Supplemental Appendix: Rapid Growth and the Emergence of Pareto Tails

Yoshiyuki Arata*
Research Institute of Economy, Trade and Industry
National Tax College

Hiroshi Yoshikawa
University of Tokyo

Shingo Okamoto
National Tax College

Date: February 9, 2026

1 Appendix

This appendix examines the robustness of our empirical findings. Section 1.1 investigates how growth rates depend on firm size and age. Section 1.2 analyzes the dependence between growth rates in consecutive periods using copula theory. Section 1.3 reports results obtained when alternative definitions of individual income are used.

1.1 Dependence of Growth Rates on Size and Age

The dependence of growth rates on firm size and age has long been discussed in the literature on firm growth. In particular, it is well established that for firms that exceed a minimum size and are no longer in the initial start-up stage, the independence of growth rates (i.e., Gibrat's law) holds reasonably well (e.g., Lotti et al. (2009); Daunfeldt and Elert (2013)). In this section, we examine the distribution of growth rates across size and age categories to assess the ranges over which this approximation is valid.

Firm sales

We use firm sales growth rates for the period 2010-11. To examine the dependence of growth rates on firm size, we divide the sample into subsamples based on firms' sales in 2010, using size brackets defined by $10^{6.0}, 10^{6.5}, \ldots, 10^{11.0}$ yen. For each subsample, we estimate the probability

*Corresponding author: y.arata0325@gmail.com

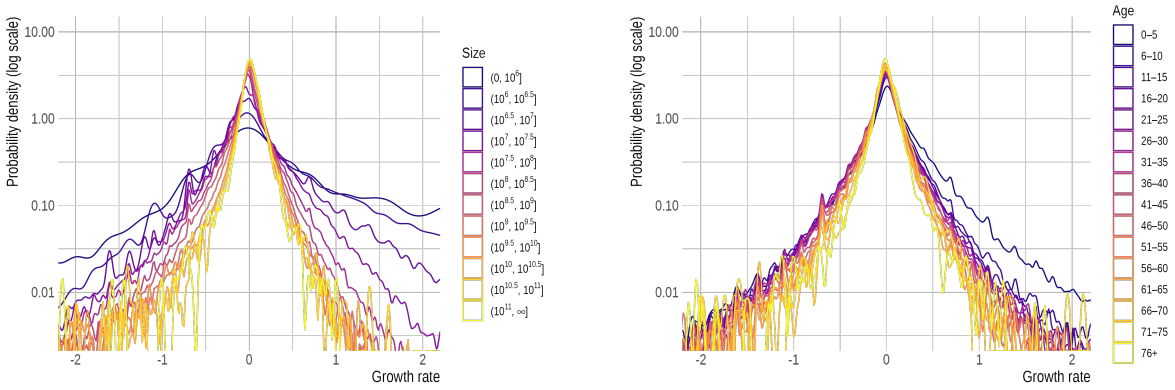

(a) Size group **(b)** Age group

Figure 1: Density estimates of growth rates by size and age groups. We use firm sales growth rates for the period 2010-11. Panel (a) divides the sample into subsamples based on firms' sales in 2010, using size brackets of $10^{6.0}, 10^{6.5}, \ldots, 10^{11.0}$ yen. Panel (b) divides the sample into subsamples based on firms' age in 2010, using five-year intervals.

density function of the growth rate, and the results are presented in **Figure 1**(a). As the figure shows, smaller firms exhibit larger variance and a higher probability of large deviations. By contrast, once firm size becomes sufficiently large, this size dependence weakens, and the estimated density functions for firms with sales of at least $10^{7.5}$ yen almost coincide. This suggests that, for firms with sales above $10^{7.5}$ yen, the size dependence of growth rates is weak, and Gibrat's law provides a better approximation.

Similarly, to examine the dependence of growth rates on firm age, we divide the sample into five-year age groups based on the firms' age in 2010 and estimate the density function of the growth rate for each group. The results are shown in **Figure 1**(b). The growth rate distribution of newly established firms—those five years old or younger—differs markedly from that of the other age groups. The density functions for the remaining age groups are nearly identical. Thus, when the sample is restricted to firms older than five years, the age dependence of growth rates is essentially negligible. These results are consistent with previous studies and suggest that, once firms exceed a certain minimum size and are no longer in their initial start-up phase, growth rates can be regarded as approximately independent of both size and age.

Finally, we confirm that the heavy tails observed in the growth rate distributions—specifically, the fact that they are heavier than exponential—are not the result of aggregating firms of different sizes. In **Figure 2**, we restrict the sample to firms that exceed a certain minimum size and are not newly established (i.e., we exclude firms with sales below $10^{7.5}$ yen and firms five years old or

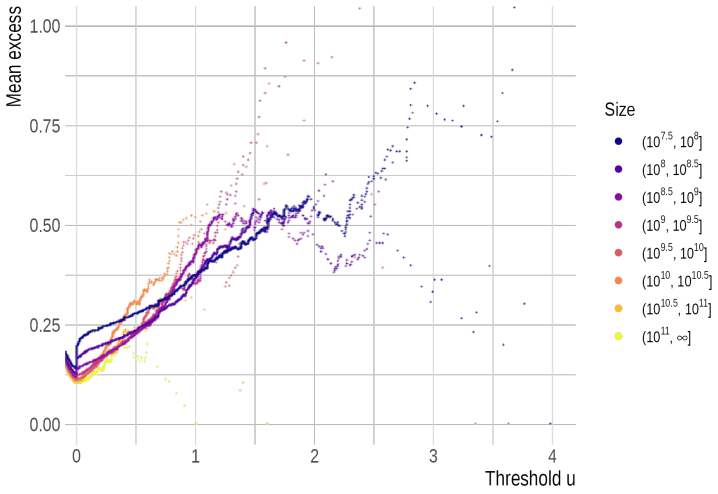

Figure 2: Mean excess function over threshold u by size group. The mean excess function of firms' sales growth rates for 2010-11 is computed by size group, restricting the sample to firms with sales of at least $10^{7.5}$ yen in 2010 and firms older than six years.

younger), and compute the mean excess function of growth rates for each size group. This figure shows that the mean excess function is increasing in the threshold u for all size groups. Thus, even when the sample is partitioned by firm size, the tails of the growth rate distribution remain heavier than exponential. These results indicate that the subexponential property is not driven by aggregation, but reflects an inherent and statistically robust characteristic of the growth rate distribution.

