

Why is Consumption More Log Normal Than Income? Gibrat's Law Revisited*

Erich Battistin
University of Padova and IRVAPP

Richard Blundell
University College London and IFS

Arthur Lewbel
Boston College

This version: August 2009

Abstract

Significant departures from log normality are observed in income data, in violation of Gibrat's law. We show empirically that the distribution of consumption expenditures across households is, within cohorts, closer to log normal than the distribution of income. We explain this empirical result by showing that the logic of Gibrat's law applies not to total income, but to permanent income and to marginal utility.

Key Words: Consumption, Gibrat Law, Income, Inequality, Lognormal.

JEL Classification: D3, D12, D91

1 Introduction

The traditional parametrization of the income distribution is log normal with a thick, Pareto upper tail. The classic explanation for log normality of income is Gibrat's (1931) law, which essentially models income as an accumulation of random multiplicative shocks. In this paper we confirm that the income distribution in countries including the United States and the United Kingdom has a shape that is close to, but not quite, log normal. We then show that the distribution of consumption is much closer to log normal than income.

This yields two puzzles: why are both consumption and income approximately log normal, and why, within cohorts, is consumption much closer to log normal than income? We show that standard models of consumption and income evolution can explain both puzzles. In particular, the usual decomposition of an individual's income evolution process into permanent and transitory components is shown to imply that Gibrat's law applies to permanent income rather than total income. Similarly, standard

*Original draft October 1999. This paper benefited from constructive comments by the editor and three anonymous referees. Funding for this research was provided by the ESRC Centre for the Analysis of Public Policy at the IFS. Data from the FES made available from the CSO through the ESRC data archive has been used by permission of the Controller of HMSO. We are responsible for all errors and interpretations.

Euler equation models make Gibrat’s law apply to marginal utility and hence to consumption. The result is that the consumption distribution is closer to log normal than the income distribution within cohorts, and observed departures from log normality in the income distribution are attributable to non log normality of the distribution of transitory income shocks across households.

Having Gibrat’s law apply to consumption within cohorts has a number of implications for welfare and inequality measurement, aggregation and econometric modeling. See the Battistin, Blundell and Lewbel (2007) working paper for further discussion of these implications. This working paper also includes results for many US Consumer Expenditure Survey household cohorts and ages in addition to those presented later, and includes similar analyses using the British Family Expenditure survey data, and using different household size adjustments, all of which further support the findings we present here.

In the next section we show why the logic behind Gibrat’s law applies to permanent income rather than total income. In Section 3 we show how standard Euler equation models of consumption also yield Gibrat’s law. The remainder of the paper is then devoted to an empirical analysis of the distributions of income and consumption by cohort based on multiple surveys of United States data.

2 The Income Process and Log Normality

For an individual that has been earning an income for τ years, let y_τ and y_τ^p be the individual’s log income and log permanent income, respectively, so $y_\tau = y_\tau^p + u_\tau$, where u_τ is defined as the transitory shock in log income and thus independent of the permanent component. Permanent income evolves as $y_\tau^p = y_{\tau-1}^p + \eta_\tau$, where η_τ is the shock to permanent income and η_1 is permanent income in the initial time period. In the above definitions it is assumed that the annuitized contributions of transitory income to future permanent income have been removed from u_τ and included in η_τ . For example, all shocks to income in the final year of a person’s life would be permanent shocks. This formalization of Friedman’s (1957) decomposition of current income into permanent and transitory components is a common model of income behavior (see, e.g., Blundell and Preston, 1998). The permanent income model implies that $y_\tau^p/\tau = \frac{1}{\tau} \sum_{s=1}^{\tau} \eta_s$, where τ is the number of time periods that the person has been earning an income, or more formally the number of periods for the income process.

Since y_τ^p/τ is a simple average of random shocks, by application of a central limit theorem (CLT) assuming standard regularity conditions (e.g., shocks η_τ that satisfy a mixing process and have moments higher than two) there exist moments μ_p and σ_p^2 such that $\tau^{1/2} (y_\tau^p/\tau - \mu_p) \rightarrow N(0, \sigma_p^2)$, so that $y_\tau^p \sim N(\tau\mu_p, \tau\sigma_p^2)$ for large τ . Therefore, the standard income generation model implies that permanent income should be close to log normally distributed, at least for individuals that are old enough to

have experienced a moderate number of permanent income shocks. In particular, if permanent income were observable, the model would imply that the distribution of permanent income across individuals in the same (working) age cohort should be close to log normal.

The CLT also immediately implies Deaton and Paxson's (1994) result that the dispersion of income within cohorts increases with the age of the cohort. This follows since $Var(y_\tau) = \tau\sigma_p^2 + Var(u_\tau)$, which grows with τ . Our derivation here shows that not only does the standard model make dispersion of log income increase with age as Deaton and Paxson (1994) observe, but that the distribution becomes more normal as well. In fact, the observation that Gibrat's law implies a growing second moment was noted as early as Kalecki (1945).

