



Laboratory of Economics and Management

Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)

Tel. +39-050-883-343 Fax +39-050-883-344

Email: lem@sssup.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

On the distributional properties of household
consumption expenditures. The case of Italy.

Giorgio Fagiolo*
Lucia Alessi*
Matteo Barigozzi*
Marco Capasso*

* Scuola Superiore Sant'Anna, Pisa, Italy

2007/24

November 2007

ISSN (online) 2284-0400

1 Introduction

In the last years, considerable effort has been devoted to the study of the distributional properties of key microeconomic variables and indicators. For example, a huge amount of contributions has explored the statistical properties of wealth and personal income distributions, both across years and countries (see, e.g., Chatterjee, Yarlagadda, and Chakrabarti, 2005, and references therein). Similarly, in the field of industrial dynamics, a large body of literature has successfully characterized the shape of cross-section firm size and growth-rate distributions, and their evolution over time (cf. among others Axtell, 2001; Bottazzi and Secchi, 2006a).

These studies show that, despite the turbulence typically detected at the microeconomic level (e.g., entry and exit of firms; positive and negative persistent shocks to personal income; etc.), there exists an incredible high level of regularity in the shape of microeconomic cross-section distributions, both across years and countries. For instance, personal income distributions appear to be characterized by a log-normal body with a Pareto upper tail in the majority of cases (Clementi and Gallegati, 2005a; Souma, 2001). Furthermore, as far as growth-rate distributions are concerned, there seems to emerge a sort of universality feature: the same family of distributions¹ is indeed able to fit growth rates for firms in different sectors, industries and even countries (both cross-sectionally and along the time-series dimension; cf. Lee et al. 1998; and Fagiolo, Napoletano and Roventini 2007) .

Notwithstanding such successful results, the above line of research has not been extensively applied, so far, to other key microeconomic variables for which detailed cross-section data are available, namely household consumption expenditures (HCEs). This is somewhat surprising for two related reasons (Attanasio, 1999). First, understanding consumption is crucial to both micro- and macro-economists, as it accounts for about two thirds of GDP and it decisively determines social welfare. Second, while we know a lot about the statistical properties of aggregate consumption time-series and microeconomic

¹That is, the exponential-power family of densities, originally introduced by Subbotin (1923). More on this below.

life-cycle profiles, our knowledge about the distributional properties of cross-section HCEs is rather poor, as we almost always limit ourselves to the first and second moments thereof.

The only exception to this trend is a recent contribution by Battistin, Blundell, and Lewbel (2007). They employ expenditure and income data from U.K. and U.S. surveys and show that HCE distributions are, within cohorts, well approximated by log-normal distributions (or, as they put it, are “more log normal than income” distributions)². Battistin, Blundell, and Lewbel (2007) show that this evidence can be accommodated by assuming that a sort of Gibrat’s Law of Proportionate Effects (Gibrat, 1931; Kalecki, 1945) holds for permanent income and, through intertemporal utility maximization, for household consumption (Hall, 1978)³.

This is a nice empirical result, because it seems to establish a stylized fact holding across cohorts and, possibly, countries, i.e. the distribution of the logarithms of HCEs is normal. In turn, it implies that all the moments of the distribution exist but we only need a two-parameter density to characterize the large majority of observed HCE patterns. The main empirical message of the paper is therefore that, as far as HCE distributions are concerned, there is no need to look at higher moments. Indeed, skewness, kurtosis, etc., of logged HCE distributions would mimic the correspondent Gaussian moments independently of cohorts, years and age classes. Furthermore, log-normality of HCE distributions implies that one can explain them by means of simple multiplicative growth models building upon the idea that consumption results from “the cumulation of random shocks to income and other variables that affect utility” (Battistin, Blundell, and Lewbel, 2007, p.4)⁴. If confirmed, this would be a powerful insight, which parallels similar ones obtained in industrial dynamics for firm size and growth rates (Ijiri and Simon, 1977; Sutton, 1997).

²Log-normality of HCE distributions in U.K. is confirmed by another early study in the econophysics domain, see Hohnisch, Pittnauer, and Chakrabarty (2002).

³Of course, one can derive log-normality of consumption directly from the hypothesis that individual consumption is approximately equal to permanent income (Friedman, 1957). Notice, however, that permanent income is not observable in practice.

⁴Notice that log-normality of individual consumption stems from log-normality of permanent income only if a long list of restrictions do indeed hold. Let aside the very hypothesis that individuals act as utility maximizers, one also needs that the random-shock process obeys some form of central-limit theorem and marginal utility is linear in log consumption.

In this paper, however, we show that log-normality is not generally the case for Italian HCE distributions. We employ the “Survey of Household Income and Wealth” (SHIW) provided by the Bank of Italy and we study HCE distributions with a parametric approach both in the aggregate and conditioned to the age of the household head (as reported in the survey) for a sequence of 8 waves from 1989 to 2004⁵. Unlike in Battistin, Blundell, and Lewbel (2007), who only check goodness of fit (GoF) employing graphical tools and Kolmogorov-Smirnov (KS) tests, we run a wider-range battery of GoF tests that overcome the well-known power limitations of the KS test.

Our main result is that in Italy, for all the waves under study and for the majority of age classes, the logs of HCE distributions are not normal but can be satisfactorily approximated by asymmetric exponential-power densities. This family of distributions features five parameters and allows one to flexibly model asymmetries in both the third and the fourth moment. Indeed, our statistical tests often reject the hypothesis that logs of HCE display zero-skewness and normal kurtosis. On the contrary, one invariably detects significant positive/negative skewness and asymmetry in the tail behavior. More specifically, the large majority of logged HCE distributions exhibit thick lower tails together with upper tails thinner than Gaussian ones. This evidence is quite robust to a series of further checks involving, e.g., estimation with robust statistics.

The basic message is that, at least for Italy, it seems impossible to come up with a statistical description of consumption data that can compress the existing heterogeneity in HCE distributions, so as to avoid a higher number of degrees of freedom in parametric characterizations. In other words, the existing, statistically-detectable, departures from log-normality prevent us from providing a simple two-parameter density that fits both aggregate and disaggregate HCE distributions. This in turn casts some doubts on the possibility to explain observable HCE distributions by means of simple, invariant models based on permanent income and Gibrat’s law hypotheses.

The rest of the paper is organized as follows. In Section 2 we describe the database

⁵Italian income distributions for SHIW data have been extensively studied in Clementi and Gallegati (2005b). They find that income is log-normal in the body and power-law in the upper tail. More on the relationships between income and consumption distributions is in Section 4.

that we employ in the analysis. Section 3 reports on our empirical results and related robustness checks. Section 4 presents a speculative discussion on the implications of our findings. Finally, Section 5 concludes.

2 Data

Our empirical analysis is based on the “Survey of Household Income and Wealth” (SHIW) provided by the Bank of Italy. The SHIW is one of the main sources of information on household income and consumption in Italy. Indeed, the quality of the SHIW is nowadays very similar to that of surveys in other comparable countries like France, Germany and the U.K.⁶.

The SHIW was firstly carried out in the 1960s with the goal of gathering data on incomes and savings of Italian households. Over the years, the survey has been widening its scopes. Household are now asked to provide, in addition to income and wealth information, also details on their consumption behavior and even their preferred payment methods. Since then, the SHIW was conducted yearly until 1987 (except for 1985) and every two years thereafter (the survey for 1997 was shifted to 1998). In 1989 a panel section consisting of units already interviewed in the previous survey was introduced in order to allow for time comparison.

The present analysis focuses on the period 1989-2004. We therefore have 8 waves. The sample used in the most recent surveys comprises about 8000 households (24000 individuals), distributed across about 300 Italian municipalities. The sample is representative of the Italian population and is based on a rotating panel targeted at 4000 units.

Available information includes data on household demographics (e.g., age of household head, number of household components, geographical area, etc.), disposable income, consumption expenditures, savings, and wealth. In this paper, we employ yearly data on aggregate HCEs⁷. We study both unconditional and age-conditioned distributions,

⁶SHIW data are regularly published in the Bank’s supplements to the Statistical Bulletin and made publicly available online at the URL <http://www.bancaditalia.it/statistiche/indcamp/bilfait>. We refer the reader to Brandolini (1999) for a detailed overview on data quality and main changes in the SHIW sample design.

⁷Data for disaggregated expenditure categories (e.g., nondurables, food, durables, etc.) are also

where age conventionally refers to the household head (on the problems related to assigning household-level data to its members, see for example Attanasio, 1999, section 2.2). Consumption figures have been converted to Euros for the entire period (1989-2000). Furthermore, HCEs have been weighted by using appropriate sample weights provided by the Bank of Italy. Finally, we deflated all consumption expenditure figures so as to obtain *real* HCE distributions.

More formally, our data structure consists of the aggregate distribution of yearly household (real) expenditure for consumption $\{C_{h,t}\}$, where $h = 1, \dots, H_t$ stands for households and $t \in T = \{1989, 1991, 1993, 1995, 1998, 2000, 2002, 2004\}$ are survey waves. Since in each wave there were many cases of unrealistic (e.g., zero or negative) consumption figures, we decided to drop such observations and to keep only strictly positive ones. We also dropped households for which consumption expenditures were larger than yearly income (as reported in the SHIW). Therefore, we ended up with a changing (but still very large) number of households in each wave (H_t). HCE distribution is complemented with information on the age of household head in wave t ($A_{h,t}$). We employ this variable to condition HCE distributions. More specifically, in line with Battistin, Blundell, and Lewbel (2007), we consider the following age breakdown: $A = \{A_1; \dots; A_8\} = \{\leq 30; 31-35; 36-40; 41-45; 46-50; 51-55; 56-60; \geq 61\}$, which generates sufficiently homogeneous subsamples as far as the number of observations is concerned. In each wave, we then build the distributions $\{C_{h,t} | A_{h,t} \in A_k\}$, with $k = 1, \dots, 8$. As usual, we will mainly employ natural logs of real consumption expenditure figures, defined as $c_{h,t} = \log(C_{h,t})$. Age-conditioned distributions will thus read $\{c_{h,t} | A_{h,t} \in A_k\}$, for $k = 1, \dots, 8$ and $t \in T$.

available. See also Section 5.

