Cardon Research Papers
in Agricultural and Resource Economics

Research Paper 2004-19
September 2004

Notes on Thick-Tailed Distributions of Wealth

Lester D. Taylor
University of Arizona

Department of Agricultural and Resource Economics
College of Agriculture and Life Sciences
The University of Arizona

This paper is available online at http://ag.arizona.edu/arec/pubs/workingpapers.html

Abstract

Notes on Thick-Tailed Distributions of Wealth

Lester D. Taylor
University of Arizona*

One of the best-known empirical regularities in economics is the 'Law of Pareto', according to which the upper tails of the distributions of income and wealth are described by the relationship,

$$P(u > x) \quad = \quad Ax^{-\alpha},$$

for income and wealth levels (x) greater than some value x_0 *and* $\alpha < 2$. *However, despite its ubiquity across economies and time, the matter of why this 'law' holds remains an open question. Is there, as Mandelbrot has suggested, simply a base "prime mover" that gives rise to scalable distributions, not only for income and wealth, but also to a large number of other economic, social, and natural phenomena? If so, then the question of "why" is, by definition, not answerable. On the other hand, if thick upper tails can somehow emerge through processes that do not assume a Pareto prime mover, then the question is not only potentially answerable, but is also of a great deal of interest. Examining the possibility of this is the purpose of the present notes.*

Three scenarios involving highly stylized, artificial economies (all of which can be given a Darwinian interpretation) are simulated under varying assumptions regarding heritability, distribution of talents, and stability of tastes. Simulations with the first two scenarios make it pretty clear that talent differentials and pure randomness of tastes cannot suffice to produce wealth distributions with sufficient thickness to be interesting. However, things change in the third scenario, in which there is an allowance for preference stability, in the sense that once an agent experiences a good, that good is consumed with a non-zero probability in subsequent periods (so long, of course, as the agent remains "alive"). What the results with the third scenario show is that, with strong preference stability and substantial productivity differences amongst agents, thick-tailed distributions of wealth can emerge that have certain Pareto features and are log-log translatable.

* I am grateful to Teodosio Perez-Amaral and Barbara Sands for comments and to the Cardon Chair Endowment in the Department of Agricultural and Resource Economics at the University of Arizona for financial assistance.

Notes on Thick-Tailed Distributions of Wealth

Lester D. Taylor
University of Arizona

I. Introduction

One of the best-known empirical regularities in economics is the 'Law of Pareto', according to which the upper tails of the distributions of income and wealth are described by the relationship,

$$P(u > x) = Ax^{-\alpha}, \quad (1)$$

for income and wealth levels (x) greater than some value x_0 and $\alpha < 2$. However, despite its ubiquity across economies and time, the matter of why this 'law' holds remains an open question. Is there, as Mandelbrot has suggested, simply a base "prime mover" that gives rise to scalable distributions, not only for income and wealth, but also to a large number of other economic, social, and natural phenomena? If so, then the question of "why" is, by definition, not answerable. On the other hand, if thick upper tails can somehow emerge through processes that do not assume a Pareto prime mover, then the question is not only potentially answerable, but is also of a great deal of interest. Examining the possibility of this is the purpose of the present notes.

II. Background

In a production/exchange/consumption economy, the distribution of wealth can be viewed as the outcome of a sequence of processes as follows: production, which generates income, which funds consumption and saving, which determines which producers prosper (and therefore go on to live another day) and the accumulation of wealth. At base, this sequence of processes involves the conflation of three factors, namely, *talents, tastes, and heritability*. *Talents*, in conjunction with a principle of remuneration, determine the distribution of income. *Tastes*, in turn, determine how income is allocated between consumption and saving, and how consumption is allocated across producers. Finally, for wealth to accumulate over time, it obviously must be *heritable*.

Since heritability is a passive (but nevertheless necessary) enabler, it follows that the real drivers of the distributions of income and wealth are the distributions of talents and tastes. If, as is the usual assumption, talents are remunerated according to marginal-productivity principles, a thick-tailed distribution of talents will yield a thick-tailed distribution of income, and (with heritability and a saving rate that does not decrease with the level of income), ultimately to a thick-tailed distribution of wealth. However, as has been noted, the interesting question is whether the latter can emerge in the presence of only mild assumptions on the distribution of talents (and tastes).