Individual incomes

We apply the same analysis to individual income growth rates for 2014–15. We divide the sample into subsamples based on individual income in 2014, using income brackets of $10^{5.0}, 10^{5.5}, \ldots, 10^{9.5}$ yen, and estimate the probability density function of income growth rates for each subsample. The results shown in **Figure 3**(a) indicate that the density functions for low-income groups exhibit greater variance and heavier tails than those for higher-income groups. However, this dependence weakens as income increases. In particular, for individuals with incomes above $10^{6.5}$ yen, the density functions almost overlap, suggesting that the dependence of the growth rate distribution on income level becomes negligible above this level.

To examine the dependence of the growth rate distribution on age, we divide the sample

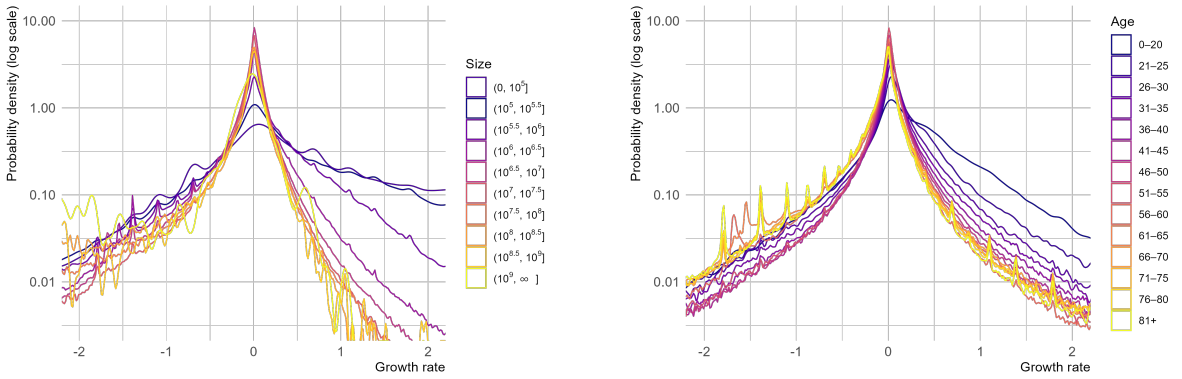

(a) Income group

(b) Age group

Figure 3: Density estimates of growth rates by size and age groups. We use individual income growth rates for the period 2014-15. Panel (a) divides the sample into subsamples based on individual income in 2014, using income brackets of $10^{5.0}, 10^{5.5}, \ldots, 10^{9.0}$ yen. Panel (b) divides the sample into subsamples based on individuals' age in 2014, using five-year intervals.

into five-year age groups based on individuals' age in 2014 and estimate the probability density function of income growth rates for each group. The results in **Figure 3**(b) show that the growth rate distribution for individuals aged 25 or younger differs substantially from that of other age groups, whereas the age dependence becomes significantly weaker for older individuals. Thus, once individuals aged 25 or younger are excluded from the sample, the age dependence of income growth rates is unlikely to pose a significant concern for our analysis. These results support the empirical validity of the independence assumption adopted in our main analysis, as long as the sample is restricted to individuals above certain thresholds of income and age.

Finally, we confirm that the subexponentiality of the growth rate distribution is not the result of aggregating individuals with different income levels. **Figure 4** shows the mean excess function of income growth rates by income group, after excluding individuals with income below $10^{6.5}$ yen and those aged 25 or younger from the sample. The mean excess function is increasing in the threshold u for each income group. That is, for all income groups, the tail of the growth rate distribution is heavier than exponential. These results indicate that the subexponentiality of the growth rate distribution is an inherent property of the distribution itself.

As shown above, in both the firm-sales and individual-income datasets, the distribution of growth rates can be reasonably approximated as independent of size and age, provided that very small or very young firms and individuals are excluded from the sample. Because the region in which this independence holds constitutes the majority of observations in both datasets, the

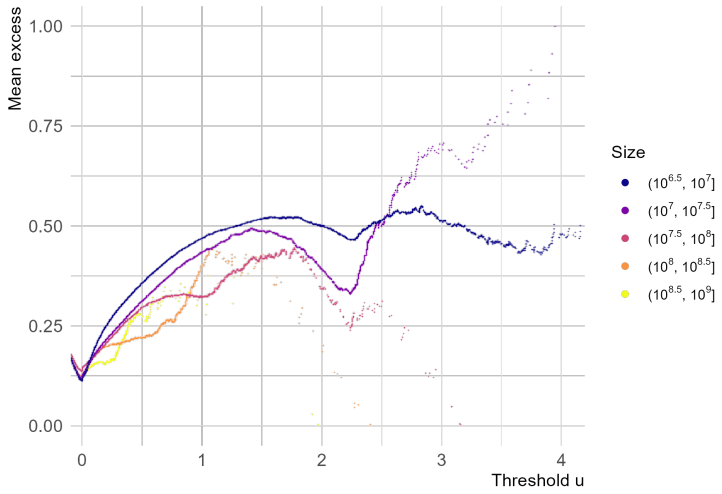

Figure 4: Mean excess function over threshold u by income group. The mean excess function of income growth rates for 2014-15 is computed by income group, restricting the sample to individuals aged between 26 and 60 and with income of at least $10^{6.5}$ yen in 2014.

assumption that growth rates are independent of size and age represents an empirically valid approximation for our analysis. Furthermore, the subexponential property of the growth rate distribution is consistently observed in both firm-sales and individual-income data and remains robust even when each dataset is further partitioned by size. This subexponentiality lies at the core of our theory, and its robustness suggests that the universality of Pareto tails arises from this distributional property.