3 Towards a Characterization of Household Consumption Expenditure Distributions

In this Section, we shall explore the statistical properties of Italian HCE distributions and their evolution over time⁸. We are interested in answering four related questions: (i) Did HCE distributions exhibit structural changes over time? (ii) Can aggregate and age-conditioned HCE distributions be well-approximated by log-normal densities? (iii) If not, which are the causes of departures from log-normality? (iv) If HCE distributions are not log-normal, can one find alternative, better statistical descriptions of HCE distributions across age classes and time?

3.1 Time-Evolution of HCE Distributions

Let us begin with a descriptive analysis of HCE distributions and their evolution over time. Table 1 reports descriptive statistics for (real) aggregate $\{C_{h,t}\}$ distributions. Simple inspection shows that HCE sample moments are quite stable over time. Such evidence is confirmed by Figure 1, where the first four sample moments of logged HCE distributions $\{c_{h,t}\}$ are plotted against time. This means that, notwithstanding many household did probably move back and forth across quantiles, HCE distributions did not dramatically change their structural properties. This is a strong result, also in light of the introduction of the Euro in 2002.

Surprisingly enough, HCE distributions appear to be quite stable over time also when one conditions to age classes. Figure 2 reports plots of sample moments vs. waves for two age classes (left: 41-45; right: ≥ 61). As it can be easily seen, also within age classes HCE distributions have been remained quite stable over the years. Furthermore, one does not detect any evident trends in the first moments of the HCE distributions when, in each wave, they are plotted against age classes (see Figure 3 for wave 2004).

Notice also that if HCE distributions were lognormal, their logs would have been normally distributed, with zero skewness and kurtosis equal to 3. On the contrary, Table 2

⁸Additional details on the statistical analyses presented in this Section are available from the Authors upon request. All exercises were performed using MATLAB[®], version 7.4.0.287 (R2007a)

(top panel) shows that for aggregate distributions some positive skewness always emerges, while kurtosis levels fluctuate below and above the normal threshold. This is true also for age-conditioned distributions, see Table 2 (bottom panel). It is interesting to note that while HCE distributions appear to be right-skewed for almost all age classes and years, kurtosis is in the majority of cases below 3 (more on that below). Of course, to decide whether these departures from the normal benchmark are significant or not, one needs a more formal battery of statistical tests. This is what we shall do in the next section.

3.2 Are HCE Distributions Log-Normal?

To check whether HCE distributions are log-normal (or equivalently if their logs are normal), we have used a battery of three normality tests: Lilliefors (Lilliefors, 1967), Jarque-Bera (Bera and Jarque, 1980, 1981) and Quadratic Anderson-Darling (Anderson and Darling, 1954). These tests are known to perform better than comparable ones (e.g., KS test) in terms of power (see D’Agostino and Stephens, 1986; Thode, 2002, for details). More specifically, the Lilliefors test adapts the KS test to the case where parameters are unknown. In this sense, it can benchmark results obtained in Battistin, Blundell, and Lewbel (2007), who, as already mentioned, only employ the less-performing KS test. Finally, the Jarque-Bera test is known to perform better in presence of outliers, which is a commonly-detected problem for consumption data (more on this in Section 3.3).

Table 3 reports GoF results for logs of aggregate and age-conditioned HCE distributions. Aggregate distributions are never log-normal, while in the age-conditioned case, only for 9 distributions (out of 64) the three tests are simultaneously unable to reject (at 5%) the null hypothesis of log-normality (in boldface). Log-normality seems to be slightly more pervasive in age classes 36-40 and 56-60, and in the years from 1995 to 2000.

The above GoF evidence casts some doubts on whether consumption distributions can be well-approximated – in Italy – by log-normal densities. This seems to be true, for all waves under study, both at the aggregate level and after one conditions to age classes.

Our statistical findings are also detectable through standard graphical analyses. As an example, Figure 4 presents for wave 2004 a size-rank plot together with a log-normal

fit (left) and a QQ-plot against exponential quantiles (right) for the aggregate HCE distribution (cf. Embrechts, Kluppelberg, and Mikosh, 1997; Adler, Feldman, and Taqqu, 1998, for details). Even a simple visual inspection suggests that the empirical HCE is far from being log-normally distributed. Departures from log-normality emerge not only in the upper tail, but also in the central-leftward part of the distribution. Furthermore, concavities and convexities displayed in the QQ-plot signal a strong mismatch between the logs of HCE distribution and log-normal upper-tail behavior (Reiss and Thomas, 2001).

A more robust way to assess the thickness of upper tails is to take a non-parametric perspective and apply the Hill estimator (Hill, 1975)⁹. As Table 4 shows, the majority of Hill’s alpha estimates are larger than 2, implying lighter upper tails. The evidence about upper-tail behavior of HCE distributions seems however quite mixed: a clear, common pattern does not emerge, thus hinting to the necessity of going beyond a simple two-parameter statistical model. We shall go back to this issues in Section 3.4.

3.3 Robustness Checks

As discussed in Battistin, Blundell, and Lewbel (2007), consumption and income data generally suffer from under reporting (especially in the tails) and outliers, and Italian data are not an exception (Brandolini, 1999). In order to minimize the effect of gross errors and outliers, we have studied distribution moments and normality GoF tests with two alternative strategies.

First, we have employed robust statistics to estimate the moments of HEC distributions (Huber, 1981). More specifically, following Battistin, Blundell, and Lewbel (2007), we have used median (MED) and mean absolute deviation (MAD) as robust estimators for location and scale parameters. Furthermore we have estimated the third moment with quartile skewness (Groeneveld and Meeden, 1984) and kurtosis using Moors’s octile-based robust estimator (Moors, 1988). Results confirm, overall, our previous findings. Robust moments for the logs of HEC are stable over time (and within each wave, across age classes). Aggregate and conditioned (logs of) HEC distributions display a significant

⁹In order to select the most appropriate value for the cutoff parameter (k^*), we have employed here the procedure discussed in Drees and Kaufmann (1998). See Lux (2001) for details.

excess skewness, while robust kurtosis values are statistically different (according to standard bootstrap tests) from their expected value in normal samples (i.e. 1.233). Again, upper tails appear to be in general relatively light. Furthermore, we have computed normality tests on logs of consumption distributions standardized using robust statistics. More formally, for any given logged consumption distribution $\{c\}$, we have computed Lilliefors, Jarque-Bera and Quadratic Anderson-Darling tests on:

$$\tilde{c} = \frac{c - MED(c)}{MAD(c)}. \quad (1)$$

Table 5 reports p-values for the three tests in the aggregate and age-conditioned cases. Results confirm the evidence obtained without robust standardization. This implies that existing statistically-significant departures from log-normality are not due to outliers.

Second, we used both our original data and robustly-standardized samples to compute normality tests on sub-samples of logged HCE distributions obtained by truncating the upper or lower tail. More precisely, we defined left- and right-truncated distributions by cutting either the lower or the upper $x\%$ of the distribution. We then ran standard truncated-GoF normality tests (Chernobay, Rachev, and Fabozzi, 2005) by allowing $x \in \{5, 10, \dots, 25, 30\}$. In all our exercises (not shown), we ended up with p-values which were *even lower* than those obtained for the full samples, irrespective of whether the $x\%$ of the lower or the upper tail was removed.

3.4 Fitting Asymmetric Exponential-Power Densities to HCE Distributions

The foregoing analysis shows that HCE distributions can hardly be described by means of log-normal distributions. In other words, the same family of two-parameter density is not able to describe the existing, statistically-detectable, heterogeneity in HCE distributions. This is true both at the aggregate level and at the age-conditioned level, across years. The underlying cause of this distributional heterogeneity is the presence/absence of: (i) positive/negative skewness; (ii) leptokurtic/platykurtic behavior of the distribution as a

whole. However, it may well be that any given HCE distribution displays tails that look different between each other. This can happen if e.g. the upper (respectively, lower) tail is thinner (respectively, thicker) than a normal one. Furthermore, HCE distributions might exhibit a variability that is larger on the right (left) of their median or modal value than it is on its left (right).

To possibly accommodate all these departures from a well-behaved log-normal statistical model, we propose here to fit *the logs of* HCE distributions with a higher-parameterized, more flexible, distribution family known as the *asymmetric exponential power* (AEP). The density of the AEP family¹⁰ reads:

$$g(x; a_l, a_r, b_l, b_r, m) = \begin{cases} K^{-1} e^{-\frac{1}{b_l} \left| \frac{x-m}{a_l} \right|^{b_l}}, & x < m \\ K^{-1} e^{-\frac{1}{b_r} \left| \frac{x-m}{a_r} \right|^{b_r}}, & x \geq m \end{cases}, \quad (2)$$

where $K = a_l b_l^{1/b_l} \Gamma(1 + 1/b_l) + a_r b_r^{1/b_r} \Gamma(1 + 1/b_r)$, and Γ is the Gamma function. The AEP features five parameters. The parameter m controls for location. The two a 's parameters control for scale to the left (a_l) and to the right (a_r) of m . Larger values for a 's imply – *coeteris paribus* – a larger variability. Finally, the two b 's parameters govern the left (b_l) and right (b_r) tail behavior of the distribution. To illustrate this point, let us start with the case of a symmetric EP, i.e. when $a_l = a_r = a$ and $b_l = b_r = b$. It is easy to check that if $b = 2$, the EP boils down to the normal distribution. In that case, the correspondent HCE distribution would be log-normal. If $b < 2$, the EP displays tails thicker than a normal one, but still not heavy. In fact, for $b < 2$, the EP configures itself as a medium-tailed distribution, for which all moments exist. In the case $b = 1$ we recover the Laplace distribution. Finally, for $b > 2$ the EP features tails thinner than a normal one and still exponential.

It is easy to see that when one allows for different left-right a - and b -parameters, the AEP can encompass a wealth of different shapes. In Figure 5 we plot the log-density¹¹ of the AEP in both the symmetric and asymmetric case for different parameter values.