With examination of this question as the goal, let us now turn to a highly stylized, artificial economy consisting of n agents producing and consuming n (possibly distinct) abstract consumption goods (whose prices are all 1) as follows:

Scenario 1:

(1). Each agent is both a producer and a consumer. Agents (for survival) consume one unit of a good during a period, but are able to produce (at constant cost) up to n units of their own good.

(2). During each period, agents "purchase" from each other (including themselves) at random according to draws from a uniform distribution.

(3). If an agent makes more than one "sale" during a period, the excess 'income' (where "excess" is defined as the number of sales minus one) is saved and carried over as 'wealth' to the next period.

(4). If an agent fails to make a sale during a period, consumption for the period can be financed by withdrawing one unit from the agent's stock of wealth. However, if the stock should be empty, the agent is assumed to "die", and is replaced by a new agent who begins from scratch.

Talents are represented in (1), tastes in (2), and heritability in (3). Talents are fixed and assumed to be the same for everyone. Tastes, on the other hand, are random, but are also assumed to be the same across consumers in the sense that the goods purchased by consumers are determined by draws from a common uniform distribution. With n consumers each consuming one unit in a period, the randomness of tastes implies that there will be a non-zero probability that some producers will have multiple sales in a period, while other producers will have no sales. The result will be a non-uniform distribution of income, and hence, through heritability, a non-uniform distribution of wealth.

An idea of the departure from non-uniformity under the conditions postulated in (1) - (3) is given in Table 1 for n equal to 2 through 5. From the numbers in this table, it is evident that the most likely outcome for a producer is to have no sales at all in a period, with a probability that increases with n. On the other hand, it is also evident that the probability of any producer becoming "super rich" during a single period is small. However, whether this small probability carries over to the distribution of wealth remains to be seen.

Table 1

n	P(0)	P(1)	P(2)	P(3)	P(4)	P(5)
2	0.33	0.33	0.33			
3	0.40	0.30	0.20	0.10		
4	0.43	0.29	0.17	0.09	0.03	
5	0.45	0.27	0.16	0.08	0.03	0.01

III. Some Initial Simulations

The exercise proceeds by simulating the economy described in rules (1) through (4) above over a number of periods and examining the resulting behavior of the distribution of "wealth". Table 2 tabulates simulation results for n equal to 100, 1000, 5000, and 10,000 for 10 periods. Two things are immediately apparent: (1) the distribution of wealth under the assumptions postulated is independent of the number of agents for $n \geq 100$ and (2) there is no evidence of a really fat upper tail. Common talents and purely random preferences (combined with heritability) thus do not appear to be sufficient for a thick upper tail.

Table 2

Simulated Wealth Distributions
Scenario 1

	Number of Agents			
P(x)	100	1000	5000	10,000
P(0)	0.333	0.308	0.285	0.296
P(1)	0.258	0.235	0.229	0.229
P(2)	0.111	0.165	0.177	0.179
P(3)	0.061	0.111	0.132	0.111
P(4)	0.091	0.083	0.075	0.082
P(5)	0.091	0.051	0.048	0.046
P(6)	0.015	0.020	0.025	0.028
P(7)	0.015	0.010	0.015	0.013
P(8)	0.000	0.010	0.006	0.010
P(9)	0.000	0.003	0.003	0.003
P(≥10)	0.000	0.001	0.004	0.002

The next steps, accordingly, will be to modify the assumptions governing the distributions of talents and tastes. This will be done in two stages. A second scenario will keep tastes random as in Scenario 1, but allows for a non-common distribution of talents, while a third scenario will inject some stability into tastes. The assumptions for Scenario 2 are as follows:

Scenario 2:

(5). Agents, as producers, are now assumed to be bifurcated into two groups, with one group having z times the "productivity" of the other (where $z > 1$). Assignment to the two groups is according to a binomial distribution with parameter pz..

(6). Producers in the "high" productivity group (i.e., whose productivity is z times that in the "low" productivity group) are assumed to have "costs of production" that are 1/z of the costs of the producers in the low productivity group. Consequently, additions to "wealth" for these agents will equal to the number of sales in the period minus 1/z, rather than the number of sales minus 1.

(7). Agents, as consumers, are (as in Scenario 1) assumed to consume one unit per period and to purchase randomly from the n producers with probability 1/n.

(8). As before, agents with negative wealth at the end of a period are assumed to "die", and are replaced with new agents (having no wealth) with productivity z or 1 as described in (5).