1.2 Serial Dependence in Growth Rates: A Copula-Based Analysis

In this section, we analyze the dependence structure of growth rates in consecutive periods using copula theory. We begin with a brief introduction to copula theory and then provide empirical results based on copula analysis.

In the existing literature, the most common measure used to quantify the degree of dependence is Pearson's correlation coefficient. However, it is well known in the statistics literature that this coefficient does not purely reflect the dependence between two random variables (see, e.g., Embrechts et al. (2002)). This is because Pearson's correlation also depends on their marginal distributions. For example, if h_1 and h_2 are strictly increasing functions, the Pearson correlation between X_1 and X_2 does not, in general, coincide with the correlation between $h_1(X_1)$ and $h_2(X_2)$.

Such sensitivity to the marginals makes it difficult to determine whether the magnitude of the coefficient reflects genuine dependence or merely properties of the marginal distributions. This issue is particularly important in our setting, where the marginal distribution of growth rates differs substantially from a Gaussian distribution.

These issues can be addressed using copula theory (see Joe (2014) for details). The key idea of copula theory is that any bivariate distribution function F can be uniquely decomposed as

$$F(x_1, x_2) = C(F_1(x_1), F_2(x_2)),$$

where C is the copula function, and F_1 and F_2 are the marginal distributions of X_1 and X_2, respectively. All information on the dependence between X_1 and X_2 is encapsulated in the copula C. The problem with Pearson's correlation coefficient arises from the fact that it is not determined solely by C, but also depends on the marginals F_1 and F_2. As a dependence measure determined only by C, Spearman's ρ_S is commonly used:

$$\rho_S := \mathrm{Corr}[F_1(X_1), F_2(X_2)]$$

Since $F_1(X_1)$ and $F_2(X_2)$ indicate the rank positions of each observation within their respective marginal distributions, Spearman's ρ_S can be interpreted as a measure of rank concordance. For this reason, we use Spearman's ρ_S in the main text instead of Pearson's correlation coefficient.

To obtain a more detailed understanding of the dependence structure, we examine the functional form of C directly. In what follows, we describe two widely used copulas: the Gaussian copula and the Student-t copula. The Gaussian copula and the Student-t copula are derived from the bivariate Gaussian distribution and the bivariate Student-t distribution, respectively. Specifically, the Gaussian copula is defined as

$$C(u_1, u_2; \rho) = \Phi_2\left(\Phi^{-1}(u_1), \Phi^{-1}(u_2); \rho\right)$$

where Φ_2 denotes the bivariate standard Gaussian distribution with correlation parameter ρ, and Φ denotes the univariate standard Gaussian distribution. In other words, the Gaussian copula extracts only the dependence structure of the bivariate Gaussian distribution by removing the effects of its marginal components. The Student-t copula is defined analogously, retaining only the dependence structure of the bivariate Student-t distribution.

The Gaussian copula and the Student-t copula each exhibit distinct characteristic shapes. The density functions of these two copulas are shown in **Figure 5**. The density of the Gaussian copula diverges at the corners $(0, 0)$ and $(1, 1)$. Note that even copulas such as the Gaussian copula—which exhibit weak tail dependence—can produce spikes at $(0, 0)$ and $(1, 1)$. Thus, spikes at these corners in a copula density plot do not necessarily indicate strong tail dependence. In contrast, the Student-t

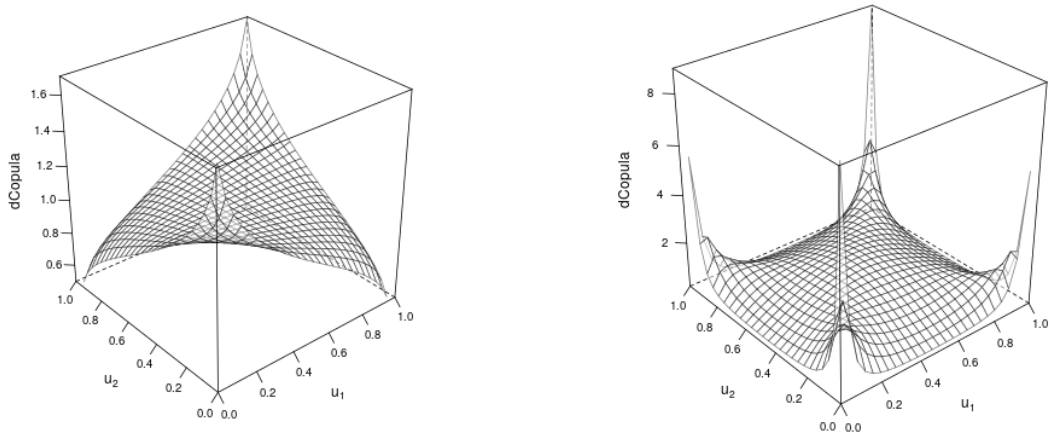

(a) Gaussian copula

(b) Student-t copula

Figure 5: Density plots of the copula function C. For the Gaussian copula, the correlation parameter ρ is set to 0.1. For the Student-t copula, the degrees of freedom are set to $\nu = 1$, and the correlation parameter is set to $\rho = 0.1$.

copula with $\rho > 0$ exhibits spikes not only near $(0, 0)$ and $(1, 1)$ but also near $(0, 1)$ and $(1, 0)$. This implies that, relative to the Gaussian copula, the Student-t copula exhibits a higher likelihood of joint extreme movements in opposite directions—even when the correlation parameter is positive.