Gibrat's original law assumed that income is determined by the accumulation of a series of proportional shocks. We have shown here that the standard permanent income model implies that it is permanent income, not total income, that is determined by an accumulation of shocks, and therefore that Gibrat's law should hold for permanent income, but not necessarily for total income. If the transitory shocks u_τ are small relative to y_τ^p then log total income will also be approximately normal, but unless transitory shocks are themselves normally distributed, log permanent income will be closer to normal than log total income. In particular, if transitory shocks have an appropriately skewed distribution (perhaps through some combination of overtime and temporary layoffs, or occasional large wealth shocks such as bequest receipts) then the total income distribution can take the classic empirical form of log normal with a Pareto upper tail.

3 Euler Equations and Log Normality of Consumption

An individual's permanent income is not directly observable. In this section we show that intertemporal utility maximization implies a similar structure for consumption, resulting from the cumulation of random shocks to income and other variables that affect utility. Traditional models of consumer behavior going at least as far back as Friedman (1957) assume that consumption is at least approximately equal to permanent income, and so the results of the previous section directly imply normality of log consumption in traditional models. In this section we obtain a similar result directly from consumption Euler equations.

Let c_τ be an individual's real consumption at age τ , and let x_τ be a vector of real income I_τ and other variables that affect utility.¹ Assume that in each time period τ the individual maximizes the

¹These other variables could include lagged consumption to permit habit effects, as well as prices, wages, demographic characteristics and stocks of durables.

expectation of the present discounted value of a time separable utility function:

$$u(c_\tau, x_\tau) + \sum_{s=\tau+1}^T \delta_{\tau+1} \dots \delta_s u(c_s, x_s),$$

subject to the expectation of the intertemporal budget constraint:

$$c_\tau - I_\tau + \sum_{s=\tau+1}^T R_\tau \dots R_s (c_s - I_s) = w_\tau,$$

where δ_τ is the individual's subjective discount rate at age τ , R_τ is the market discount rate when the individual is aged τ , and w_τ is accumulated wealth at age τ (which can include a desired bequest, appropriately time discounted). Budget constrained maximization of this utility function yields the standard Euler equation model for consumption (see Deaton, 1992), which is:

$$\phi(c_\tau, x_\tau) = b_\tau \phi(c_{\tau-1}, x_{\tau-1}) + e_\tau^*.$$

Here $\phi_\tau \equiv \phi(c_\tau, x_\tau) \equiv \partial u(c_\tau, x_\tau) / \partial c_\tau$ is the marginal utility of consumption and e_τ^* is the shock to consumption resulting from new information at age τ . Following Hall (1978), this new information could just be shocks to income I_τ , but could also include new information regarding interest rates R_τ and $b_\tau = \delta_\tau / R_\tau$. By defining $e_1^* = \phi_1$, $\varepsilon_{\tau\tau} = e_\tau^*$ and $\varepsilon_{\tau s} = b_\tau b_{\tau-1} \dots b_{s+1} e_s^*$, for $s = 0, \dots, \tau - 1$, there is $\phi(c_\tau, x_\tau) = \sum_{s=1}^{\tau} \varepsilon_{\tau s}$.

Assuming that the $\varepsilon_{\tau s}$ terms satisfy the conditions required for a triangular array CLT, there exist moments μ_ϕ and σ_ϕ^2 such that: $\phi(c_\tau, x_\tau) \sim N(\tau\mu_\phi, \tau\sigma_\phi^2)$ for large τ . There are many alternative regularity conditions that will yield a CLT here (see, e.g., Wooldridge and White 1988). This derivation shows that marginal utility ϕ should be close to normal, so if $\phi(c, x)$ is approximately log-linear in c , then logged consumption will also be close to normal. This derivation allows the risk free rate to be time varying, and also permits some dependence in the Euler equation errors, as would arise if individuals are sometimes liquidity constrained.

Though this derivation delivers asymptotic normality of the marginal utility of consumption, it does not imply in general that consumption itself is log normally distributed. Thus, it is worth considering conditions that are sufficient for exact asymptotic log normality of consumption data. One set of sufficient conditions is to assume that $\delta_\tau = R_\tau$, the shocks e_τ^* are independently distributed with finite moments higher than two, and the utility function given above has $u(c_\tau, x_\tau) = \alpha c_\tau + \beta c_\tau \ln(c_\tau) + \gamma(x_\tau)$ for some function γ and constants α and β . This then makes $b_\tau = 1$ and marginal utility $\phi_\tau = (\alpha + \beta) + \beta \ln(c_\tau)$, so the Euler equation yields the sample average $\beta \ln(c_\tau) / \tau = \sum_{s=1}^{\tau} e_\tau^* / \tau$, which is asymptotically normal at rate root τ by the Lindeberg Feller CLT.

4 Detecting Departures from Log Normality

We examine the closeness of observed data to log normality by comparing different features of the empirical distributions of log income and log expenditures to their theoretical normal counterparts. To visually depict departures from normality we construct quantile-quantile (QQ) plots as well as histograms of the sample, overlaid with a $N(\mu, \sigma^2)$ density function.