¹⁰The AEP density turns out to be a very good statistical model for many economic variables, like firm growth-rates and market-price returns. See Bottazzi and Secchi (2006b) and references cited therein for details.

¹¹Thus, in the normal case one will end up with a parabolic log-density shape.

Notice how the AEP can easily pick up across-distribution heterogeneity in skewness and kurtosis, but also within-distribution heterogeneity concerning variance and tail-thickness behaviors.

In what follows, we fit AEP densities to both aggregate and age-conditioned *logged* HCE distributions. Parameters are estimated via maximum likelihood¹².

To check whether logs of HCE distributions can be satisfactorily described by AEP densities, we firstly estimate AEP parameters via ML and then we employ a battery of GoF tests based on empirical distribution function (EDF) statistics. More specifically, we run three widely-used EDF-based GoF tests: Kuiper (KUI), Cramér-Von Mises (CVM) and Quadratic Anderson-Darling (AD2), with small-sample modifications usually considered in the literature (the AD2 test statistic employed here is analogous to that used above to test for normality: for more formal definitions, see D’Agostino and Stephens, 1986, Chapter 4, Table 4.2). We also compute a KS test to benchmark our results to those in Battistin, Blundell, and Lewbel (2007). Notice, however, that the former three tests are known to perform better – in terms of e.g. power – than the KS in all practical situations (Thode, 2002). All p-values for the test statistics are computed by running Monte-Carlo simulations (1000 replications) under the null hypothesis that the empirical sample comes from an AEP with unknown parameters; see Capasso et al. (2007) for a discussion.

Table 6 reports test statistics and Monte-Carlo p-values for the four GoF tests. Notice how AEP fits perform dramatically better than normal ones. P-values are almost always larger than 0.05, meaning that (if one takes 5% as the relevant significance level) in almost all the cases the logs of HCE distributions can be statistically described by AEP densities (and not by normal distributions). There are only three exceptions to this rule (in boldface). Indeed, for the age class 41-45, all four tests are rejected at 5% for the 1989 and 1995 waves. Borderline cases (still at 5%) are represented by the aggregate distribution and the 51-55 age class in 1989. However, if one lowers the significance level at 1% in almost every situation all four tests pass.

¹²See Agrò (1995) and Bottazzi and Secchi (2006b) for technical details. Estimation has been carried out with the package SUBBOTOOLS, available online at <http://www.cafed.eu>.

The main reason why AEP distributions are able to better explain, from a statistical perspective, the logs of HCE distributions appears to be evident from Table 7, where ML estimates of AEP parameters are shown. Let us focus on \hat{a} 's and \hat{b} 's parameter estimates (the estimate for the location parameter m is not relevant, as it closely tracks the mode of the distribution). It is easy to see that in about 79% of cases the right tail of the distribution appears to be thinner than the left tail ($\hat{b}_r > \hat{b}_l$). Furthermore, about 89% of all distributions display a right tail which is thinner than a normal one ($\hat{b}_r > 2$). Conversely, in 47% of times $\hat{b}_l < 2$, meaning that the left tail is thicker than a normal one. Right thin tails are also associated to higher \hat{a} 's parameters. In fact, about 90% of distributions obey the condition $\hat{a}_r > \hat{a}_l$.

This evidence seems to confirm our conjecture about the existence of a double-faced heterogeneity in HCE distributions. On the one hand, HCE distributions, even if they belong to the same family, are characterized, as expected, by very different parameter values. On the other hand, each HCE distribution displays very different structural properties as far as its left and right tails are concerned: the right tail is typically thinner but more dispersed than the left tail. This implies that a two-parameter, log-normal distribution is not enough to statistically model HCE distributions. As a consequence, one needs to employ higher-parameterized distributions that, as happens with the AEP, are able to account for these two levels of heterogeneity.

To further elucidate this point, Figure 6 plots AEP fits to aggregate HCE distributions in the period considered (similar patterns emerge also for age-conditional HCE distributions). Notice how AEP distributions are able to satisfactorily characterize the data. This is true especially as far as asymmetry in tail thickness is concerned, a feature that could never be accommodated by a symmetric, two-parameter density like the log-normal one.

4 Discussion

The foregoing findings convey two methodological messages. First, moments estimators (whether they are robust statistics or not) are not always able to shed light on within-distribution heterogeneity. Whereas HCE kurtosis levels hinted to the presence of tails

thinner than normal as a whole, EP fits have shown that this could be probably due to the presence of a medium-thick left tail coexisting with a thin right tail. Therefore, our study not only suggests that moments higher than the second one do indeed matter, but it also stresses the importance of investigating the existence of within-distribution asymmetries in higher moments. Second, and relatedly, the use of too generic statistical models (like the log-normal) or of non-parametric tools designed to look only at one side of the coin (e.g., the Hill's estimator, originally developed to study the right tail behavior only) might hinder the exploration of more subtle statistical properties, as the co-existence of tail asymmetries.

In fact, the AEP turned out to be an extremely flexible, but still parsimonious, family of densities capable of accounting for the extreme heterogeneity found in the data. This is important because without the AEP one should have employed two different distributions in order to account for different left and right tail behaviors. Having a single distribution that does the job is not only more elegant, but also more parsimonious from a statistical viewpoint.

Furthermore, our exercises show that alternative families of densities like the Levy-stable (Nolan, 2006) and the generalized hyperbolic (Barndorff-Nielsen, 1977), which can only account for tails thicker than normal ones, are not able to statistically describe our data, in the sense that GoF tests are always rejected. This implies, once again, that a key feature of the logs of HCE distributions is within- and cross-distribution heterogeneity in tail behavior.

It must be noted, however, that the interpretation of AEP fits discussed above should be scrutinized vis-à-vis a number of additional statistical tests. In fact, a deeper inspection of Table 7 reveals that \hat{b}_l and \hat{b}_r estimates are often close to each other (and to the normal threshold). As a rule of thumb (Agrò, 1995), one can use confidence intervals of the form $\hat{b} \pm 2s(\hat{b})$, where $s(\hat{b})$ are the standard errors of the estimates. In our case, given the large number of observations, standard errors are always very small. This means that in almost all cases differences between \hat{b}_l and \hat{b}_r are statistically significant. More generally, however, one should employ likelihood-ratio tests to check whether tail

asymmetries are statistically significant and whether \hat{b} estimates are different from 2 (see Fagiolo, Napoletano, and Roventini, 2007, for an application). Indeed, it might well turn out that in some cases (e.g., aggregate HCE, wave 1989) a more parsimonious symmetric EP density (i.e., $a_l = a_r = a$, $b_l = b_r = b$) can satisfactorily describe the data. This would still represent a departure from log-normality if the tail parameter b is statistically different from 2.

No matter whether a symmetric or asymmetric exponential-power fits the data, the existence of tails different from Gaussian ones for the logs of HCE distributions implies that the results in Battistin, Blundell, and Lewbel (2007) do not apply for Italy. Furthermore, even their weaker statement that consumption is more log-normal than income seems to be rejected by our data. Indeed, it is well-established (Clementi and Gallegati, 2005b) that Italian household income is log-normal in the body and power-law in the tail. Our results about thick-left and thin-right tails in HCE distributions suggest that, if any, it is income that is more log-normal than consumption. In other words, income and consumption distributions seem to belong to two different density families. Log of income displays a thick upper tail but features a standard-normal lower tail. On the contrary, consumption displays a thinner-than-normal upper tail and features a thicker-than-normal lower tail.

What does the foregoing statistical description add to our understanding of consumption? To begin with, a parametric approach can shed light on the nature of the existing heterogeneity in cross-section HCE distributions. The fact that the logs of HCE distributions can be robustly characterized – over the years – by AEP densities with thick lower tails and thin upper tails, enable us to better understand how consumption is distributed across households and to go beyond standard representative-hypothesis assumptions where only the first (and sometimes the second) moment matters. Since heterogeneity has been shown to be of crucial importance for aggregation (Forni and Lippi, 1997), a deeper knowledge of HCE distributional properties might hopefully help to better grasp the statistical properties of consumption dynamics at the macro level.

A second important reason why knowledge of cross-section HCE distributional prop-

erties may be important is that they can be considered as stylized facts that theoretical models should be able to replicate and explain. Of course, their “unconditional-object” status prevents us from univocally finding *the* generating process: as discussed in Brock (1999), there can be many alternative generating processes that could generate a given unconditional object such as a HCE distribution as their long-run equilibrium.

Nevertheless, knowledge about distributional properties of unconditional objects may help us in restricting the scope of the analysis. In other words, knowing that logs of HCE are not log-normal but AEP-distributed gives us some information on which generating processes one may exclude. For example, rejection of the null hypothesis that consumption is log-normally distributed irrespectively of disaggregation casts some serious doubt on whether some sort of Gibrat’s law is at work. If one indeed assumes – like in Battistin, Blundell, and Lewbel (2007) – that a simple multiplicative process applies to permanent income and marginal utility, the resulting limit HCE distributions would be log-normal. If this were the stylized fact to be replicated, one would have come very close to Champernowne’s indication about the strategy to follow in modeling economic phenomena (in his case, income dynamics):

The forces determining the distribution of incomes in any community are so varied and complex, and interact and fluctuate so continuously, that any theoretical model must either be unrealistically simplified or hopelessly complicated. We shall choose the former alternative but then give indications that the introduction of some of the more obvious complications of the real world does not seem to disturb the general trend of our conclusions. [Champernowne, 1953, p.319]

On the contrary, it appears that our results suggest to leave the “unrealistically simplified” domain: if stylized facts to be replicated are much more complicated than what a Gibrat’s law would have implied, also their theoretical explanation in terms of Gibrat’s dynamics could hardly hold.