The results from simulating this scenario for 10 periods, for z = 2, pz = 0.2, and again for n equal to 100, 1000, 5000, and 10,000 (a single simulation for each n), are tabulated in Table 3.

Table 3

Simulated Wealth Distributions
Scenario 2

z = 2 pz = 0.2

	Number of Agents			
P(x)	100	1000	5000	10,000
P(0)	0.316	0.311	0.319	0.315
P(1)	0.165	0.182	0.171	0.177
P(2)	0.114	0.143	0.148	0.142
P(3)	0.152	0.116	0.108	0.112
P(4)	0.114	0.079	0.085	0.088
P(5)	0.063	0.053	0.056	0.055
P(6)	0.038	0.044	0.042	0.040
P(7)	0.000	0.026	0.024	0.029
P(8)	0.000	0.014	0.019	0.018
P(9)	0.025	0.012	0.014	0.012
P(≥10)	0.011	0.018	0.012	0.014

Comments:

(i). As with the first scenario, the distribution of wealth appears to stabilize for agents equal to 1000 or more.

(ii). The shape of the distribution, however, has changed. While the probability of no wealth is increased slightly (≅0.32 vs. ≅0.30), the really interesting thing is the transfer of "mass" from wealth stocks of 1 and 2 units to stocks of 6 or more units. This would appear to be a reflection of an increased mean stock of wealth that is allowed for by the differential in talents.

(iii). Variations in the values of z and pz will be examined below.

Scenario 3:

In this scenario, the assumption of purely random tastes is replaced with an assumption that, once formed, tastes tend to persist from one period to the next (provided, of course, that the agent continues to live). Specifically, the assumptions of this scenario are as follows:

(9). As in Scenario 2, agents, as producers, are assumed to be bifurcated into two groups, with one group have z times the "productivity" of the other (where $z > 1$). Assignment to the two groups is according to a binomial distribution with parameter pz.

(10). Agents, as consumers, are still assumed to consume one unit per period, but now with some stability in tastes as follows:

(i). In the initial period, agents purchase randomly from the n producers with probability 1/n.

(ii). However, in subsequent periods, agents purchase from the producers acquired from in the preceding period (provided that the producers are still "alive") with probability q and from the other n-1 producers with probability (1-q)/(n-1).

(iii). If the producers "acquired from" in the preceding period are no longer "alive", then agents purchase from all n agents with probability 1/n.

(11). Producers in the "high" productivity group (i.e., whose productivity is z times that in the "low" productivity group) are assumed to have "costs of production" that are 1/z. Consequently, additions to "wealth" for these agents are the number of sales in the period minus 1/z, rather than the number of sales minus 1.

(12). As before, agents with negative wealth at the end of a period are assumed to "die", and are replaced with new agents (having no wealth) with productivity z or 1 as described in (5).

The simulation results (again for 10 periods) for this scenario for z = 2, pz = 0.2, and q = 0.8 are given in Table 4.

Table 4

Simulated Wealth Distributions
Scenario 3
z = 2 pz = 0.2 q = 0.8

	Number of Agents			
P(x)	100	1000	5000	10,000
P(0)	0.414	0.423	0.404	0.404
P(1)	0.103	0.081	0.104	0.096
P(2)	0.034	0.063	0.076	0.085
P(3)	0.034	0.073	0.076	0.070
P(4)	0.103	0.061	0.061	0.057
P(5)	0.080	0.055	0.050	0.051
P(6)	0.046	0.043	0.040	0.043
P(7)	0.046	0.046	0.035	0.035
P(8)	0.011	0.023	0.029	0.030
P(9)	0.023	0.020	0.023	0.026
P(≥10)	0.103	0.112	0.101	0.102

Comments:

(i). With this scenario, the wealth distribution appears more to stabilize at 5000 agents, rather than at 1000 as with the first two scenarios.

(ii). Once again, the distribution is seen to change shape, with a shift in mass away from stocks of 1 to 3 units. This time, however, the shift is in both directions -- back to zero, as well as into the tail. The distribution has become much more (one-sided) leptokurtic. In short, preference stability (or what in the discussion below will be referred to as ‘habit formation’) clearly appears to be a condition allowing for the emergence of “large” fortunes.