To construct graphical comparisons or formal test statistics for normality requires estimation of the location and scale parameters μ and σ . Standard estimates of these and higher moments can be very sensitive to outliers, and both income and consumption data may well contain reporting errors, particularly topcoding, underreporting or misreporting by high and low income households. We therefore use estimates and tests based on robust statistics, which mitigate the impact of gross errors and outliers in the data (see, e.g., Hampel *et al.*, 1986). Consequently, in our application we will use the median $M(Y)$ and the population median absolute deviation $MAD(Y) \equiv M(|Y - M(Y)|)$ as our robust measures of location and scale.² We provide histograms of the data, and superimposed on each histogram is a normal density function that uses these robust mean and variance estimates.

Given location and scale estimates, tests for departure from normality can be implemented. We first construct Kolmogorov-Smirnov tests based on the distance between the empirical distributions of income and expenditure and the corresponding normal distributions. To account for estimation error in $\hat{\mu}$ and $\hat{\sigma}$, we obtained p-values for this test using 10,000 random samples generated under the null hypothesis of normality, $N(\hat{\mu}, \hat{\sigma}^2)$, and counted the number of replicate samples that produced a test statistic greater than or equal to that calculated for the actual data.

We also construct two additional tests based on robust indicators of skewness and kurtosis. Groeneveld and Meeden (1984) suggest skewness measures of the form:

$$\frac{[Q_{1-p}(Y) - M(Y)] - [M(Y) - Q_p(Y)]}{Q_{1-p}(Y) - Q_p(Y)}, \quad (1)$$

where $Q_\alpha(Y)$ is the α -th percentile of the distribution of Y . In our application we use quartile skewness, which takes $p = 0.25$ and is zero for normal distributions. The resulting expression is analogous to estimating skewness by first using the median to center the data and scaling with the interquartile range. Positive (negative) values of this statistic indicate right (left) skewness. Additionally, this coefficient will take values in the interval $(-1, 1)$, with 1 (-1) representing extreme right (left) skewness.

²For normal distributions $M(Y)$ and $MAD(Y)$ are related to the mean and variance by $M(Y) = \mu$ and $MAD(Y) \simeq 0.6745\sigma$ (where the approximation \simeq is just due to the number of decimal places used). The corresponding robust estimators of the location and scale parameters for a normal distribution are $\hat{\mu} = \hat{M}(Y)$ and $\hat{\sigma} = \frac{\hat{MAD}(Y)}{0.6745}$, where $\hat{M}(Y)$ and $\hat{MAD}(Y)$ denote the sample median and sample median absolute deviation.

Analogous to these other moments, for kurtosis we follow Moors (1988) and use:

$$\frac{[O_7(Y) - O_5(Y)] + [O_3(Y) - O_1(Y)]}{O_6(Y) - O_2(Y)}, \quad (2)$$

where $O_\alpha(Y)$ is the α -th octile of the distribution of Y . This statistic is non-negative and robust to the extreme tails of the distribution, and for normal distributions it equals 1.233. We computed the sample analogues of both the skewness coefficient (1) and the kurtosis coefficient (2), and compare them to their theoretical values under the assumption of normality. P-values under the null hypothesis of normality were computed from 10,000 pseudo-samples as before.

5 The Consumption and Income Data

Most of our empirical analysis is based on expenditure and income data from the US Consumer Expenditure (CEX) Interview Survey. We used quarterly expenditures published by the Bureau of Labor Statistics (BLS) between 1980 and 2003 to derive annual aggregate measures of expenditure at the household level.³ For income, we use before tax figures as reported in the fifth interview by households who were classified as complete income reporters. This nominal income and expenditure data are converted to real by deflating using the Consumer Price Index.

We complemented information on income from the CEX with data from the Panel Study of Income Dynamics (PSID). Unlike the CEX, the PSID collects longitudinal annual data on a sample of households followed on a consistent basis since 1968. We examine family disposable income in the PSID for a sample of couples with and without children as described in Blundell, Pistaferri and Preston (2008).

We focus on a sample of married couples (with or without children) and define cohorts based on the year of birth of the head, which we conventionally take to be the husband. The two panels of Table 1 provide the cohort definitions and sample size for the CEX and the PSID samples.

6 The Empirical Distributions of Consumption and Income

Figures 1 and 2 show the distribution in the CEX of log expenditure and log income across the life-cycle for two birth decade cohorts. The left column of figures shows a log real expenditure distribution that is very close to normal. In contrast, the right column of figures shows that log real income for these households is much further from normal with the upper tail skewness that is typical of income

³We used only households who participated in the survey for all interviews (representing about 75-80 percent of the original sample) and sum their quarterly expenditures over the year covered by the four interviews. We considered the measure of total expenditure as published by the BLS after excluding ‘cash contributions’ and ‘personal insurance and pensions’, thus using a definition that includes expenditures for food, alcohol, housing, transportation, apparel, medical care, entertainment, and other miscellaneous items (such as personal care services, reading, education and tobacco products).

distributions, and greater kurtosis as well. A similar pattern holds across all age groups. The log income distributions all present a long lower tail. We expect that at least some of this observed lower tail behavior is due to measurement error, possibly due to underreporting of income at these levels.