While empirical copula plots are commonly used to examine the shape of a copula, the associated density functions can diverge near the corners $(0, 0)$ and $(1, 1)$. To address this issue, Joe (2014) recommends transforming the variables to normal scores. Specifically, each variable X_j is transformed as $Z_j = \Phi^{-1} \circ F_j(X_j)$, and we consider the resulting distribution F_N of (Z_1, Z_2), whose marginals are standard normal:

$$F_N(Z_1, Z_2) := C(\Phi(Z_1), \Phi(Z_2))$$

If the copula C is Gaussian, then F_N reduces to a bivariate Gaussian distribution. Therefore, if the plot of (Z_1, Z_2) departs from the elliptical shape of a bivariate Gaussian distribution, this deviation indicates a departure from the Gaussian copula. **Figure 6** shows the normal-score plots for the Gaussian and Student-t copulas. The Gaussian copula yields an elliptical pattern (i.e., the shape of the bivariate Gaussian distribution), whereas the Student-t copula produces a diamond-shaped pattern that extends toward the four corners, reflecting its stronger tail dependence.

Furthermore, we consider upper and lower semi-correlation coefficients as measures of de-

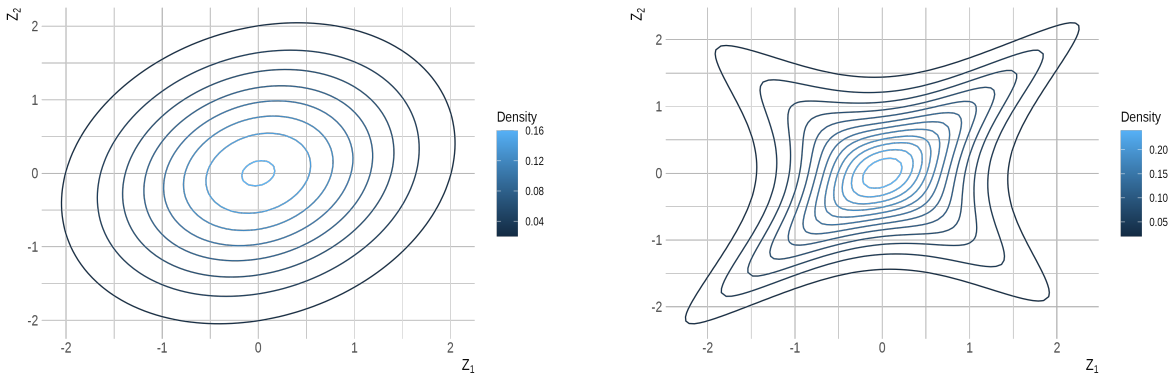

(a) Gaussian copula **(b)** Student-t copula

Figure 6: Density plots of F_N obtained via normal scores. The parameters used here are the same as those employed in **Figure 5**.

pendence based on the transformed variables Z_1 and Z_2:

$$\rho_N^+ = \text{Corr}\left[Z_1, Z_2 \mid Z_1 > 0, Z_2 > 0\right],$$

$$\rho_N^- = \text{Corr}\left[Z_1, Z_2 \mid Z_1 < 0, Z_2 < 0\right]$$

These coefficients correspond to the correlation in the upper-right and lower-left quadrants of the scatter plot of the transformed variables. Since the theoretical values of ρ_N^+ and ρ_N^- under the Gaussian copula can be derived analytically, comparing them with the empirical values allows us to evaluate deviations from the Gaussian copula.

The parameters of the copula are estimated using the pseudo maximum likelihood method proposed by Genest et al. (1995). This method consists of two steps: (1) obtaining nonparametric estimates of the marginal distributions, and (2) estimating the parameters of the copula. To examine whether the Student-t copula provides a statistically significantly better fit to the data than the Gaussian copula, we employ Vuong's method (Vuong (1989)). This method compares the log-likelihoods of the two models to assess whether the difference in fit is statistically significant. In our analysis, we compare the Student-t copula with the Gaussian copula.

Firm sales

We apply the above methods to firm sales data using the growth rates for 2010-11 and 2011-12. **Figure 7** presents the empirical copula and the scatter plot of the normal scores. The empirical copula density exhibits spikes at the four corners—$(1,1)$, $(0,0)$, $(1,0)$, and $(0,1)$—which is a characteristic feature of the Student-t copula. Likewise, the normal-scores plot displays a diamond-