III. Variations

We now turn to an examination of how sensitive Scenarios 2 and 3 are to the parameters governing the distribution of talents and tastes. Since the simulations involving 5000 and 10,000 agents are fairly computer intensive, the examination will be confined to economies with 1000 agents.[1] Table 5 presents results for Scenario 2 for different combinations of z and pz, while Table 6 presents results for Scenario 3 for different combinations of z, pz, and q. "Anchors" for comparison will be the columns for 1000 agents from Tables 3 and 4 (that is z = 2 and pz = 0.2 for Scenario 2, and z = 2, pz = 0.2, and q = 0.8 for Scenario 3).

For Scenario 2, in which only talents and their distribution across agents vary, the first interesting result appears to be the effect on the probability of zero wealth, especially with respect to the size of the binomial parameter governing the distribution of talents. In columns (1) to (3), in which the differential in productivity varies from 2 to 4, but pz remains constant at 0.2, the only noticeable effect is a mild flattening of the wealth distribution's tail. Variation in pz, holding z constant at 2 (cf., columns (4)-(7)), however, presents a totally different picture. Here, as pz varies from 0.4 to 0.05, the probability of a stock of wealth of zero increases monotonically from about 0.25 to a value approaching 0.4. Especially interesting is the fact that a tripling of the productivity differential in the last column (i.e., an increase of z from 2 to 6, but holding pz constant at 0.05) appears to have little impact. The conclusion that emerges from Table 5, accordingly, is that it is *how* talents are distributed, rather than *differences* in talents, *per se*, that affects the distribution of wealth.

Table 5

Simulated Wealth Distributions
Parameter Variations
Scenario 2

	Parameters							
P(x)	z pz	z pz	z pz	z pz	z pz	z pz	z pz	z pz
	2 0.2	*3 0.2*	*4 0.2*	*2 0.3*	*2 0.4*	*2 0.1*	*2 0.05*	*6 0.05*
P(0)	0.311	0.323	0.322	0.269	0.245	0.371	0.387	0.394
P(1)	0.182	0.170	0.146	0.160	0.162	0.182	0.171	0.201
P(2)	0.143	0.126	0.102	0.154	0.153	0.153	0.142	0.132
P(3)	0.116	0.093	0.084	0.114	0.126	0.092	0.101	0.094
P(4)	0.079	0.075	0.073	0.097	0.092	0.053	0.076	0.049
P(5)	0.053	0.053	0.085	0.058	0.068	0.056	0.062	0.039
P(6)	0.044	0.041	0.042	0.050	0.052	0.031	0.028	0.037
P(7)	0.026	0.036	0.043	0.039	0.038	0.025	0.020	0.017
P(8)	0.014	0.028	0.032	0.029	0.019	0.016	0.009	0.006
P(9)	0.012	0.024	0.024	0.017	0.017	0.007	0.000	0.007
P(≥10)	0.018	0.028	0.046	0.013	0.028	0.013	0.004	0.015
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)

[1]Simulations (on a standard PC with Pentium III chip) for 10 periods for Scenario 2 take about 8 minutes for 5000 agents and 40 minutes for 10,000 agents. For Scenario 3, the time required is about 10 minutes for 1000 agents, an hour for 5000 agents, and more than 10 hours for 10,000 agents. All coding and simulations have been done in SAS.

Scenario 3, it will be recalled, allows not only for talents to vary across agents, but also for the emergence of a stability in tastes, in the sense that agents (assuming that they survive) tend, with a certain probability, to purchase from the same producers from one period to the next. In Table 4, we saw how this "habit formation" lead to a noticeable fattening of the tail of the wealth distribution. In Table 6, we see how sensitive the shape of the distribution, especially the tail, is to the value of q. A q of 1 implies that purchases of surviving agents are all from producers of their first purchase, while a q of 0 of course implies that preferences are purely random from period to period, as in Scenarios 1 and 2. In columns 1 and 2, we see that increasing q from 0.8 to 1 leads to an increase in the probability of zero wealth from 0.42 to 0.48, combined with an increase in the tail probability (i.e., the probability of a stock of wealth of 10 units or more) from 0.11 to 0.14. Columns (2) and (3) present this sensitivity in even sharper relief, in that a halving of q from 1 to 0.5 is seen to lead to a decrease in the probability of zero wealth from 0.48 to 0.35, together with a decrease in the tail probability from 0.14 to 0.06. In short, stable preferences appear to be a powerful force in making for skewed distributions of income and wealth.