However, the fact that our data reject simple multiplicative models *à la* Gibrat does not necessarily mean that the alternative explanation is “hopelessly complicated”. The idea is that a fruitful strategy to replicate and explain what we have found in Italian data is to build simple-enough stochastic equilibrium models like those developed in industrial dynamics to explain the emergence of power-laws and Laplace distributions for firm size

and growth rates (Bottazzi and Secchi 2006a; Fu et al. 2005). Along similar lines, one might attempt to single out the basic forces driving the generating process of cross-section distributions (e.g., imitation, innovation, fulfilment of basic needs, optimal choice under budget constraints, etc.) and find necessary conditions for the emergence of the observed stylized facts ¹³.

5 Conclusions

In this paper we have studied the statistical properties of Italian household consumption expenditure (HCE) distributions. We have found that, contrary to Battistin, Blundell, and Lewbel (2007), HCE distributions, both in the aggregate and within homogeneous age classes, are not log-normally distributed. Goodness-of-fit tests allow us to conclude that HCE distributions can be well approximated by asymmetric exponential-power densities. We have shown that departures from log-normality are due to a pervasive heterogeneity that is present at two different levels. First, given the same parametric fit, consumption distributions differ, as expected, in their parameters. Second, within each given distribution, the lower tail behaves differently from the upper tail. More specifically, our results indicate that, in the majority of cases, HCE distributions display a thick lower tail co-existing with a thin upper tail. These results hold vis-à-vis a series of further checks involving e.g. robust-statistic estimation.

The fact that the AEP performs better than a normal density in statistically describing HCE distributions suggests not only that higher moments matter, but also that within-distribution asymmetries in tail behavior can be important. It must be also noticed that the AEP density has been extensively used to statistically model many economic variables and indicators related to growth rates or returns expressed as differences between log levels (e.g., growth rates of firm size, returns of log prices, etc.; see Bottazzi and Secchi, 2006b). To our best knowledge, this is the first time that the AEP density is shown to provide a good approximation for the logs of the levels of a given variable (i.e., consumption

¹³This is quite in tune with the evolutionary agenda on consumption, see for example Aversi et al. (1999). Complementary ideas are in Witt (2001, 2007).

expenditures).

Several extensions to the present study can be conceived. First, it would be interesting to apply the same methodology employed above to U.K. and U.S. data to investigate whether one could detect in other countries the same departures found in Italian HCE distributions.

Second, one might explore the distributional properties of HCE data disaggregated over consumption categories (durables, non-durables, etc.) and study whether the tail behavior observed in the aggregate can be traced back to some particular consumption category or it is the mere effect of aggregation. Furthermore, since HCE distributions for consumption categories are likely to be correlated, it would be interesting to characterize the distributional properties of the joint G -dimensional HCE distribution (where G is the number of observable consumption categories).

Finally, another issue worth to be addressed concerns the statistical characterization of consumption budget-share (CBS) distributions (where for any given household the consumption budget-share for good g is simply a number in the unit interval defined as the share of expenditure for good g over total household consumption expenditure). Indeed, the fact that HCE distributions – disaggregated over consumption categories – are not statistically independent makes very hard to make predictions about the shape of CBS distributions, even if we knew how HCE marginals are distributed.

ACKNOWLEDGMENTS

Thanks to Angelo Secchi for his support with the SUBBOTOOLS package. All usual disclaimers apply.

References

- Adler, R., R. Feldman and M. Taquq (1998), *A Practical Guide to Heavy Tailed Data*. Boston, Birkhäuser.
- Agrò, G. (1995), “Maximum Likelihood Estimation for the Exponential Power Function Parameters”, *Communications in Statistics*, 24: 523–536.
- Anderson, T. and D. Darling (1954), “A test for goodness-of-fit”, *Journal of the American Statistical Association*, 49: 765–769.
- Attanasio, O. (1999), “Consumption Demand”, in Taylor, J. and M. Woodford (eds.), *Handbook of Macroeconomics*. Amsterdam, Elsevier Science.

- Aversi, R., G. Dosi, G. Fagiolo, M. Meacci and C. Olivetti (1999), “Demand Dynamics with Socially Evolving Preferences”, *Journal of Evolutionary Economics*, 8: 353–408.
- Axtell, R. (2001), “Zipf Distributions of U.S. Firm Sizes”, *Science*, 293: 1818–20.
- Barndorff-Nielsen, O. (1977), “Exponentially decreasing distributions for the logarithm of particle size”, *Proceedings of the Royal Society London A*, 353: 401–419.
- Battistin, E., R. Blundell and A. Lewbel (2007), “Why is Consumption More Log Normal Than Income? Gibrat’s Law Revisited”, IFS Working Papers WP08/07, London, U.K., The Institute for Fiscal Studies.
- Bera, A. and C. Jarque (1980), “Efficient tests for normality, homoscedasticity and serial independence of regression residuals”, *Economics Letters*, 6: 255–259.
- Bera, A. and C. Jarque (1981), “Efficient tests for normality, homoscedasticity and serial independence of regression residuals: Monte Carlo evidence”, *Economics Letters*, 7: 313–318.
- Bottazzi, G. and A. Secchi (2006a), “Explaining the Distribution of Firm Growth Rates”, *RAND Journal of Economics*, 37: 235–256.
- Bottazzi, G. and A. Secchi (2006b), “Maximum Likelihood Estimation of the Symmetric and Asymmetric Exponential Power Distribution”, Working Paper 2006/19, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies, Pisa, Italy.
- Brandolini, A. (1999), “The Distribution of Personal Income in Post-War Italy: Source Description, Data Quality, and the Time Pattern of Income Inequality”, Temi di Discussione (Economics Working Papers), No. 350, Bank of Italy, Rome.
- Brock, W. A. (1999), “Scaling in Economics: A Reader’s Guide”, *Industrial and Corporate Change*, 8(3): 409–46.
- Capasso, M., L. Alessi, M. Barigozzi and G. Fagiolo (2007), “On approximating the distributions of goodness-of-fit test statistics based on the empirical distribution function: The case of unknown parameters”, LEM Working Paper, 2007-23, Laboratory of Economics and Management, Sant’Anna School of Advanced Studies, Pisa, Italy.
- Champernowne, D. (1953), “A Model of Income Distribution”, *Economic Journal*, 63: 318–351.
- Chatterjee, A., S. Yarlagadda and B. Chakrabarti (eds.) (2005), *Econophysics of Wealth Distributions*. Milan, Springer-Verlag Italia.
- Chernobay, A., S. Rachev and F. Fabozzi (2005), “Composites goodness-of-fit tests for left-truncated loss samples”, Tech. Rep., University of California at Santa Barbara.
- Clementi, F. and M. Gallegati (2005a), “Pareto’s Law of Income Distribution: Evidence for Germany, the United Kingdom, and the United States”, in Chatterjee, A., S. Yarlagadda and B. Chakrabarti (eds.), *Econophysics of Wealth Distributions*. Milan, Springer-Verlag Italia.

- Clementi, F. and M. Gallegati (2005b), “Power Law Tails in the Italian Personal Income Distribution”, *Physica A: Statistical Mechanics and Theoretical Physics*, 350: 427438.
- D’Agostino, R. and M. Stephens (1986), *Goodness of Fit Techniques*. New York, Marcel Dekker.
- Drees, H. and E. Kaufmann (1998), “Selecting the optimal sample fraction in univariate extreme value estimation”, *Stochastic Processes and their Applications*, 75: 149–172.
- Embrechts, P., C. P. Kluppelberg and T. Mikosch (1997), *Modelling Extremal Events*. Berlin, Springer–Verlag.
- Fagiolo, G., M. Napoletano and A. Roventini (2007), “Are Output Growth–Rate Distributions Fat–Tailed? Some Evidence from OECD Countries”, *Journal of Applied Econometrics*, forthcoming.
- Forni, M. and M. Lippi (1997), *Aggregation and the Microfoundations of Dynamic Macroeconomics*. Oxford, Clarendon Press.
- Friedman, M. (1957), *A Theory of the Consumption Function*. Princeton, Princeton University Press.
- Fu, D., F. Pammolli, S. Buldyrev, M. Riccaboni, K. Matia, K. Yamasaki and H. Stanley (2005), “The growth of business firms: Theoretical framework and empirical evidence”, *Proceedings of the National Academy of Science*, 102: 18801–6.
- Gibrat, R. (1931), *Les Inégalités Economiques*. Paris, Librairie du Recueil Sirey.
- Groeneveld, R. and G. Meeden (1984), “Measuring Skewness and Kurtosis”, *The Statistician*, 33: 391–399.
- Hall, R. (1978), “Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence”, *Journal of Political Economy*, 86: 971–987.
- Hill, B. (1975), “A simple general approach to inference about the tail of a distribution”, *Annals of Statistics*, 3: 1163–1173.
- Hohnisch, M., S. Pittnauer and M. Chakrabarty (2002), “Empirical Regularities in Distributions of Individual Consumption Expenditure”, arxiv:cond–mat/0201195v1.
- Huber, P. (1981), *Robust Statistics*. New York, John Wiley Sons.
- Ijiri, Y. and H. Simon (1977), *Skew Distributions and the Sizes of Business Firms*. North–Holland, New York.
- Kalecki, M. (1945), “On the Gibrat Distribution”, *Econometrica*, 13: 161–170.
- Lee, Y., L. A. N. Amaral, D. Canning, M. Meyer and H. E. Stanley (1998), “Universal Features in the Growth Dynamics of Complex Organizations”, *Physical Review Letters*, 81: 3275–8.
- Lilliefors, H. (1967), “On the Kolmogorov–Smirnov tests for normality with mean and variance unknown.”, *Journal of the American Statistical Association*, 62: 399–402.