In accordance with Kalecki (1945) and Deaton and Paxson (1992), and consistent with Gibrat’s law, the figures show as every birth cohort ages their distributions of income and consumption become more disperse. Also, comparing people aged 41 – 45 in both cohorts shows that the younger cohort has a higher dispersion of income and consumption. A similar pattern holds for other cohorts and age groups.

The departures from log normality of consumption are very small and do not seem to systematically decrease with age, which suggests that by relatively early in one’s working life enough shocks have accumulated to get close to asymptotic normality. However, in even younger cohorts (21 – 25) the distributions are further from log normal than for the older groups, which is again consistent with our inter-temporal consumer theory interpretation of Gibrat’s law.⁴

Our theory suggests that consumption should be closer to log normal than income, because income contains a potentially large transitory component in addition to a log normal permanent income component. This is what we found in the CEX, but one might worry that departures from log normality in CEX income data could be due measurement error, because income may be measured less precisely than consumption in that data set. As a check, in Figure 3 we examine income by birth cohort and age but this time for log family disposable income from the PSID data set, which measures income more carefully than the CEX. We find significant deviations from normality of log income in this data, similar to the departures from log normality found in the CEX.

7 Conclusions

The income distribution has long been known to be approximately log normal. We have shown that the consumption distribution is also close to log normal, and that within demographically homogeneous groups, the distribution of consumption is much closer to log normal than is the distribution of income. We also demonstrate that these empirical regularities are implications of traditional models of the evolution of income and consumption, specifically, that the theory which motivates Gibrat’s law should apply to permanent income and consumption (via Euler equations), rather than to total

⁴Our data includes households with varying numbers of children, because sub-populations sorted by household size would not be comparable across age brackets. For example, households at age 40 with three children are more representative of the general population than households at age 20 that have three children. However, numbers of children correlates with income, and affects the propensity to consume out of current income. So as further check on the robustness of our results, we recalculated distributions after dividing each household’s income and consumption by \sqrt{n} where n is family size, thereby following a common practice of using \sqrt{n} as an equivalence scale. These results remain consistent with our other findings.

income as originally formulated.

We would not expect perfect normality for a variety of reasons. Traditional permanent income and Euler equation models are implausibly simplistic, so we should not expect them to hold exactly. Also, the CLT is an asymptotic property while individuals only have finite lifespans. Even when permanent income is close to log normal for some individuals, their consumption may depart from log normality if marginal utility differs substantially from log consumption, or if liquidity constraints, precautionary savings, or purchases of large durables produce enough dependence in Euler equation innovations to violate the conditions required for a CLT. More generally, normality may not hold for some individuals because their time series of shocks may possess features such as the discount ratios b_τ 's far from one or long memory, that violate the regularity conditions required for a CLT. Despite these possible problems, we find that the observed distributions of consumption and income are broadly consistent with the distribution implications of these models, across cohorts, over time, and across data sets.

Other explanations for the observed consumption and income distributions may exist. For example, if consumption is very badly measured, then its observed distribution could be dominated by measurement errors that happen to be log normal. Another possibility is based on the observation that higher income households tend to consume a smaller fraction of income than lower income households, resulting in a consumption distribution that has a thinner upper tail than the income distribution. If the income distribution is close to log normal except for a thick (Pareto) upper tail, the consumption distribution should then have a thinner upper tail, which could by coincidence be almost the same size as its lower tail, resulting in a near normal distribution. These alternative explanations for consumption log normality require coincidences that we find less plausible than our derivations based on permanent income and Euler equation models, though these alternatives could be contributing factors in the observed distributions.

The finding that Gibrat's law applies to consumption within cohorts has many important implications for welfare and inequality measurement, aggregation, and econometric model analysis, and results in additional regularities in the distributions of related variables. It would be interesting to test if other economic variables that are determined either by Euler equations or decompositions into permanent and transitory components display a similar conformity to Gibrat's law.

References

- [1] Battistin, E. Blundell, R. and A. Lewbel (2007), 'Why is Consumption More Log Normal Than Income? Gibrat's Law Revisited', Working Paper 08/07, Institute for Fiscal Studies, London, and Working Paper 671, Department of Economics, Boston College.

- [2]Blundell, R., L. Pistaferri and I. Preston (2008), ‘Consumption Inequality and Partial Insurance’, *American Economic Review*, 98(5), 1887-1921.
- [3]Blundell, R., and I. Preston (1998), ‘Consumption inequality and income uncertainty’, *Quarterly Journal of Economics* 113, 603-640.
- [4]Deaton, A. (1992), *Understanding Consumption*. Baltimore: John Hopkins University Press.
- [5]Deaton, A., and C. Paxson (1994), ‘Intertemporal choice and inequality’, *Journal of Political Economy*, 102, 384-94.
- [6]Friedman, M. (1957), *A Theory of the Consumption Function*, Princeton: Princeton University Press.
- [7]Gibrat, R. (1931). *Les Inegalites Economiques*. Librairie du Recueil Sirey: Paris.
- [8]Groeneveld, R.A. and Meeden, G. (1984), ‘Measuring Skewness and Kurtosis’, *The Statistician*, Vol. 33, No. 4, 391-399.
- [9]Hall, R. E. (1978), ‘Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence’, *The Journal of Political Economy*, 86, 971-987.
- [10]Hampel, F.R., Ronchetti, E.M., Rousseeuw, P.J., and Stahel, W.A. (1986), *Robust Statistics*, John Wiley and Sons: New York.
- [11]Kalecki, M. (1945), ‘On the Gibrat Distribution’, *Econometrica*, 13, 161-170.
- [12]Moors, J.J.A. (1988), ‘A Quantile Alternative to Kurtosis’, *The Statistician*, Vol. 37, No. 1, 25-32.
- [13]Wooldridge, J. M. and H White, (1988), ‘Some Invariance Principles and Central Limit Theorems for Dependent Heterogeneous Processes’, *Econometric Theory*, 4, 210-230.