Table 6

Simulated Wealth Distributions
Parameter Variations
Scenario 3

	Parameters								
P(x)	z pz q	z pz q	z pz q	z pz q	z pz q	z pz q	z pz q	z pz q	z pz q
	2 0.2 0.8	*2 0.2 1*	*2 0.2 0.5*	*4 0.2 1*	*6 0.2 1*	*2 0.4 0.8*	*2 0.6 0.8*	*4 0.6 1*	*6 0.2 0.1*
P(0)	0.423	0.480	0.353	0.412	0.395	0.350	0.261	0.277	0.250
P(1)	0.081	0.079	0.135	0.092	0.101	0.099	0.095	0.097	0.121
P(2)	0.063	0.063	0.115	0.082	0.086	0.075	0.091	0.068	0.137
P(3)	0.073	0.050	0.090	0.050	0.045	0.073	0.069	0.039	0.106
P(4)	0.061	0.051	0.068	0.044	0.050	0.054	0.064	0.028	0.073
P(5)	0.055	0.037	0.045	0.047	0.035	0.056	0.068	0.035	0.072
P(6)	0.043	0.030	0.041	0.030	0.021	0.050	0.049	0.027	0.052
P(7)	0.046	0.028	0.035	0.023	0.026	0.047	0.065	0.038	0.043
P(8)	0.023	0.019	0.032	0.025	0.020	0.037	0.039	0.046	0.034
P(9)	0.020	0.023	0.028	0.033	0.034	0.024	0.047	0.057	0.038
P(≥10)	0.112	0.141	0.059	0.162	0.188	0.135	0.151	0.289	0.084
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)

However, what is especially interesting about the simulations in Table 6 is the interaction between the parameters governing talents and the amount of preference stability. This is particularly evident in column (8), which combines a high differential in productivity (z = 4) with a fairly even distribution of talents (pz = 0.6) and extreme preference stability (q = 1). The contrast with the combination of parameters in this column with the combination in column (5), which has a even higher differential in productivity (z = 6), but a much smaller presence of this differential (pz = 0.2), and a relatively minute amount of preference stability (q = 0.1), is striking.[2] The conclusion is thus the same as before: stable preferences are a potent enabler in the creation of heritable stocks of wealth.

IV. Interpretation of Parameters and Scenarios

Despite the fact that the economies being simulated are highly stylized and artificial, the parameters governing production, consumption, and the distribution of talents have reasonably straightforward interpretations, and indeed not implausible connections to the real world. To begin with, Scenario 1 in effect represents a totally subsistence economy in which everyone has to consume one unit of "bread" in order to survive. Everyone has the same talent, so that there is no accumulation from differential productivity. The only way that wealth can accumulate is by agents being lucky enough through the randomness of preferences to sell multiple units of output. Experience in consumption does not affect choice in this scenario. If an agent is fortunate enough to survive to the next period, the agent purchased from in the next period is totally random. The numbers in Table 2 suggest that agents do not survive long in this Scenario, and that the probability of becoming (and staying!) "super rich" is virtually nil. In short, there is little scope for a thick-tailed distribution of wealth ever arising in this scenario.

Things are different in Scenario 2, for by positing variation in talents, this scenario allows for *economic growth* to occur. Growth will be faster, the larger is z, and also the larger is pz. This scenario accordingly allows for accumulation to take place, not only from luck, but also from differential productivity. The result is wealth distributions with decidedly thicker tails. Nevertheless, since preferences are still purely random in this scenario, being able to hang onto a fortune, once acquired, remains a chancy proposition.

The introduction of preference stability through the parameter q in Scenario 3 posits, as has been noted at a couple of points, a probabilistic form of *habit formation*. At one extreme, a q of 1 implies that habit formation is complete -- i.e., that subsequent consumption is completely determined by initial experience -- while, at the other extreme, a q of 0 implies that preferences remain completely without structure, as in Scenarios 1 and 2. From the numbers in Table 6, it is clear that preference stability has a major impact on the shape of the wealth distribution, and indeed might even be sufficient (in conjunction with variation in talents) to give the distribution its real-world Pareto twist. We now turn to an examination of the extent to which this might be the case.