- Lux, T. (2001), “The limiting extremal behaviour of speculative returns: an analysis of intra-daily data from the Frankfurt Stock Exchange”, *Applied Financial Economics*, 11: 299–315.
- Moors, J. (1988), “A Quantile Alternative to Kurtosis”, *The Statistician*, 37: 25–32.
- Nolan, J. (2006), *Stable Distributions. Models for Heavy Tailed Data*. Birkhäuser: Boston.
- Reiss, R.-D. and M. Thomas (2001), *Statistical Analysis of Extreme Values*. Basel, Birkhäuser Verlag.
- Souma, W. (2001), “Universal Structure of the Personal Income Distribution”, *Fractals*, 9: 463–470.
- Subbotin, M. (1923), “On the Law of Frequency of Errors”, *Matematicheskii Sbornik*, 31: 296–301.
- Sutton, J. (1997), “Gibrat’s Legacy”, *Journal of Economic Literature*, 35: 40–59.
- Thode, H.C., J. (2002), *Testing for Normality*. New York, Marcel Dekker.
- Witt, U. (2001), “Special issue: Learning to consume – A theory of wants and the growth of demand”, *Journal of Evolutionary Economics*, 11: 23–36.
- Witt, U. (2007), “Evolutionary Economics and Psychology”, Papers on Economics and Evolution 2006–13, Max Planck Institute of Economics, Evolutionary Economics Group, Jena, Germany.

Table 1: Moments of aggregate HCE distributions vs. waves.

Stats	Waves							
	1989	1991	1993	1995	1998	2000	2002	2004
Mean	194.044	182.233	181.363	180.051	175.263	182.445	182.685	187.132
Std Dev	250.789	251.174	226.596	190.695	190.301	200.845	220.376	213.916
Skewness	3.357	3.732	3.294	2.655	3.100	2.830	3.215	3.253
Kurtosis	18.503	22.892	17.486	12.655	17.777	14.304	17.709	17.546
N Obs	7424.000	7208.000	6241.000	6274.000	5606.000	6292.000	6376.000	6277.000

Table 2: Skewness (top panel) and kurtosis (bottom panel) of HCE distributions vs. age classes and waves.

Skewness		Waves							
Age Classes	1989	1991	1993	1995	1998	2000	2002	2004	
≤ 30	0.452	0.186	0.358	0.203	0.048	0.016	0.130	0.228	
31-35	0.440	0.320	0.398	0.083	0.011	0.079	0.196	0.427	
36-40	0.329	0.289	0.403	0.030	0.164	0.060	0.289	0.372	
41-45	0.254	0.161	0.388	0.224	0.089	0.156	0.249	0.051	
46-50	0.258	0.107	0.132	0.095	0.080	0.176	0.408	0.251	
51-55	0.278	0.136	0.265	0.038	-0.189	0.156	0.083	0.263	
56-60	0.207	0.100	0.270	0.054	0.031	0.016	0.127	-0.074	
≥ 61	0.110	0.113	0.237	0.159	0.122	0.163	0.313	0.289	

Kurtosis		Waves							
Age Classes	1989	1991	1993	1995	1998	2000	2002	2004	
≤ 30	3.066	2.583	2.698	2.318	2.519	2.170	2.254	2.322	
31-35	2.959	2.898	2.735	2.602	2.553	2.611	2.531	2.577	
36-40	2.753	2.866	2.909	3.440	2.828	2.748	2.597	2.839	
41-45	2.558	2.687	2.798	2.530	2.492	2.501	2.628	3.022	
46-50	2.631	2.628	2.648	2.649	2.650	2.575	2.617	2.854	
51-55	2.791	2.915	2.711	2.486	2.596	2.529	2.687	2.950	
56-60	2.750	3.011	2.764	2.712	2.622	2.975	2.797	2.872	
≥ 61	2.632	2.794	2.862	2.609	2.670	2.812	2.632	2.873	

Table 3: Goodness-of-fit tests for the null hypothesis of normal logs of HCE distributions. Monte-Carlo (one-tailed) p-values in round brackets. JB: Jarque-Bera test; AD2: Anderson-Darling test. Boldface entries: The null hypothesis is never rejected (at 5%).

Age Class	Test	Waves									
		1989	1991	1993	1995	1998	2000	2002	2004		
Aggregate	Lilliefors	0.025 (0.000)	0.023 (0.000)	0.025 (0.000)	0.019 (0.000)	0.017 (0.002)	0.019 (0.000)	0.036 (0.000)	0.023 (0.000)		
	JB	65.371 (0.000)	28.831 (0.000)	62.612 (0.000)	42.774 (0.000)	35.355 (0.000)	39.652 (0.000)	100.242 (0.000)	57.291 (0.000)		
≤ 30	AD2	8.181 (0.000)	6.751 (0.000)	7.756 (0.000)	4.208 (0.000)	2.545 (0.000)	4.689 (0.000)	14.509 (0.000)	5.722 (0.000)		
	Lilliefors	0.0518 (0.0000)	0.0415 (0.0920)	0.0474 (0.1060)	0.0644 (0.0140)	0.0373 (0.0287)	0.0622 (0.0400)	0.0760 (0.0030)	0.0552 (0.0080)		
31-35	JB	20.5164 (0.0010)	5.0113 (0.0690)	7.5121 (0.0160)	6.9433 (0.0310)	2.0264 (0.0514)	6.5595 (0.0290)	6.2203 (0.0450)	6.0806 (0.0490)		
	AD2	0.0000 (0.0000)	0.0920 (0.0573)	0.1060 (0.0220)	0.0140 (0.0003)	0.7087 (0.0688)	0.0400 (0.0063)	0.0030 (0.0033)	0.1080 (0.0129)		
36-40	Lilliefors	0.0417 (0.0110)	0.0390 (0.0350)	0.0555 (0.0010)	0.0306 (0.0352)	0.0285 (0.0671)	0.0405 (0.0164)	0.0561 (0.0060)	0.0615 (0.0080)		
	JB	20.4006 (0.0000)	10.3319 (0.0080)	12.3648 (0.0070)	7.2821 (0.0168)	7.0633 (0.0183)	2.6383 (0.2315)	5.6840 (0.0570)	10.9717 (0.0150)		
41-45	AD2	0.0110 (0.0000)	0.0350 (0.0005)	0.0010 (0.0000)	0.4350 (0.0975)	0.6781 (0.0840)	0.1624 (0.0019)	0.0060 (0.0026)	0.0080 (0.0004)		
	Lilliefors	0.0454 (0.0010)	0.0566 (0.0000)	0.0351 (0.1357)	0.0305 (0.1112)	0.0280 (0.4525)	0.0254 (0.5482)	0.0454 (0.0120)	0.0642 (0.0000)		
46-50	JB	14.2769 (0.0030)	8.8226 (0.0140)	14.6828 (0.0030)	4.0984 (0.1148)	2.8605 (0.2334)	1.7654 (0.3968)	11.3930 (0.0090)	11.2770 (0.0090)		
	AD2	0.0010 (0.0000)	0.0000 (0.0000)	0.1357 (0.0005)	0.3112 (0.1687)	0.4525 (0.4312)	0.5482 (0.5391)	0.0120 (0.0003)	0.0000 (0.0001)		
51-55	Lilliefors	0.0402 (0.0030)	0.0347 (0.0240)	0.0544 (0.0000)	0.0415 (0.0210)	0.0318 (0.1780)	0.0374 (0.0400)	0.0514 (0.0030)	0.0249 (0.5942)		
	JB	15.8387 (0.0010)	6.7070 (0.0330)	16.9326 (0.0020)	10.0249 (0.0120)	6.6776 (0.0360)	8.6106 (0.0170)	9.2720 (0.0100)	0.2449 (0.8786)		
56-60	AD2	0.0030 (0.0000)	0.0240 (0.0020)	0.0532 (0.0000)	0.0210 (0.0006)	0.1780 (0.0365)	0.0400 (0.0014)	0.0030 (0.0006)	0.5942 (0.6104)		
	Lilliefors	0.0360 (0.0170)	0.0322 (0.0550)	0.0251 (0.0295)	0.0274 (0.0243)	0.0275 (0.2702)	0.0394 (0.0130)	0.0616 (0.0000)	0.0389 (0.0370)		
≥ 61	JB	13.5895 (0.0050)	5.6765 (0.0560)	5.1119 (0.0710)	4.4812 (0.1050)	4.1129 (0.1185)	8.5301 (0.0210)	21.6100 (0.0010)	6.6417 (0.0460)		
	AD2	0.0170 (0.0008)	0.0550 (0.0116)	0.4395 (0.2849)	0.2743 (0.0784)	0.2702 (0.1643)	0.0130 (0.0031)	0.0000 (0.0000)	0.0370 (0.0074)		
≥ 61	Lilliefors	0.0424 (0.0000)	0.0234 (0.3407)	0.0374 (0.0160)	0.0258 (0.3412)	0.0370 (0.0230)	0.0362 (0.0240)	0.0366 (0.0250)	0.0512 (0.0000)		
	JB	11.9020 (0.0050)	2.7986 (0.2322)	10.7778 (0.0090)	7.5082 (0.0320)	8.6477 (0.0240)	9.4538 (0.0130)	3.6450 (0.1551)	7.6226 (0.0280)		
≥ 61	AD2	0.0000 (0.0001)	0.3407 (0.1141)	0.0160 (0.0008)	0.3412 (0.0293)	0.0230 (0.0005)	0.0240 (0.0047)	0.0250 (0.0177)	0.5061 (0.0001)		
	Lilliefors	0.0469 (0.0010)	0.0326 (0.0570)	0.0429 (0.0060)	0.0240 (0.4385)	0.1255 (0.0429)	0.0363 (0.0350)	0.0249 (0.4611)	0.0267 (0.0281)		
≥ 61	JB	7.5154 (0.0310)	1.2201 (0.5210)	9.1354 (0.0150)	2.7416 (0.2278)	3.6515 (0.1605)	0.0458 (0.9760)	2.7667 (0.2425)	5.0774 (0.1187)		
	AD2	0.0010 (0.0006)	0.0570 (0.0162)	0.0060 (0.0002)	0.4385 (0.5161)	0.4629 (0.0208)	0.0350 (0.0445)	0.4611 (0.2329)	0.3881 (0.0878)		
≥ 61	Lilliefors	0.0167 (0.1210)	0.0296 (0.0000)	0.0255 (0.0020)	0.0276 (0.0000)	0.0230 (0.0060)	0.0176 (0.0540)	0.0469 (0.0000)	0.0295 (0.0000)		
	JB	17.3600 (0.0000)	9.8501 (0.0110)	24.1645 (0.0010)	26.1708 (0.0000)	14.3017 (0.0010)	14.9425 (0.0030)	58.9140 (0.0000)	41.7881 (0.0000)		
≥ 61	AD2	0.1210 (0.0032)	0.2410 (0.0000)	0.0020 (0.0000)	0.0650 (0.0000)	0.0060 (0.0022)	0.0540 (0.0032)	0.3256 (0.0000)	0.0782 (0.0000)		

Table 4: Estimates for Hill's α tail statistic. Aggregate and age-conditioned HCE distributions vs. waves.

