980	12394	350599	0.212	0.306	0.400	133.5	0.729	0.789	11.021	1.222	1.307	6.857
1985	11136	444770	0.181	0.269	0.362	110.3	0.759	0.812	11.150	1.244	1.400	7.398
1990	9372	560622	0.162	0.255	0.381	110.3	0.779	0.828	11.279	1.264	1.460	7.605
1995	7662	773340	0.228	0.344	0.482	169.1	0.809	0.850	11.465	1.271	1.448	7.882
2000	6698	1178917	0.361	0.473	0.584	337.2	0.831	0.872	11.695	1.321	1.334	7.292
2005	6347	1593805	0.438	0.526	0.625	473.7	0.838	0.874	11.920	1.357	1.204	6.690
2010	5800	2068722	0.519	0.603	0.699	662.1	0.858	0.888	12.083	1.325	1.236	7.186

Notes to Table 1

Asset size is measured in USD thousands, 2010 prices, conversion using IMF US GDP deflator.

CR_5, CR_{10}, CR_{20} are the five-, ten- and twenty-firm concentration ratios, respectively. HHI is the Herfindahl index.

Top 2.5% and Top 5% are the percentages of industry assets held by banks in the upper 2.5% and 5% of the asset size distribution.

Table 2 Estimation results: Commercial banks, assets size measure

	k (value,	*percentile*)	μ_k	σ_k	Θ_k	α_k	$s(\alpha_k)$	CR_k	K-S
All commercial banks									
1980	609	*95.8*	10.938	0.957	609223	0.898	0.036	0.719	0.059**
1985	605	*95.8*	11.058	0.969	752914	0.841	0.034	0.719	0.064**
1990	520	*95.8*	11.182	0.988	939823	0.767	0.034	0.735	0.073**
1995	611	*93.8*	11.303	0.947	797323	0.752	0.030	0.803	0.081**
2000	325	*96.1*	11.515	1.031	1554390	0.704	0.039	0.827	0.074**
2005	267	*96.4*	11.742	1.081	2113786	0.715	0.044	0.839	0.064**
2010	180	*97.2*	11.910	1.078	2994218	0.670	0.050	0.856	0.071**
Ownership groups: BHC and independent commercial banks									
1980	728	*94.1*	10.820	0.901	425544	0.773	0.029	0.788	0.075**
1985	504	*95.5*	10.977	0.937	638739	0.684	0.030	0.818	0.074**
1990	356	*96.2*	11.122	0.972	950579	0.663	0.035	0.835	0.077**
1995	283	*96.3*	11.312	0.984	1148926	0.656	0.039	0.874	0.079**
2000	260	*96.1*	11.534	1.036	1547326	0.694	0.043	0.868	0.070**
2005	220	*96.5*	11.773	1.098	2290864	0.700	0.047	0.865	0.063**
2010	170	*97.1*	11.954	1.087	2994218	0.685	0.053	0.868	0.067**

Notes to Table 2

k (value) is the fitted threshold rank, below which the log size distribution is truncated normal, and above which the size distribution is Pareto. k (*percentile*) is k (value) expressed as a percentile of the entire size distribution.

μ_k and σ_k are the mean and standard deviation of the fitted truncated normal distribution for log size (asset size measured in USD thousands, 2010 prices, conversion using IMF US GDP deflator).

Θ_k is the fitted threshold asset size (measured in USD thousands).

α_k and $s(\alpha_k)$ are the fitted upper-tail shape parameter in the Pareto distribution, and the standard error of this fitted parameter.

CR_k is the k-firm concentration ratio (k defined as above).

K-S is the Kolmogorov-Smirnov test of the null hypothesis that the entire log size distribution is normal.