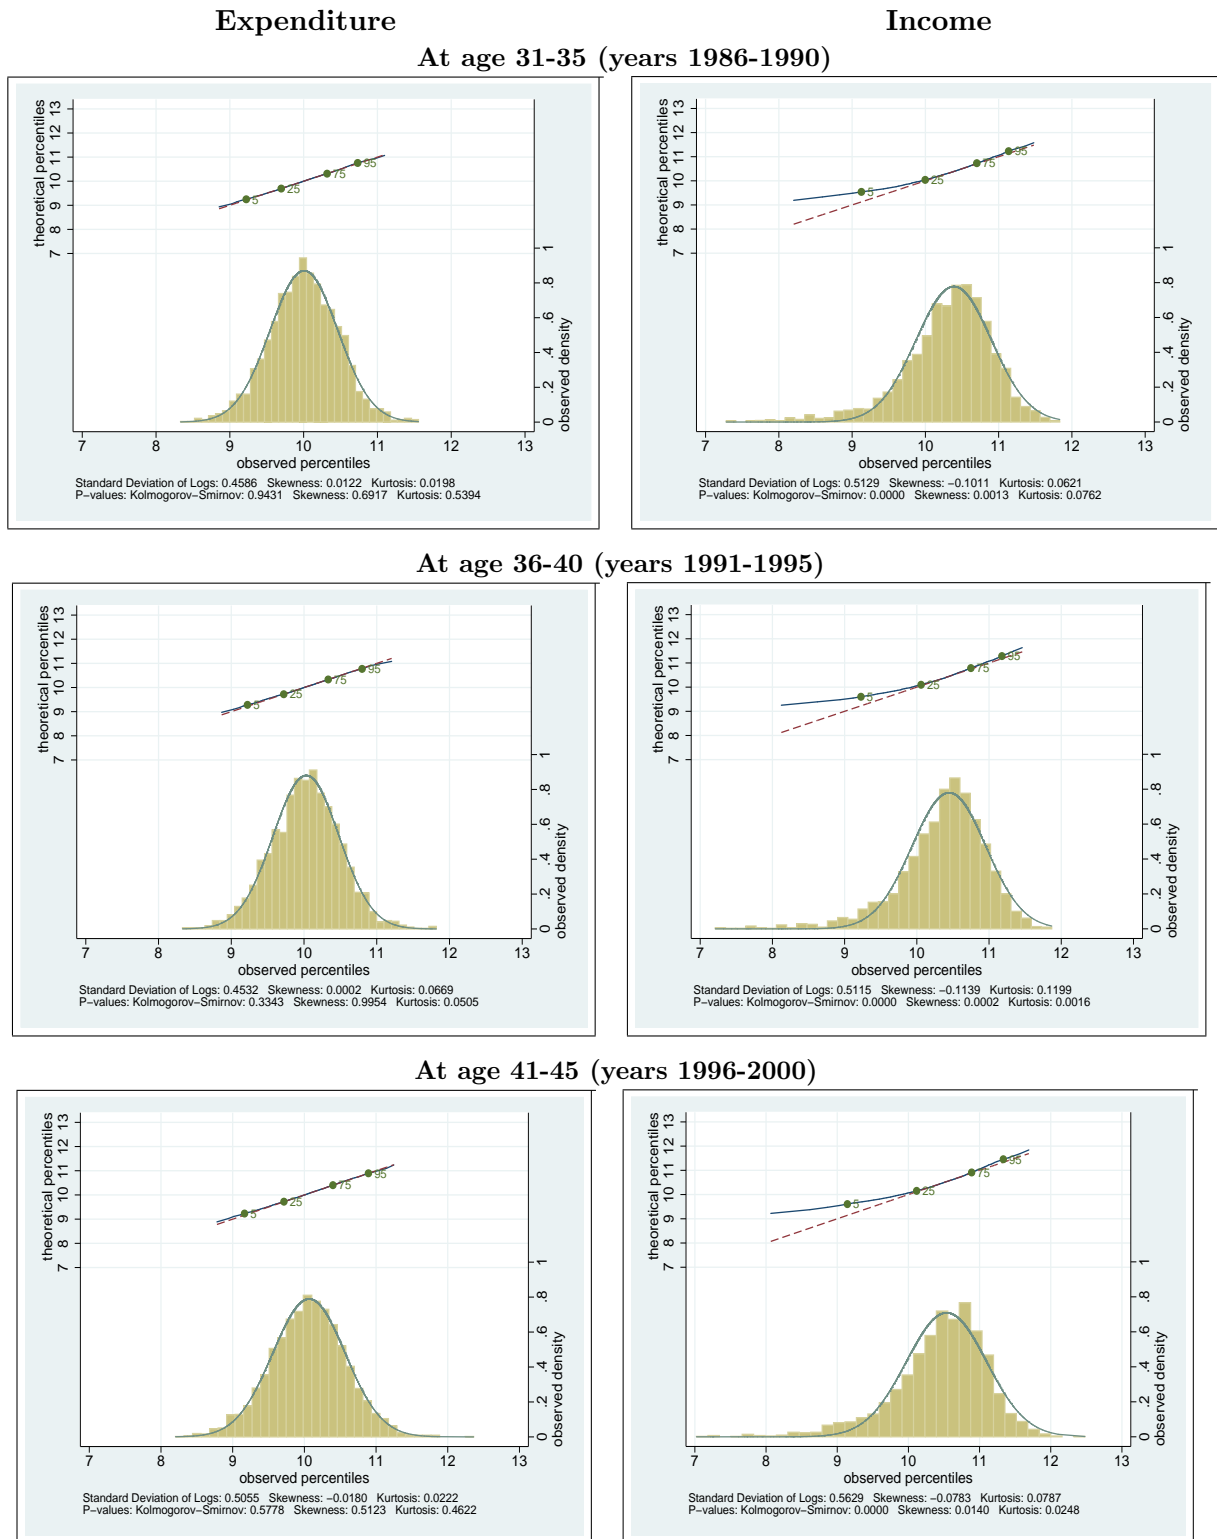
TABLE 1. Sample size by cohort and interview year, separately for data from the Consumer Expenditure Surveys (CEX) and the Panel Study of Income Dynamics (PSID).