Age Class	1989	1991	1993	1995	1998	2000	2002	2004
Aggregate	3.296	3.913	3.478	3.099	2.524	3.546	3.583	3.151
≤ 30	4.029	2.635	3.434	3.084	3.578	3.037	2.696	2.800
31-35	2.387	2.122	2.027	2.277	2.934	2.329	2.455	2.480
36-40	2.467	2.848	5.725	3.562	2.462	3.581	3.082	3.144
41-45	2.326	3.491	2.511	2.924	3.995	2.702	3.116	3.263
46-50	2.755	2.743	2.253	2.412	2.365	2.707	2.468	2.876
51-55	2.868	2.410	3.002	3.043	2.595	2.579	2.705	3.218
56-60	2.410	2.427	2.653	2.971	2.316	2.709	2.482	3.025
≥ 61	2.081	2.517	3.395	2.227	2.811	3.292	2.568	3.238

Table 5: Robustly-standardized distributions. Monte-Carlo (one-tailed) p-values for goodness-of-fit tests. Null hypothesis: normal logs of HCE distributions. JB: Jarque-Bera test; AD2: Anderson-Darling test. Boldface entries: The null hypothesis is never rejected (at 5%).

Age Class	Test	Waves							
		1989	1991	1993	1995	1998	2000	2002	2004
Aggregate	Lilliefors	0.000	0.000	0.000	0.000	0.002	0.000	0.000	0.000
	JB	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	AD2	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
≤ 30	Lilliefors	0.000	0.114	0.112	0.008	0.071	0.027	0.003	0.101
	JB	0.000	0.074	0.023	0.027	0.032	0.039	0.042	0.055
	AD2	0.000	0.057	0.022	0.000	0.119	0.006	0.003	0.013
31-35	Lilliefors	0.008	0.030	0.003	0.438	0.068	0.065	0.006	0.010
	JB	0.001	0.009	0.006	0.164	0.016	0.029	0.047	0.011
	AD2	0.000	0.001	0.000	0.097	0.041	0.074	0.003	0.000
36-40	Lilliefors	0.001	0.000	0.118	0.333	0.459	0.566	0.008	0.000
	JB	0.006	0.014	0.000	0.108	0.213	0.405	0.011	0.010
	AD2	0.000	0.000	0.001	0.169	0.431	0.539	0.000	0.000
41-45	Lilliefors	0.000	0.021	0.001	0.020	0.210	0.048	0.001	0.602
	JB	0.002	0.035	0.002	0.012	0.032	0.024	0.017	0.875
	AD2	0.000	0.002	0.000	0.001	0.037	0.001	0.001	0.610
46-50	Lilliefors	0.018	0.077	0.045	0.058	0.268	0.018	0.000	0.026
	JB	0.003	0.051	0.082	0.101	0.116	0.017	0.001	0.029
	AD2	0.001	0.012	0.028	0.079	0.164	0.003	0.000	0.007
51-55	Lilliefors	0.001	0.358	0.020	0.361	0.023	0.019	0.019	0.000
	JB	0.010	0.234	0.010	0.026	0.018	0.011	0.149	0.029
	AD2	0.000	0.114	0.001	0.029	0.001	0.005	0.018	0.000
56-60	Lilliefors	0.000	0.064	0.003	0.449	0.465	0.036	0.474	0.029
	JB	0.027	0.140	0.012	0.240	0.138	0.081	0.233	0.057
	AD2	0.001	0.016	0.000	0.516	0.521	0.044	0.233	0.022
≥ 61	Lilliefors	0.143	0.000	0.001	0.000	0.010	0.043	0.000	0.000
	JB	0.003	0.010	0.000	0.000	0.003	0.001	0.000	0.000
	AD2	0.003	0.000	0.000	0.000	0.002	0.003	0.000	0.000

Table 6: Goodness-of-fit tests for the null hypothesis of AEP-distributed logs of HCEs. KS=Kolmogorow-Smirnov test; KUI=Kuiper test; CVM: Cramér - Von Mises test; AD2: Quadratic Anderson-Darling test. Monte-Carlo (one-tailed) p-values in round brackets. Boldface entries: The null hypothesis is always rejected (at 5%).

Age Class	Test	Waves									
		1989	1991	1993	1995	1998	2000	2002	2004		
Aggregate	KS	1.146 (0.061)	0.722 (0.094)	0.713 (0.104)	0.666 (0.157)	0.533 (0.457)	1.016 (0.064)	0.489 (0.599)	0.640 (0.198)		
	KUI	1.904 (0.084)	1.146 (0.137)	1.212 (0.075)	1.257 (0.052)	1.064 (0.228)	1.634 (0.082)	0.950 (0.443)	1.249 (0.053)		
	CVM	0.250 (0.003)	0.074 (0.054)	0.080 (0.037)	0.062 (0.101)	0.055 (0.139)	0.157 (0.052)	0.039 (0.325)	0.071 (0.063)		
	AD2	1.586 (0.074)	0.438 (0.039)	0.792 (0.062)	0.428 (0.054)	0.385 (0.072)	0.920 (0.071)	0.243 (0.348)	0.570 (0.098)		
≤ 30	KS	0.408 (0.868)	0.520 (0.501)	0.401 (0.889)	0.857 (0.024)	0.844 (0.551)	0.531 (0.461)	0.688 (0.129)	0.667 (0.226)		
	KUI	0.789 (0.786)	0.971 (0.399)	0.798 (0.777)	1.236 (0.061)	1.271 (0.422)	0.988 (0.356)	1.186 (0.097)	1.151 (0.159)		
	CVM	0.018 (0.861)	0.041 (0.280)	0.030 (0.509)	0.135 (0.004)	0.132 (0.688)	0.026 (0.608)	0.053 (0.154)	0.085 (0.131)		
	AD2	0.204 (0.504)	0.278 (0.242)	0.191 (0.560)	0.646 (0.004)	0.682 (0.571)	0.206 (0.498)	0.280 (0.240)	0.316 (0.226)		
31-35	KS	0.544 (0.416)	0.499 (0.577)	0.495 (0.584)	0.603 (0.272)	0.742 (0.345)	0.768 (0.059)	0.749 (0.958)	0.411 (0.860)		
	KUI	0.984 (0.366)	0.887 (0.598)	0.980 (0.373)	1.018 (0.302)	1.396 (0.218)	1.383 (0.079)	1.410 (0.819)	0.788 (0.791)		
	CVM	0.035 (0.391)	0.025 (0.640)	0.037 (0.350)	0.046 (0.217)	0.101 (0.314)	0.075 (0.053)	0.088 (0.795)	0.018 (0.864)		
	AD2	0.237 (0.371)	0.225 (0.416)	0.240 (0.356)	0.257 (0.296)	0.445 (0.198)	0.436 (0.059)	0.461 (0.684)	0.154 (0.771)		
36-40	KS	0.621 (0.235)	0.665 (0.158)	0.457 (0.702)	0.481 (0.626)	0.663 (0.163)	0.515 (0.518)	0.686 (0.132)	0.875 (0.050)		
	KUI	1.139 (0.143)	1.270 (0.050)	0.893 (0.586)	0.941 (0.468)	1.072 (0.215)	0.793 (0.784)	1.215 (0.070)	1.455 (0.067)		
	CVM	0.074 (0.054)	0.056 (0.135)	0.026 (0.612)	0.022 (0.753)	0.039 (0.323)	0.031 (0.486)	0.063 (0.098)	0.068 (0.080)		
	AD2	0.511 (0.015)	0.326 (0.143)	0.248 (0.328)	0.168 (0.687)	0.208 (0.483)	0.208 (0.488)	0.367 (0.090)	0.443 (0.048)		
41-45	KS	0.855 (0.024)	0.520 (0.504)	0.502 (0.568)	0.853 (0.025)	0.328 (0.982)	1.416 (0.994)	0.694 (0.126)	0.423 (0.828)		
	KUI	1.562 (0.004)	0.858 (0.643)	0.880 (0.609)	1.337 (0.030)	0.650 (0.973)	2.150 (1.001)	1.012 (0.313)	0.741 (0.880)		
	CVM	0.129 (0.004)	0.028 (0.559)	0.028 (0.566)	0.082 (0.035)	0.015 (0.940)	0.308 (1.021)	0.033 (0.428)	0.023 (0.702)		
	AD2	0.785 (0.002)	0.158 (0.747)	0.256 (0.302)	0.506 (0.016)	0.147 (0.809)	1.668 (0.774)	0.208 (0.486)	0.216 (0.451)		
46-50	KS	0.624 (0.227)	0.470 (0.667)	0.835 (0.700)	0.849 (0.781)	0.428 (0.815)	0.687 (0.132)	0.833 (0.178)	0.810 (0.241)		
	KUI	1.140 (0.143)	0.822 (0.732)	1.263 (0.719)	1.272 (0.623)	0.825 (0.720)	1.115 (0.163)	1.211 (0.145)	1.242 (0.205)		
	CVM	0.065 (0.089)	0.024 (0.691)	0.109 (0.667)	0.179 (0.827)	0.021 (0.795)	0.085 (0.025)	0.091 (0.056)	0.125 (0.056)		
	AD2	0.391 (0.066)	0.168 (0.690)	0.600 (0.679)	0.683 (0.730)	0.197 (0.537)	0.454 (0.034)	0.618 (0.087)	0.670 (0.029)		
51-55	KS	0.829 (0.033)	0.460 (0.691)	0.533 (0.458)	0.438 (0.774)	0.650 (0.180)	0.476 (0.649)	0.498 (0.579)	0.671 (0.153)		
	KUI	1.268 (0.047)	0.816 (0.745)	0.921 (0.513)	0.855 (0.649)	1.159 (0.125)	0.899 (0.571)	0.942 (0.467)	1.093 (0.186)		
	CVM	0.073 (0.048)	0.021 (0.792)	0.038 (0.337)	0.028 (0.559)	0.064 (0.095)	0.035 (0.394)	0.029 (0.525)	0.046 (0.215)		
	AD2	0.489 (0.022)	0.230 (0.397)	0.308 (0.175)	0.245 (0.337)	0.392 (0.064)	0.230 (0.397)	0.238 (0.366)	0.346 (0.116)		
56-60	KS	0.893 (0.057)	0.796 (0.042)	0.478 (0.637)	0.401 (0.889)	0.372 (0.944)	0.870 (0.022)	0.381 (0.925)	0.365 (0.950)		
	KUI	1.316 (0.038)	1.411 (0.015)	0.858 (0.644)	0.781 (0.808)	0.714 (0.912)	1.211 (0.075)	0.755 (0.852)	0.729 (0.894)		
	CVM	0.094 (0.068)	0.076 (0.051)	0.028 (0.549)	0.023 (0.723)	0.022 (0.732)	0.052 (0.154)	0.017 (0.906)	0.019 (0.852)		
	AD2	0.515 (0.074)	0.390 (0.066)	0.250 (0.321)	0.140 (0.854)	0.172 (0.667)	0.275 (0.247)	0.133 (0.891)	0.150 (0.797)		
≥ 61	KS	0.741 (0.078)	0.517 (0.514)	0.464 (0.683)	0.448 (0.730)	0.715 (0.103)	0.345 (0.972)	0.636 (0.202)	0.786 (0.068)		
	KUI	1.401 (0.056)	0.990 (0.353)	0.823 (0.729)	0.884 (0.602)	1.363 (0.054)	0.675 (0.956)	1.144 (0.139)	1.202 (0.083)		
	CVM	0.082 (0.034)	0.042 (0.265)	0.028 (0.559)	0.024 (0.692)	0.079 (0.042)	0.017 (0.898)	0.043 (0.260)	0.058 (0.118)		
	AD2	0.542 (0.098)	0.287 (0.217)	0.207 (0.491)	0.173 (0.662)	0.405 (0.058)	0.138 (0.869)	0.249 (0.326)	0.361 (0.095)		