** denotes rejection of the null at the 0.01 level; * denotes rejection at the 0.05 level.

Table 3 Summary estimation results: Commercial banks, alternative size measures

	Loans			*Deposits*			*Employees*		
	k (value, *percentile*)		α_k	k (value, *percentile*)		α_k	k (value, *percentile*)		α_k
All commercial banks									
1980	526	*96.4*	0.898	572	*96.0*	0.934	573	*96.0*	0.927
1985	534	*96.3*	0.822	568	*96.0*	0.880	690	*95.2*	0.877
1990	573	*95.3*	0.768	523	*95.7*	0.809	588	*95.2*	0.824
1995	565	*94.3*	0.728	614	*93.8*	0.805	398	*96.0*	0.806
2000	341	*95.9*	0.714	311	*96.2*	0.754	325	*96.1*	0.770
2005	217	*97.1*	0.724	224	*97.0*	0.742	266	*96.4*	0.789
2010	190	*97.1*	0.723	185	*97.1*	0.712	199	*96.9*	0.800
Ownership groups: BHC and independent commercial banks									
1980	615	*95.0*	0.760	723	*94.2*	0.768	529	*95.7*	0.788
1985	625	*94.4*	0.698	601	*94.6*	0.705	547	*95.1*	0.722
1990	398	*95.8*	0.678	520	*94.5*	0.701	529	*94.4*	0.739
1995	390	*94.9*	0.681	267	*96.5*	0.649	264	*96.6*	0.691
2000	247	*96.3*	0.691	265	*96.0*	0.691	186	*97.2*	0.695
2005	195	*96.9*	0.718	201	*96.8*	0.697	192	*97.0*	0.771
2010	173	*97.0*	0.737	182	*96.9*	0.713	130	*97.8*	0.752

Notes to Table 3

k (value) is the fitted threshold rank, below which the log size distribution is truncated normal, and above which the size distribution is Pareto. k (*percentile*) is k (value) expressed as a percentile of the entire size distribution. α_k is the fitted upper-tail shape parameter in the Pareto distribution.

Table 4 Estimation results: Credit unions, assets size measure

Year	Number	Mean	CR_5	CR_{10}	CR_{20}	HHI	μ_0	σ_0	K-S
All credit unions									
1995	11746	36132	0.061	0.088	0.124	19.2	8.830	1.822	0.007
2000	10314	53866	0.067	0.096	0.138	21.9	9.141	1.878	0.008
2005	8691	88071	0.085	0.119	0.165	31.4	9.519	1.939	0.016^{*}
2010	7334	124611	0.108	0.143	0.191	46.5	9.780	2.004	0.018^{*}

Notes to Table 4

μ_0 and σ_0 are the mean and standard deviation of the fitted normal distribution for log size (asset size measured in USD thousands, 2010 prices, conversion using IMF US GDP deflator).
K-S is the Kolmogorov-Smirnov test of the null hypothesis that the entire log size distribution is normal.
** denotes rejection of the null at the 0.01 level; * denotes rejection at the 0.05 level.