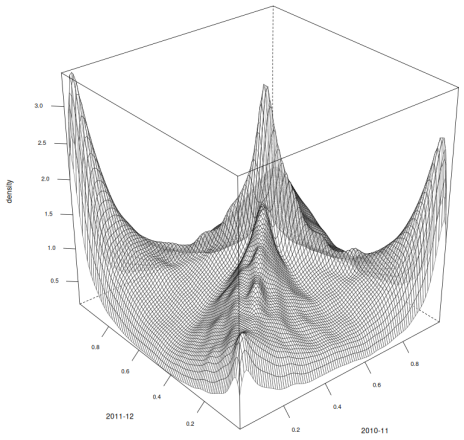

(a) Empirical copula

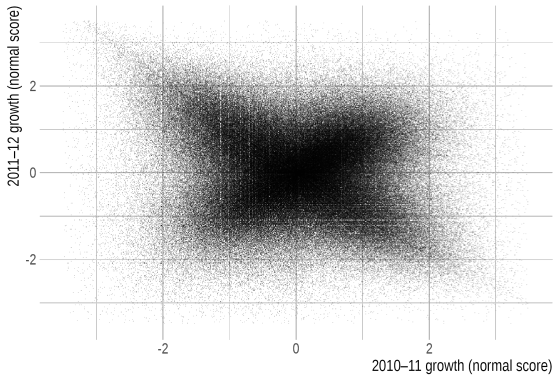

(b) Normal scores

Figure 7: Empirical copula and normal-score plots. Firm sales growth rates for 2010-11 and 2011-12 are used. Panel (a) presents the density estimate of the empirical copula. Panel (b) shows the growth rates transformed into normal scores. The semi-correlation coefficients are $\rho_N^+ = 0.239$ and $\rho_N^- = 0.306$.

shaped pattern rather than an elliptical one. These results suggest that the underlying copula is closer to a Student-t copula than to a Gaussian copula.

This tendency is also reflected in the correlation coefficients. While Spearman's ρ_S is close to zero or slightly negative, as shown in the main text, the semi-correlation coefficients are $\rho_N^+ = 0.239$ and $\rho_N^- = 0.306$. These values are substantially larger than those expected under the Gaussian copula ($\rho_N^+ = \rho_N^- = -0.042$). This indicates that, although there is no strong positive or negative correlation over the entire domain, relatively strong correlations emerge when attention is restricted to the first and third quadrants.

We compare the Gaussian copula and the Student-t copula to determine which provides a better fit to the data. The estimation results for each parametric copula are reported in **Table 1**. Based on the AIC, the Student-t copula with degrees of freedom $\nu = 2$ provides the best fit. We further assess whether the improvement in fit of the Student-t copula over the Gaussian copula is statistically significant using Vuong's method. The test rejects the null hypothesis in favor of the Student-t copula with a p-value of 0.01. This indicates that the Student-t copula provides a significantly better fit to the data than the Gaussian copula.

Finally, we directly examine how well the estimated Student-t copula fits the data. Contour plots of the empirical copula and the estimated copula are presented in **Figure 8**. The Student-t copula provides a close approximation to the empirical copula. We also assess whether the estimated Student-t copula can replicate the observed semi-correlation coefficients of the normal

Parameter estimates and model fit for parametric copulas

Data: Tokyo Shoko Research

copula	shape	maximum log-likelihood	Vuong's test statistic (s.e.)	semi-correlation coefficient
Gaussian	–0.127	4427	-	–0.042
Student-t with df 1	–0.028	-970	-	0.607
Student-t with df 2	–0.070	33496	0.0530(0.00042)	0.375
Student-t with df 3	–0.091	32625	-	0.256
Student-t with df 4	–0.103	29351	-	0.187
Student-t with df 5	–0.110	26392	-	0.142
Student-t with df 10	–0.122	17927	-	0.049

Table 1: Parametric estimates and model fit for parametric copulas. Firm sales growth rates for 2010–11 and 2011–12 are used. Since all copula families considered have the same number of parameters, model selection based on the AIC is equivalent to selection based on the maximum likelihood. The test statistic used in Vuong's method is defined as $\widehat{D} := \log(f^S(y_i;\widehat{\theta}^S)/f^G(y_i;\widehat{\theta}^G))$, where f^S and $\widehat{\theta}^S$ denote the likelihood function and the maximum likelihood estimate for the Student-t copula, and f^G and $\widehat{\theta}^G$ denote those for the Gaussian copula. The last column reports the values of ρ_N^+ and ρ_N^- implied by the estimated copulas.

scores. Under the estimated Student-t copula with degrees of freedom $\nu = 2$, the semi-correlation coefficients are $\rho_N^+ = \rho_N^- = 0.375$. These results further reinforce the finding that the Student-t copula captures the empirical pattern in which stronger dependence emerges when the sample is restricted to the first or third quadrant.

Individual incomes

We apply the above method to individual income growth rates for 2014-15 and 2015-16. **Figure 9** presents the estimated empirical copula and the scatter plot of the normal scores. The empirical copula density exhibits spikes at all four corners, which is characteristic of the Student-t copula. Likewise, the scatter plot of the normal scores forms a diamond shape rather than an ellipse, another feature indicative of the Student-t copula.

Furthermore, this characteristic feature of the Student-t copula is also reflected in the correlation coefficients. While the overall correlation coefficient over the entire domain is close to zero, the semi-correlation coefficients are $\rho_N^+ = 0.392$ and $\rho_N^- = 0.399$. This indicates that, compared with the Gaussian copula, the empirical data exhibit stronger dependence when attention is restricted to the first and third quadrants.