Table 7: Maximum-likelihood estimates for AEP parameters. Logs of aggregate and age-conditioned HCE distributions vs. waves.

Age Class	Par Est	Waves							
		1989	1991	1993	1995	1998	2000	2002	2004
Aggregate	b_l	2.1100	1.4800	1.9600	2.0600	2.8000	2.0000	1.7200	2.2300
	b_r	2.1600	2.8400	2.2300	2.7600	2.2200	2.6000	3.0000	2.4500
	a_l	0.8790	0.6840	0.7650	0.7540	1.0700	0.7640	0.5590	0.8140
	a_r	1.2000	1.6000	1.2100	1.2100	0.9340	1.2100	1.5300	1.0200
	m	4.4800	3.9300	4.3500	4.4400	4.8300	4.4400	4.0100	4.6400
≤ 30	b_l	1.8264	2.6838	3.4784	2.8409	2.3229	2.0270	1.6824	1.4698
	b_r	1.9099	2.5343	2.1067	4.7809	5.5946	5.9859	4.7128	4.6543
	a_l	0.6375	1.0720	0.9439	0.3336	0.4333	0.5746	0.4023	0.3939
	a_r	1.1933	1.3895	1.1355	1.9250	1.9339	1.8933	1.8632	1.8796
	m	4.0828	4.3805	4.6277	3.6510	3.8154	3.9950	3.8786	3.9322
31-35	b_l	2.5369	1.4394	1.1079	1.4773	1.6434	1.7550	2.0736	2.8661
	b_r	1.8686	2.5804	2.8514	3.7219	3.3895	2.8100	3.0101	2.8389
	a_l	0.8025	0.5577	0.3579	0.5240	0.5804	0.7265	0.6399	0.4790
	a_r	1.1243	1.6217	1.6403	1.4988	1.3814	1.2432	1.3074	1.3745
	m	4.4728	3.8696	3.9560	4.2383	4.5930	4.6114	4.6097	4.4728
36-40	b_l	2.1609	0.9380	3.2251	1.0805	2.9043	1.8151	2.2951	2.3470
	b_r	2.0893	3.0883	1.6099	2.8363	1.6710	2.5732	2.5396	1.8137
	a_l	0.7643	0.3794	0.9303	0.4707	1.0649	0.7106	0.7197	0.7209
	a_r	1.1940	1.7964	0.8543	1.3487	0.7860	1.0961	1.2988	0.9504
	m	4.5031	3.5896	4.9635	4.3117	5.0316	4.7428	4.5627	4.9490
41-45	b_l	2.5612	1.4617	1.6333	1.8705	3.8072	0.6505	1.2828	1.6956
	b_r	2.5086	3.1394	2.3128	3.0148	2.0921	5.3062	3.3669	2.2922
	a_l	0.8341	0.6264	0.5325	0.5753	1.3160	0.2537	0.3950	0.7260
	a_r	1.2709	1.7432	1.3212	1.3335	0.8500	2.1415	1.6184	1.0751
	m	4.6130	3.9111	4.3436	4.4315	5.2057	3.6109	4.0068	4.8672
46-50	b_l	2.8096	1.4084	1.3657	1.7577	1.9047	2.0853	2.1231	1.7792
	b_r	2.2515	3.4740	3.4019	3.0167	2.8639	2.7611	2.7404	3.1136
	a_l	0.9484	0.5992	0.6125	0.6877	0.7266	0.7180	0.7567	0.6186
	a_r	1.1956	1.7433	1.7715	1.4712	1.2562	1.2382	1.1886	1.3936
	m	4.7793	4.0615	4.1525	4.0397	4.3729	4.4334	4.4656	4.4622
51-55	b_l	2.1609	1.7359	2.0888	2.1000	3.2167	2.0871	1.3787	1.4694
	b_r	1.9891	2.3749	2.2832	3.4695	1.5723	3.1234	3.5366	2.0241
	a_l	0.8357	0.7851	0.7684	0.7527	1.4808	0.6823	0.5149	0.6065
	a_r	1.1406	1.3103	1.2501	1.3319	0.5256	1.3621	1.5937	1.1296
	m	4.6808	4.4028	4.4891	4.5970	5.4916	4.3501	3.8658	4.5630
56-60	b_l	1.8000	1.0879	1.2046	2.0186	2.2544	1.2993	2.1298	2.6124
	b_r	2.3144	2.9803	2.8425	2.6568	2.6424	3.1377	2.3666	1.6653
	a_l	0.7367	0.5377	0.4703	0.7786	0.9005	0.5924	0.8440	1.1678
	a_r	1.2507	1.6693	1.5856	1.1368	1.1019	1.4766	1.1392	0.6601
	m	4.3758	3.8507	3.7869	4.5036	4.6072	4.1829	4.3743	5.2561
≥ 61	b_l	2.2891	1.2890	2.1234	2.6800	2.5814	2.5327	1.5162	2.6018
	b_r	2.7515	3.1477	2.2968	2.8594	2.7732	2.9858	3.0719	2.9253
	a_l	1.1169	0.6045	0.8207	0.8811	0.9528	0.9494	0.4568	0.8324
	a_r	1.1290	1.7444	1.1313	1.0749	1.0168	0.9387	1.6372	0.9385
	m	4.4717	3.5217	4.2002	4.4132	4.5531	4.5833	3.7797	4.4634

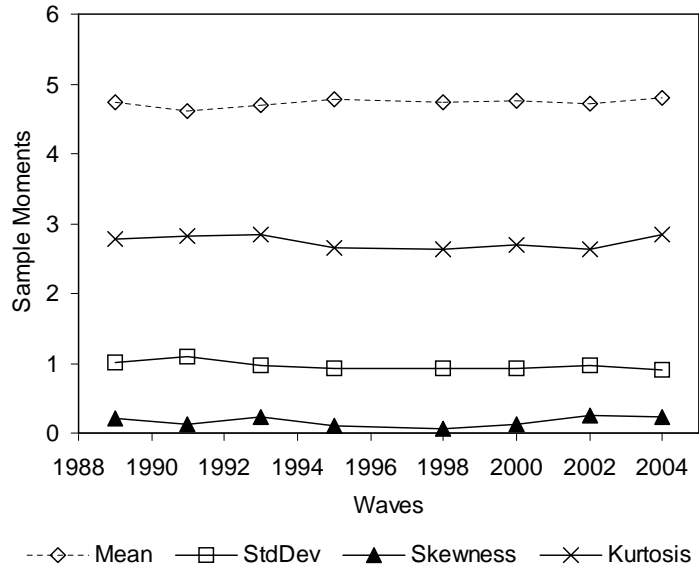


Figure 1: Moments of aggregate logs of HCE distributions vs. waves.




Figure 2: Moments of logs of age-conditioned HCE distributions vs. waves. Left: Class 4 (41-45). Right: Class 8 (≥ 61)




Figure 3: Moments of age-conditioned logs of HCE distributions vs. age classes in wave 2004.




Figure 4: Aggregate 2004 HCE distribution. Left: Size-rank plot (dots) vs. log-normal fit (dashed). Right: QQ-plot (dots) against exponential quantiles (solid line).