Figure 1 Estimation results, upper-tail threshold rank k

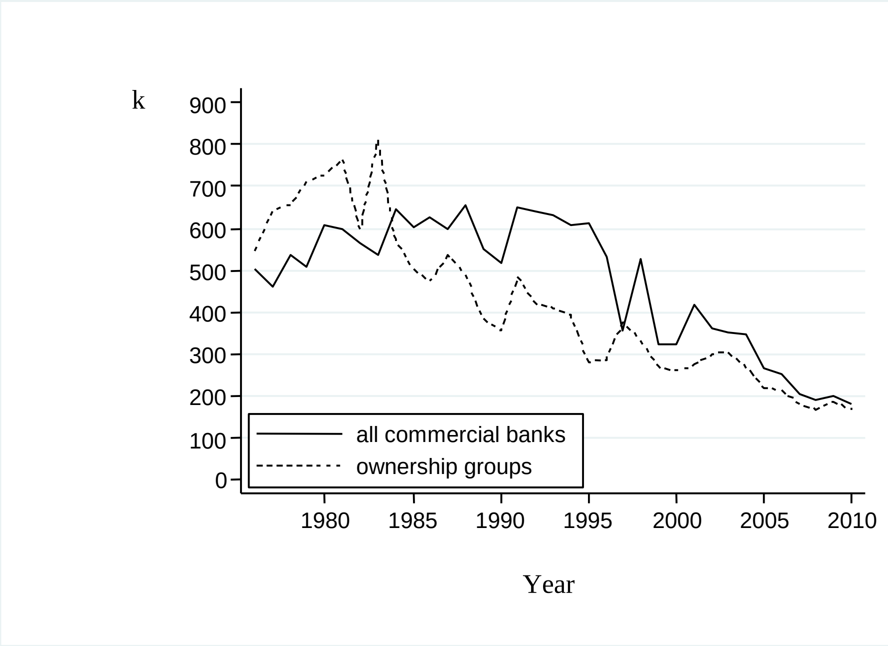

Figure 2 Estimation results, upper-tail shape parameter α_k

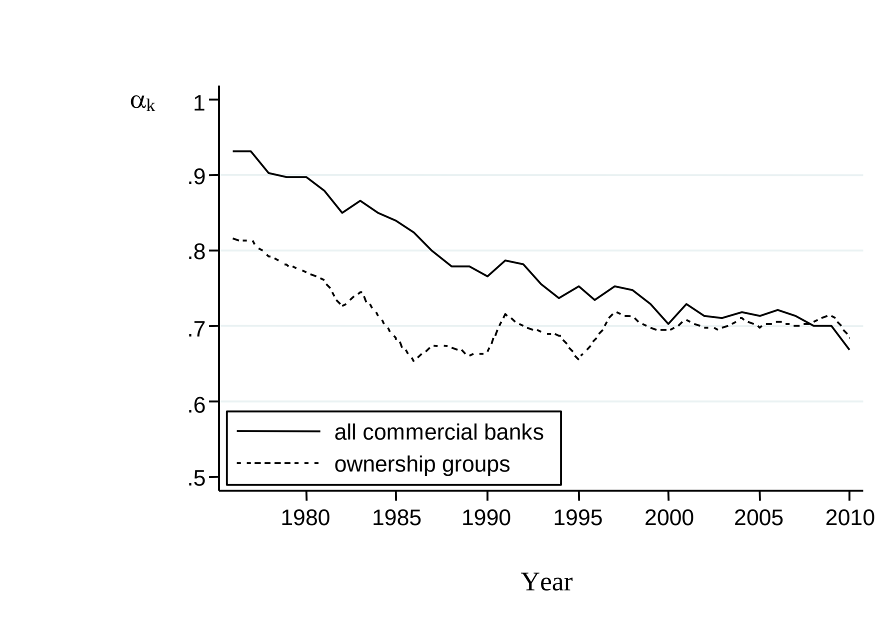