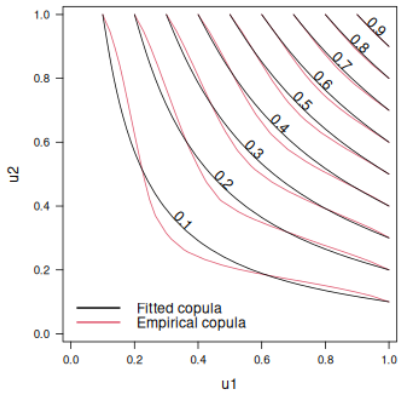

(a) Gaussian copula

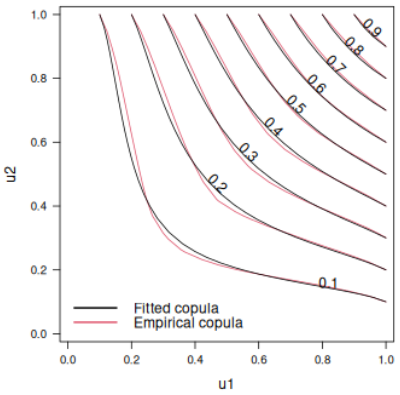

(b) Student-t copula with $\nu = 2$

Figure 8: Comparison between the empirical copula and model-estimated parametric copulas. Firm sales growth rates for 2010-11 and 2011-12 are used. Contour plots of the empirical copula (red lines) and the parametric copulas estimated from the data (black lines) are presented. Panel (a) uses the Gaussian copula, and Panel (b) uses the Student-t copula with $\nu = 2$.

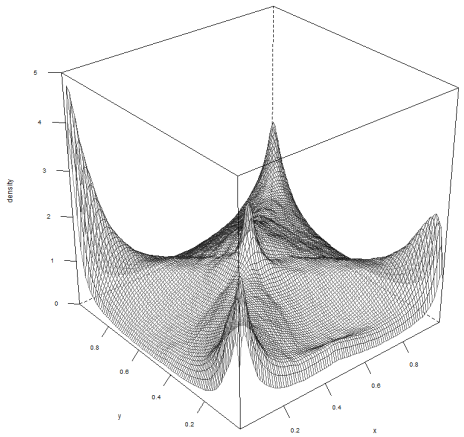

(a) Empirical copula

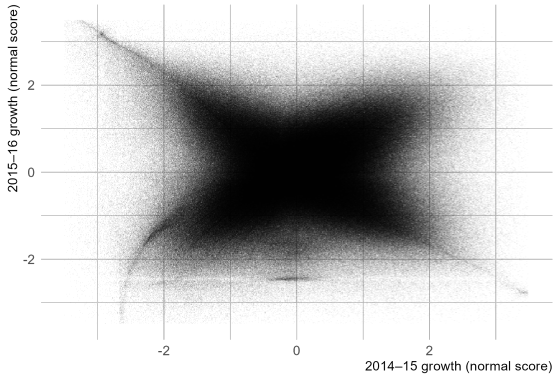

(b) Normal scores

Figure 9: Empirical copula and plots of normal scores. Individual income growth rates for 2014-15 and 2015-16 are used. Panel (a) displays the density estimate of the empirical copula. Panel (b) presents the growth rates transformed into normal scores. The semi-correlation coefficients are $\rho_N^+ = 0.392$ and $\rho_N^- = 0.399$.

Estimates for parametric copulas

Data: Tax return data from National Tax College from 2014 to 2020

copula	shape	maximum log-likelihood	Vuong's test	semi-correlation coefficient
Gaussian	0.013	8	-	0.005
Student-t with df 1	0.051	4277	-	0.618
Student-t with df 2	0.057	8352	0.0834(0.0009)	0.401
Student-t with df 3	0.055	7366	-	0.295
Student-t with df 4	0.052	6321	-	0.232
Student-t with df 5	0.048	5496	-	0.191
Student-t with df 10	0.037	3317	-	0.102

Table 2: Parametric estimates and model fit for parametric copulas. Individual income growth rates for 2014-15 and 2015-16 are used. Due to computational constraints, we do not use the full dataset; instead, we randomly draw $100,000$ observations from the sample to conduct the estimation. For other details, refer to the description in **Table 1**.

The above characteristics of the copula can also be confirmed by the parameter estimates of the Gaussian copula and the Student-t copula. As shown by the estimation results in **Table 2**, the Student-t copula with degrees of freedom $\nu = 2$ provides the best fit to the data according to the AIC. Moreover, Vuong's method confirms that the difference between the Student-t copula with $\nu = 2$ and the Gaussian copula is statistically significant, with a p-value of 0.01.

Finally, to assess how well the estimated copulas fit the data, **Figure 10** presents contour plots of the empirical copula and the estimated copulas. The Student-t copula with degrees of freedom $\nu = 2$ provides a close approximation to the empirical copula. Moreover, under the estimated Student-t copula with $\nu = 2$, the semi-correlation coefficients are $\rho_N^+ = \rho_N^- = 0.401$. These values are close to those computed from the data, $\rho_N^+ = 0.392$ and $\rho_N^- = 0.399$. These results indicate that the Student-t copula captures the dependence structure between growth rates as reflected in the empirical copula.

The above analysis shows that the statistical properties of firm sales growth rates and individual income growth rates are remarkably similar. Specifically, not only are their marginal growth rate distributions and tail dependencies similar, but the copulas summarizing their dependence structures are also closely aligned. In particular, the results indicate that when the current-period growth rate is positive (denoted by $+$), the next period is more likely to exhibit either positive growth $(+,+)$ or negative growth $(+,-)$, whereas cases in which the next period shows approximately zero growth

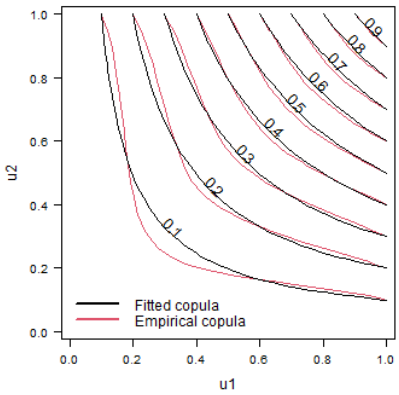

(a) Gaussian copula

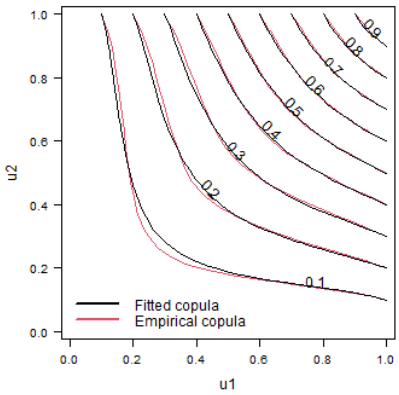

(b) Student-t copula with $\nu = 2$

Figure 10: Comparison between empirical and model-estimated copulas. Individual income growth rates for 2014-15 and 2015-16 are used. Contour plots of the empirical copula (red lines) and the parametric copula estimated from the data (black lines) are presented. Panel (a) uses the Gaussian copula, while Panel (b) uses the Student-t copula with degrees of freedom $\nu = 2$.