V. Law-of-Pareto Tests

---

[2] So striking, in fact, that it is useful to report a variation on the parameters in column (8), in which z and q are kept at 4 and 1, respectively, but pz is reduced to 0.2. The result (for the simulation undertaken) is an increase in P(0) to 0.423, and a reduction in $P(\geq 10)$ to 0.172. The point to note is that the latter remains high in relation to simulations in which q is less than 1.

In his classic paper in *The Journal of Political Economy* that introduced the study of long-tailed distributions of the Pareto type into economics, Mandelbrot (1963) made use of a simple graph to make many of his points. Taking logarithms of both sides of equation (1):

$$\ln P(u > x) = \ln A - \alpha \ln x, \qquad (2)$$

we see that a basic implication of the Pareto distribution is that the Pr (u > x) and x are linear in logarithms, with slope equal to minus α. In short, if the distribution of wealth is Pareto in its upper tail, then the graph of the upper tail should be a straight line when plotted on a double logarithmic grid.

Figure 1 provides graphs for four of the Scenario 3 simulations from Table 6. The blue graph is for the simulation in Column (1) in Table 6 (i.e., z = 2, pz = 0.2, and q = 0.8), pink corresponds to Column (4), while yellow and green correspond to Columns (8) and (9).[3] However, before drawing any conclusions from these graphs, several comments are in order:

(i). The Law-of-Pareto is an abstract distribution that refers to a large number of agents involving birth, death, consumption, production, and heritability processes that are (in some sense) in equilibrium. Since real-world distributions are necessarily finite, the Law cannot explain the very upper reaches of the distributions of income and wealth (since these necessarily drop to zero, in a possibly precipitous fashion). What this means, accordingly, is that we should not expect to find linearity (or at least linearity with slope greater than minus 2) at the very extreme of the upper tail of the distribution of wealth, but rather, at best, only for an interval of values leading to the extreme.

(ii). The distributions depicted in the figure are generated from simulations of economies that are extremely "small" in relation to the real-world, "small" both in the sense of the number of agents (1000) and the "age" of the economies (10). While the results in Tables 2 - 4 suggest that, for the scenarios being investigated, the wealth distribution stabilizes fairly quickly in terms of agents for a *given* number of periods, this may not be the case for the number of periods, for, if nothing else, the distribution will almost certainly "move to the right" as the number of periods increases.

Looking first at the "anchor" case for Scenario 3 (i.e., the second graph from the left in Figure 1), we find perhaps a hint of linearity for stocks of wealth between 9 and 19 units. However, the

[3] The graphs in these figures are for different realizations than those shown in Table 6. In the captions of the figures, npds refers to the number of periods in the simulations. Graphs have been constructed in Excel.