CEX Data						
Cohort	Expenditure data			Income data		
	1986-1990	1991-1995	1996-2000	1986-1990	1991-1995	1996-2000
Born in 1960-69	952	1,483	2,802	846	1,279	2,226
Born in 1950-59	2,883	2,641	3,458	2,530	2,193	2,639
Born in 1940-49	2,813	2,192	2,681	2,348	1,746	1,964
Born in 1930-39	2,025	1,476	1,831	1,667	1,177	1,419

PSID Data						
Cohort	Expenditure data			Income data		
	1986-1990	1991-1995	1996-2000	1986-1990	1991-1995	1996-2000
Born in 1960-69						
Born in 1950-59				10,164		
Born in 1940-49				5,642		
Born in 1930-39				3,366		

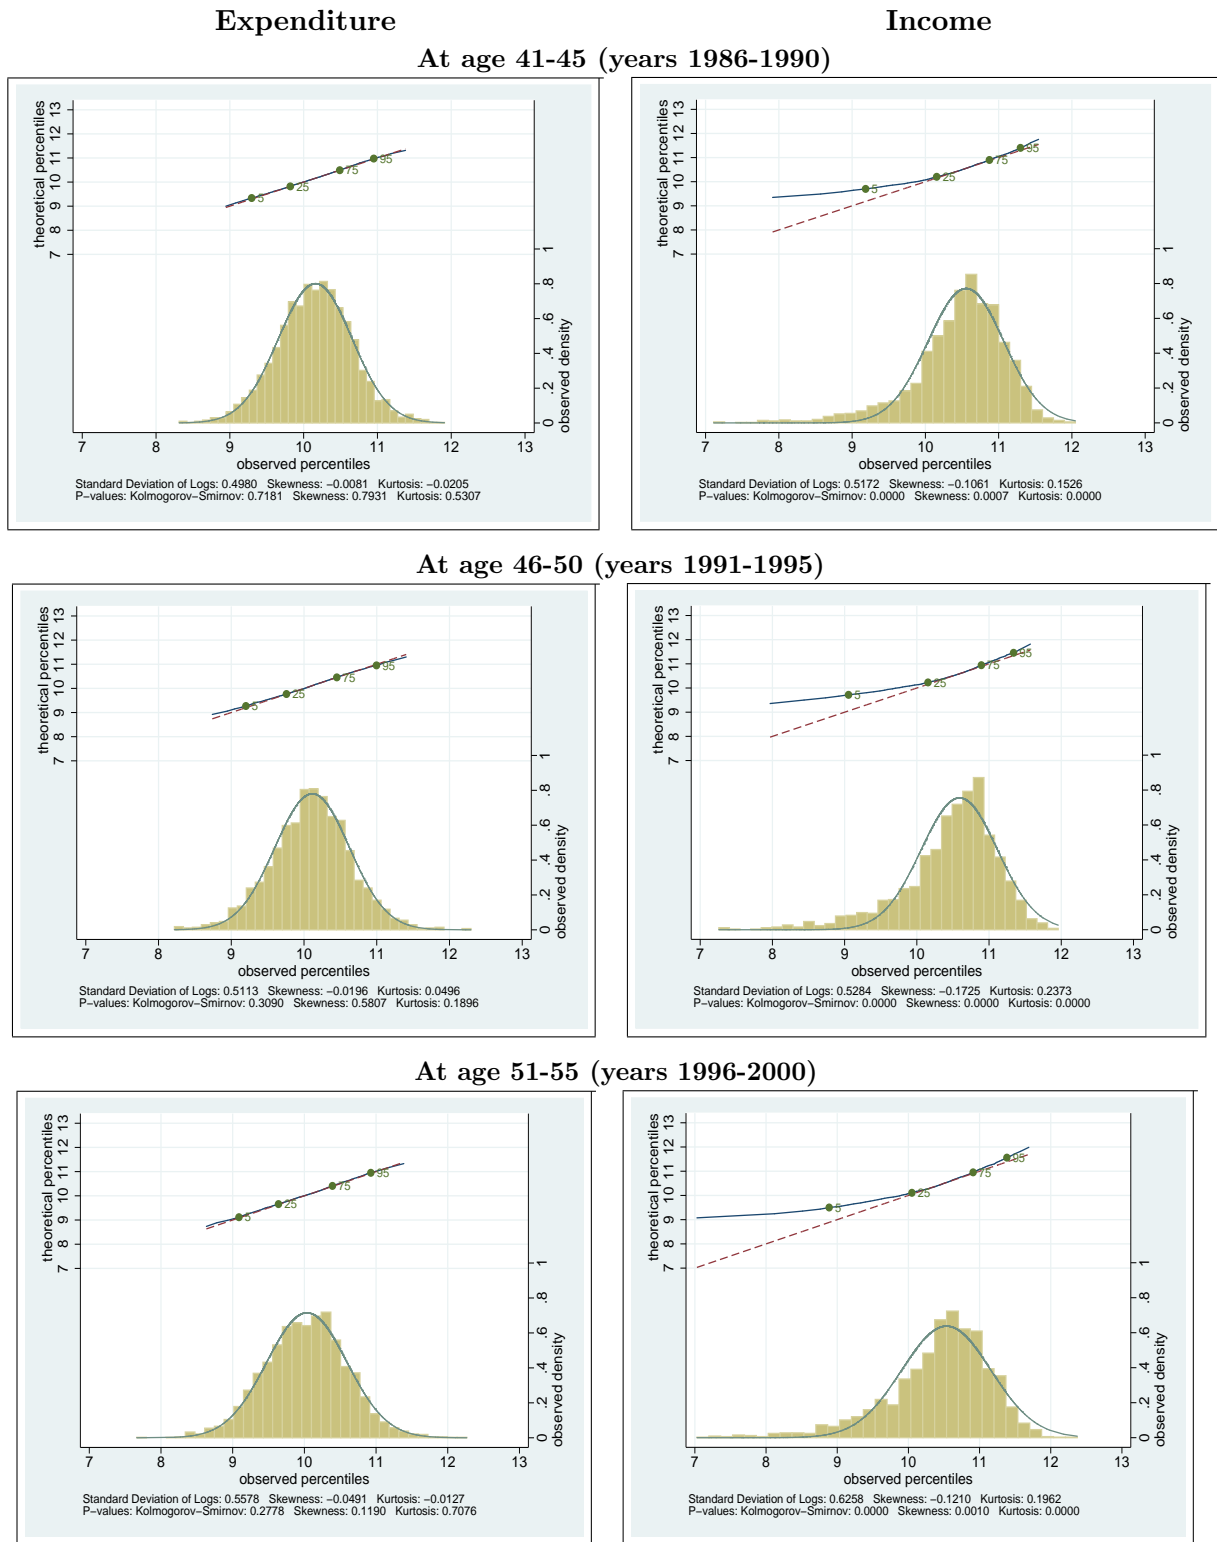
NOTE. Only married couples (with or without children) are considered. The definition of cohorts is based on the year of birth of the head, which we conventionally take to be the husband. CEX data: figures for *total expenditure* (as published by the BLS, excluding “cash contributions” and “personal insurance and pensions”) and *total family income before tax* for complete income reporters in the second interview are considered. Only households who completed all interviews are considered. PSID data: the measure of *income* considered excludes income from financial assets and subtracts federal taxes on nonfinancial income (see Blundell, Pistaferri and Preston, 2008, for further details).