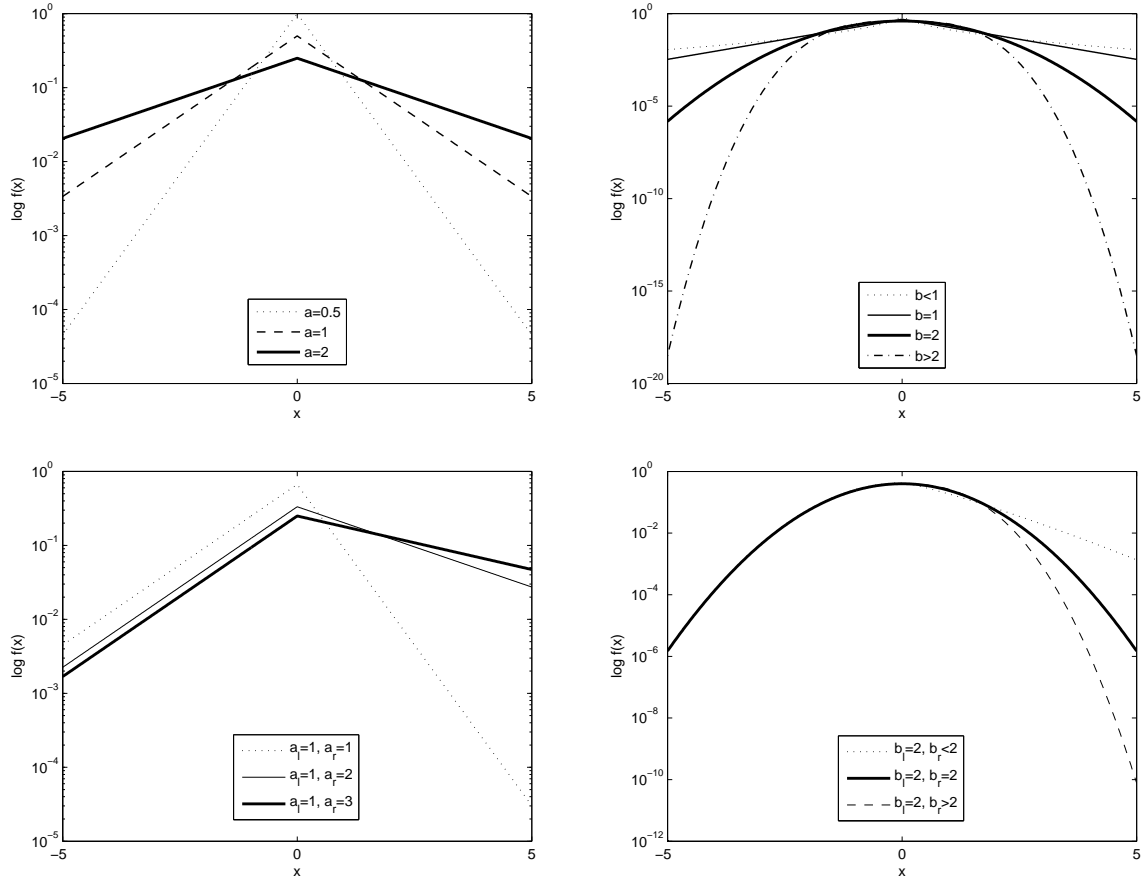


Figure 5: The asymmetric exponential-power density (logs) for different parameter values. Top-Left: Symmetric case for increasing shape-parameter (a). Other parameter values: $m = 0, b = 1$. Top-Right: Symmetric case for increasing tail-parameter (b). Other parameter values: $m = 0, a = 1$. Bottom-Left: Asymmetric case for increasing right shape-parameter (a_r). Other parameter values: $m = 0, b_l = b_r = 1$. Bottom-Right: Asymmetric case for increasing tail-parameter (b_r). Other parameter values: $m = 0, a_l = a_r = 1$.

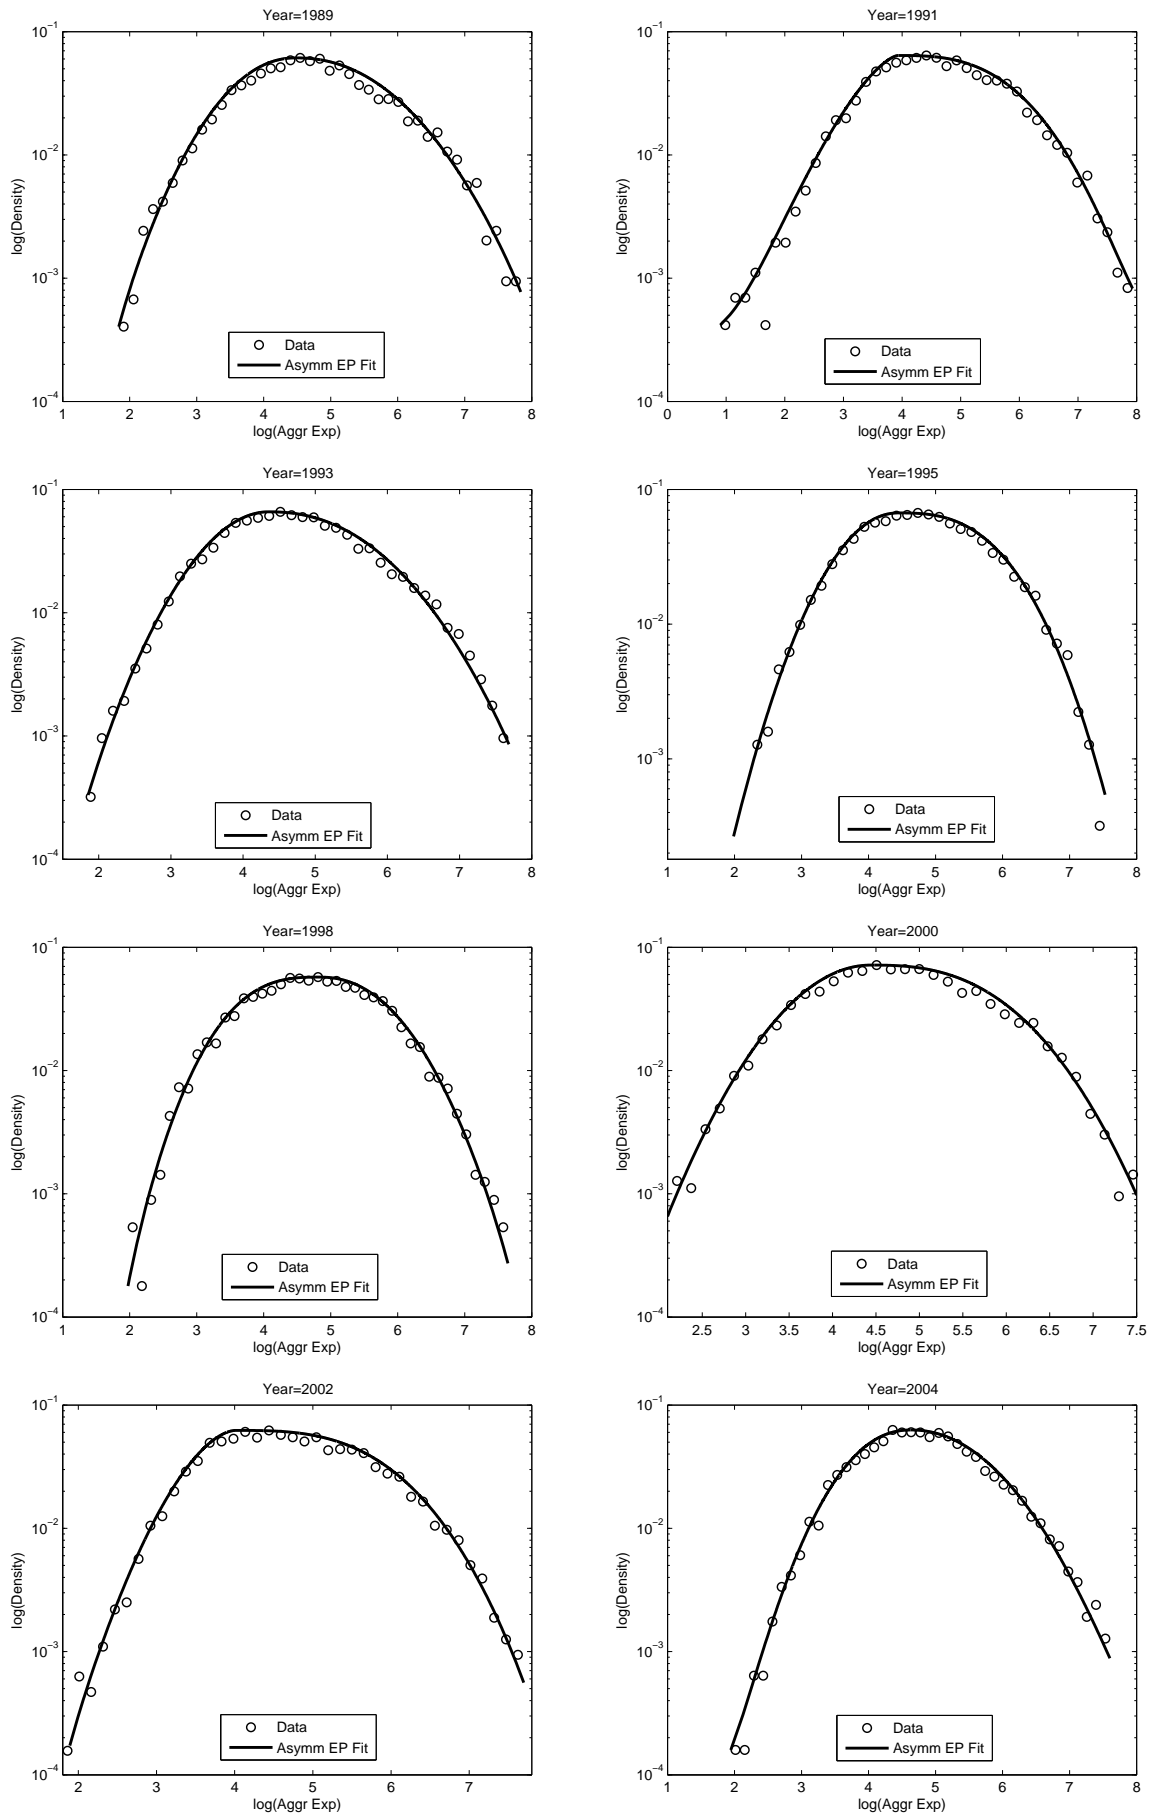


Figure 6: AEP fits for aggregate HCE distributions.