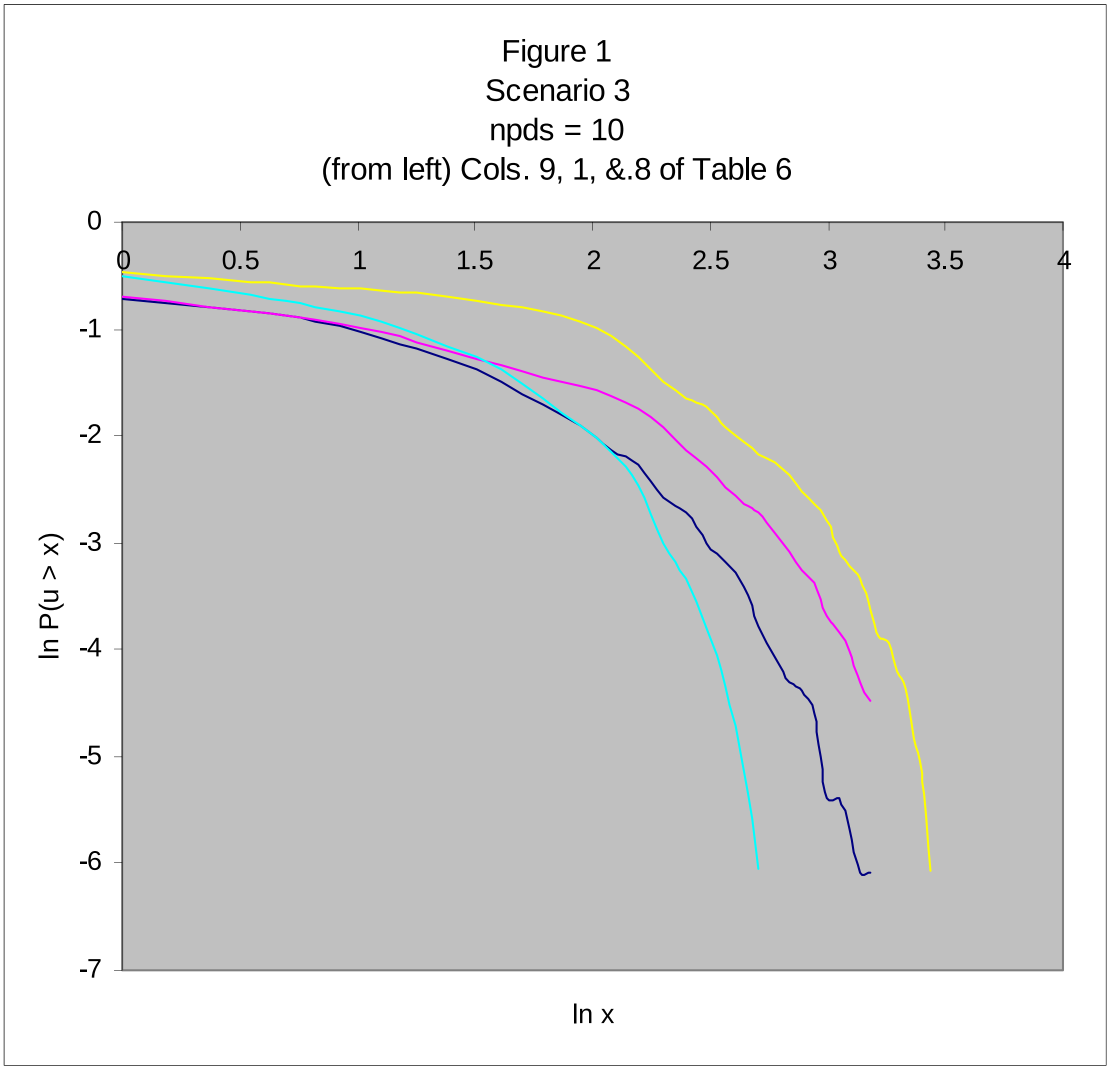

slope of the segment is clearly less than −2. In short, there is no evidence of a Pareto tail for this case. On the other hand, for the third graph from the left, for which the Scenario 3 parameters are z = 4, pz = 0.2, and q = 1, we see that are essentially three connected regions of linearity. The first is for 2 through 7 units, the second for 8 though 18, and the third is for 19 on. Of the three segments, the second is clearly of most interest, especially if the slope is greater than -2. However, a regression of ln P(y > x) on ln x yields a slope of -2.15 (with a t-ratio of -35 and an R^2 of 0.99). Turning now to the graph at the far right, a near linear segment somewhat similar to the second segment in the third graph from the left, is defined for x between 7 and 24. This time, a regression equation yields a slope of -2.07, with a t-ratio of -20 and an R^2 of 0.96 -- in short, we are getting closer to possibly some form of the Law-of-Pareto, but are not yet quite there! Finally, in the first graph on the left, we see that, in the face of extreme (but isolated) talent differences (z = 6, pz = 0.2), but little preference stability (q = 0.1), there is no evidence at all of a Law-of-Pareto.

If we were to stop at this juncture, the conclusion would pretty much have to be that, while strong habit formation (i.e., preference stability), combined with variations in talents and heritibility, yield wealth distributions that are increasing long-tailed, the processes stop short of the Law-of-Pareto. However, as noted in comment (ii) above, an economy of only 10 periods is scarcely out of a "state of nature", and it may be that the Law-of-Pareto emerges only with the passage of many generations. As a partial test of this, a simulation entailing 100 periods has been undertaken, using what from the earlier simulations appear to be particularly sensitive values for the parameters, namely, z = 6, pz = 0.4, and q = 1. The resulting graph of ln P(y > x) against ln x is given in Figure 2. Bingo! For what from the earlier simulations appear to be particularly sensitive values for the parameters, namely, z = 6, pz = 0.4, and q = 1. The resulting graph of ln P(y > x) against ln x is given in Figure 2. Bingo! For the regression slope for the segment for x between 9 and 89 is -0.86 (with a t-ratio of -37 and an R^2 of 0.94)! Not only is this consistent with the Law-of-Pareto, but would appear to be a thick-tailed distribution with a vengeance -- for a slope of -0.86 for such segment (being greater than -1) -- would be consistent with a distribution that does not even have a mean! Whatever, further investigation of this scenario seems clearly to be in order.