FIGURE 1. Expenditure and income distributions for the 1950-59 cohort (CEX data)



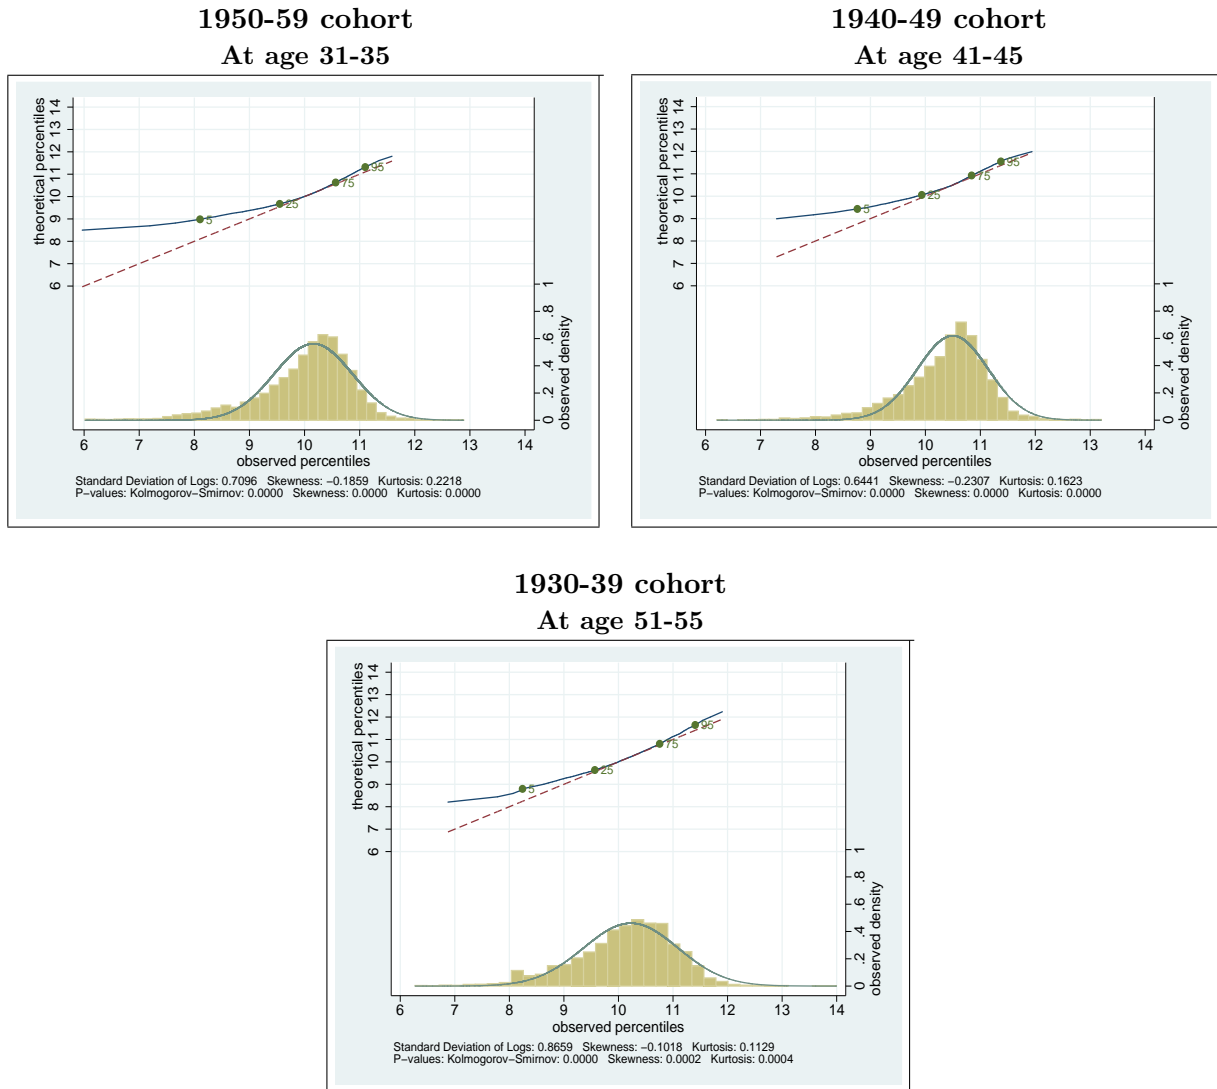
NOTE. The sample size of each panel is given in Table 1. *Total expenditure* excludes “cash contributions” and “personal insurance and pensions”, income refers to *total family income before tax*. All figures are deflated by the CPI. Each panel reports: (a) the histogram of the data with a normal density superimposed calculated at robust mean and variance estimates, (b) the QQ plot of observed vis-à-vis theoretical quantiles under normality (the 5th, 25th, 50th, 75th and 95th percentiles are superimposed), and (c) the Kolmogorov-Smirnov statistic, robust skewness and kurtosis coefficients and the p-value of their difference from theoretical values under normality (see Section 4 for further details).

FIGURE 2. Expenditure and income distributions for the 1940-49 cohort (CEX data)



NOTE. See note to Figure 1.

FIGURE 3. Selected income distributions from PSID data



NOTE. The sample size of each panel is given in Table 1. The measure of *income* considered excludes income from financial assets and subtracts federal taxes on nonfinancial income (as in Blundell, Pistaferri and Preston, 2008). All figures are deflated by the CPI. Each panel reports: (a) the histogram of the data with a normal density superimposed calculated at robust mean and variance estimates, (b) the QQ plot of observed vis-à-vis theoretical quantiles under normality (the 5th, 25th, 50th, 75th and 95th percentiles are superimposed), and (c) the Kolmogorov-Smirnov statistic, robust skewness and kurtosis coefficients and the p-value of their difference from theoretical values under normality (see Section 4 for further details).