Figure 3 Rank-size plot, assets size measure, all commercial banks

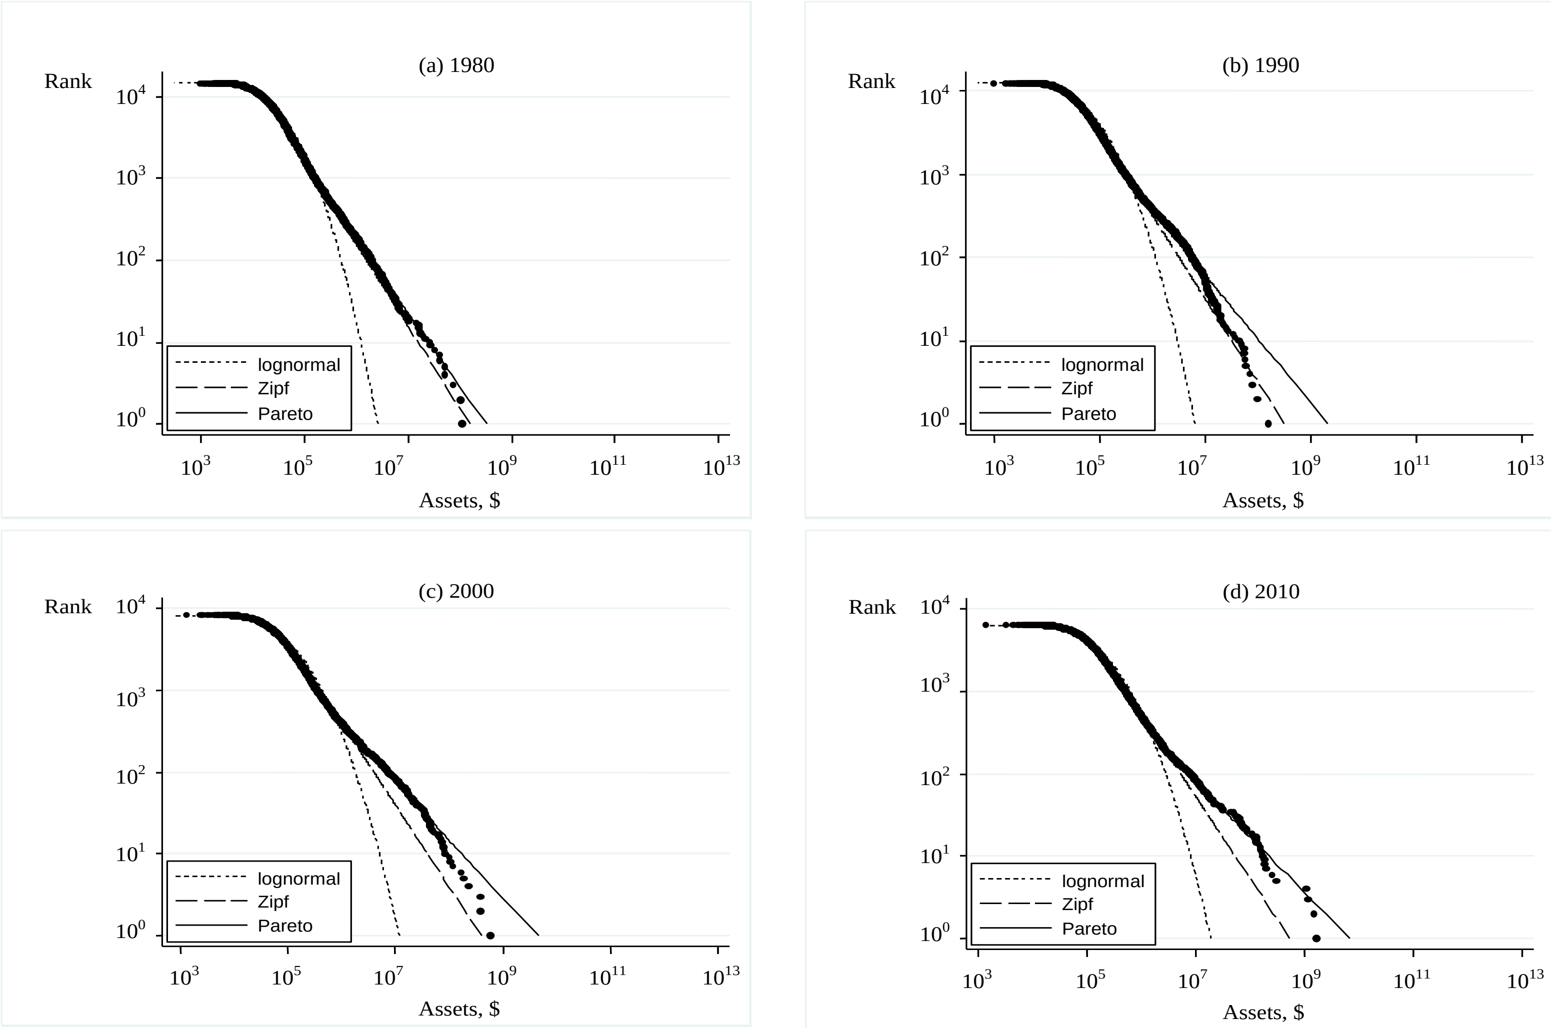

Figure 4 Rank-size plot, assets size measure, credit unions, 2010