$(+, 0)$ are relatively rare. While a full characterization of the dependence structure lies beyond the scope of this paper, the similarity in the statistical properties of firm sales and individual income growth—despite differences in their underlying economic mechanisms—may help explain why macro-level regularities such as Pareto tails emerge in both settings.

1.3 Robustness to Alternative Income Definitions

Here, we verify that the main findings of our analysis remain robust even when alternative definitions of individual income are used. We repeat the same analysis using two alternative income definitions: net income (after deducting expenses and tax deductions) and wage and salary income. Using these alternative definitions, we find that our main results are replicated:

- Pareto tails emerge in the size distribution by age;
- The share of the tail probability remains nearly constant over a wide range of x, and this pattern holds across all age groups;
- The growth rate distribution exhibits heavier tails than an exponential function;
- Apart from jump episodes, the distribution of period-by-period growth rates, conditional on cumulative growth over n periods, differs little between individuals with high cumulative growth and those without.

We conduct the analysis by defining individual income as net income, i.e., income after

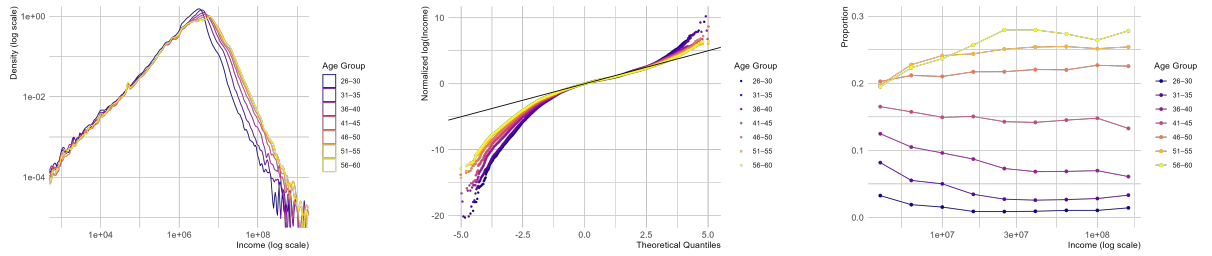

(a) Size distributions **(b)** QQ-plots **(c)** Age-group proportions

Figure 11: Size distributions by age group. We use net income as the measure of individual income. We restrict the sample to individuals aged 26 to 60 and divide them into five-year age groups. Panel (a) presents the estimated density functions of size S for each age group. Panel (b) shows the QQ-plots of the size distributions for each age group. Panel (c) reports the proportion of each age group in the tail probability $\mathbb{P}(S > x)$.

deducting expenses and tax deductions. Results using wage and salary income as an alternative definition are provided in **Figure 13** and **Figure 14**, and the interpretation is analogous to that for the net-income case. **Figure 11**(a) presents the estimated densities of the size distributions across age groups. All age groups exhibit a Pareto tail, and the tail slopes are similar across groups. **Figure 11**(b) shows the QQ-plots of these age-specific size distributions. Consistent with the density estimates, the distributions deviate from the Gaussian distribution, indicating heavy-tailed behavior. **Figure 11**(c) reports the share of each age group in the tail probability $\mathbb{P}(S > x)$. These shares remain stable as x increases, implying that the upper tail is not disproportionately driven by older age groups.

Next, we present the results for growth rates. **Figure 12**(a) shows the density estimates of the n-year growth rate distributions for $n = 1, 2, \ldots, 6$. In all cases, the distributions deviate clearly from the Gaussian distribution. **Figure 12**(b) presents the mean excess function $e(u)$ for the n-year growth rates. Consistent with the density estimates, $e(u)$ is an increasing function of u, indicating that the growth rate distributions have heavier tails than an exponential function. Moreover, the shape of $e(u)$ is largely invariant to n, exhibiting a similar pattern across all horizons, which implies that the distributions are subexponential.

Finally, to examine the characteristic patterns of growth, we compare the histograms of the minimum growth rate within the n-year period for individuals who achieved $\widetilde{S}_n > x$ and those who did not. For $n = 6$, the 97th percentile of $\widetilde{S}_n$ is 0.814. If the distribution were Gaussian, the distribution of growth rates for those who achieved $\widetilde{S}_n > x$ (i.e., the conditional distribution) would be expected to shift to the right by approximately $0.814/6 = 0.136$. The results are shown

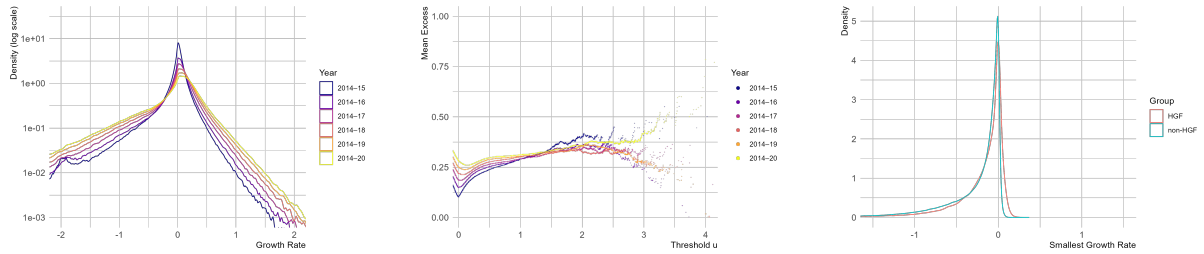

(a) Growth rate distribution **(b)** Mean excess function **(c)** Minimum growth rate

Figure 12: Results based on n-year growth rates We use net income as the measure of individual income. In computing growth rates, we restrict the sample to individuals aged 25 to 60. Panel (a) reports the density estimates of the n-year growth rates ($n = 1, 2, \ldots, 6$), computed using 2014 as the starting year. The vertical axis is shown on a logarithmic scale. Panel (b) presents the mean excess function for the n-year growth rates. Panel (c) compares the histograms of the minimum growth rate within the n-year period for individuals who achieved $\widetilde{S}_n > x$ and those who did not. The 97th percentile of $\widetilde{S}_n$ is 0.814. The medians of the two histograms are -0.0867 and -0.121, respectively.

in **Figure 12**(c). The two histograms are similar: their medians are -0.0867 and -0.121, with a difference of 0.0343. This difference is much smaller than the 0.136 gap expected under a Gaussian distribution, providing evidence in favor of jump-driven growth.

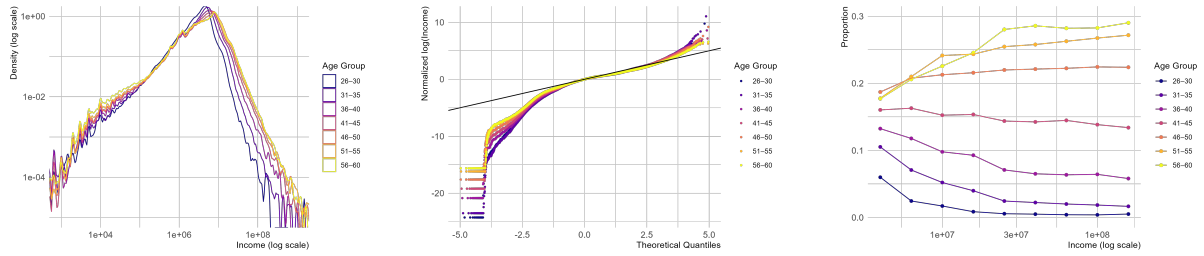

(a) Size distributions **(b)** QQ-plots **(c)** Age-group proportions

Figure 13: Size distributions by age group. We use wage and salary income as the measure of individual income. We restrict the sample to individuals aged 26 to 60 and divide them into five-year age groups. Panel (a) presents the estimated density functions of size S for each age group. Panel (b) shows the QQ-plots of the size distributions for each age group. Panel (c) reports the proportion of each age group in the tail probability $\mathbb{P}(S > x)$.

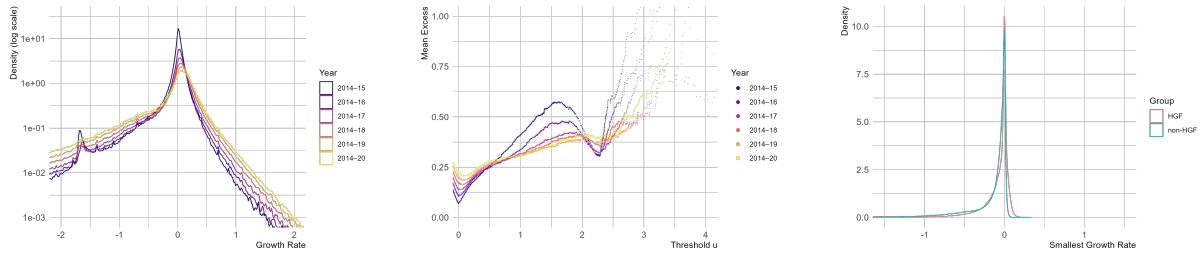

(a) Growth rate distribution **(b)** Mean excess function **(c)** Minimum growth rate

Figure 14: Results based on n-year growth rates We use wage and salary income as the measure of individual income. In computing growth rates, we restrict the sample to individuals aged 26 to 60. Panel (a) reports the density estimates of the n-year growth rates ($n = 1, 2, \ldots, 6$), computed using 2014 as the starting year. The vertical axis is shown on a logarithmic scale. Panel (b) presents the mean excess function for the n-year growth rates. Panel (c) compares the histograms of the minimum growth rate within the n-year period for individuals who achieved $\widetilde{S}_n > x$ and those who did not. The 97th percentile of $\widetilde{S}_n$ is 0.652. The medians of the two histograms are -0.0237 and -0.0559, respectively.

References

Daunfeldt, S.-O. and Elert, N. (2013). When is Gibrat's law a law? *Small Business Economics*, 41:133–147.

Embrechts, P., McNeil, A., and Straumann, D. (2002). Correlation and dependence in risk management: properties and pitfalls. *Risk management: value at risk and beyond*, 1:176–223.

Genest, C., Ghoudi, K., and Rivest, L.-P. (1995). A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3):543–552.

Joe, H. (2014). *Dependence modeling with copulas*. CRC press.

Lotti, F., Santarelli, E., and Vivarelli, M. (2009). Defending Gibrat's Law as a long-run regularity. *Small Business Economics*, 32(1):31–44.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica*, pages 307–333.