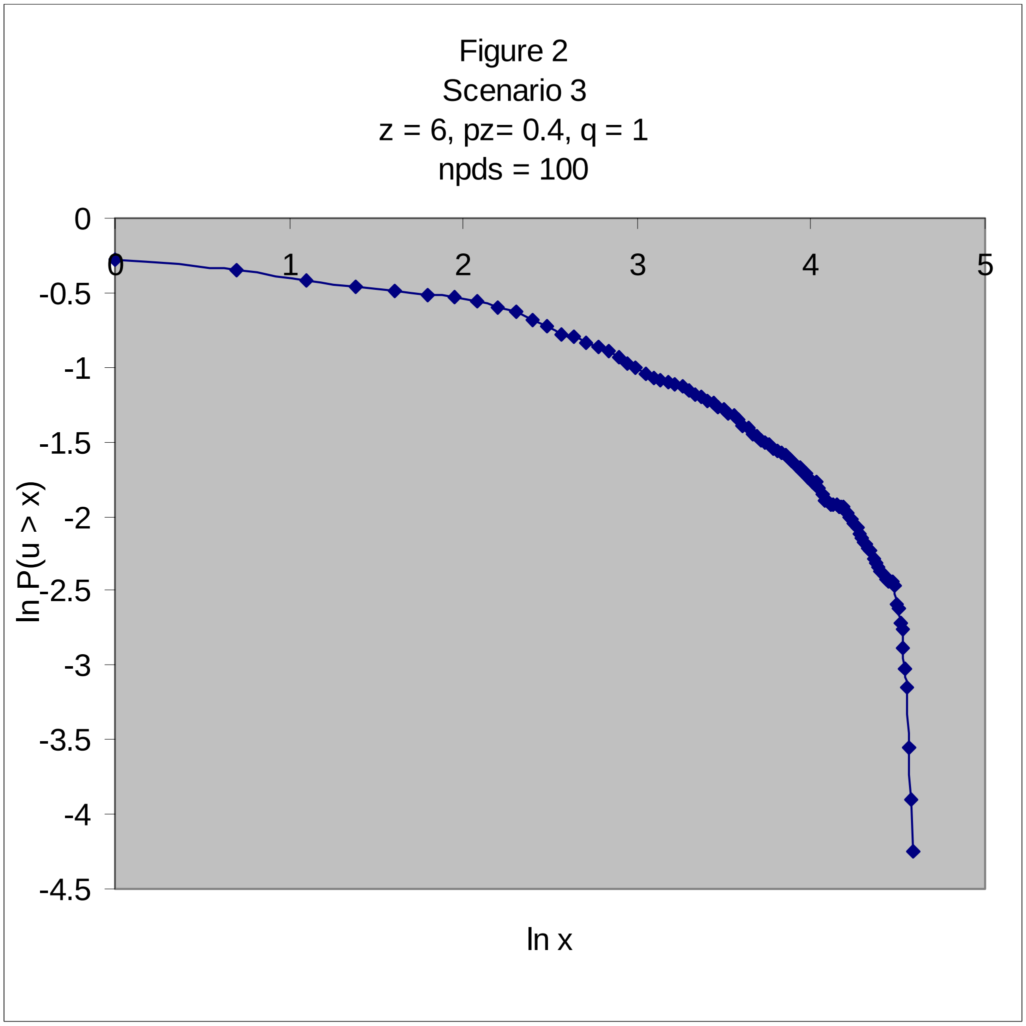

Figure 3 provides comparative graphs for z = 6, pz = 0.4, and q = 1 for periods of 100, 200, 300, and 400. First, we see that the log-log graphs, particularly the last three, are pretty much right-ward translations of one another. As in Figure 2, the middle segments of the graphs are strikingly linear, and with what appears to be a common slope. The slope coefficients, t-ratios, and R^2's for linear least-squares regression lines fitted to these segments are given in Table 7. The slopes are seen to vary from -0.31 for the period of 100 to -0.51 for 200. The slopes for 300 and 400 are virtually the same as for 200.[4] The t-ratios for the estimated slope coefficients are huge, and R^2's are 0.9760 or higher.[5]

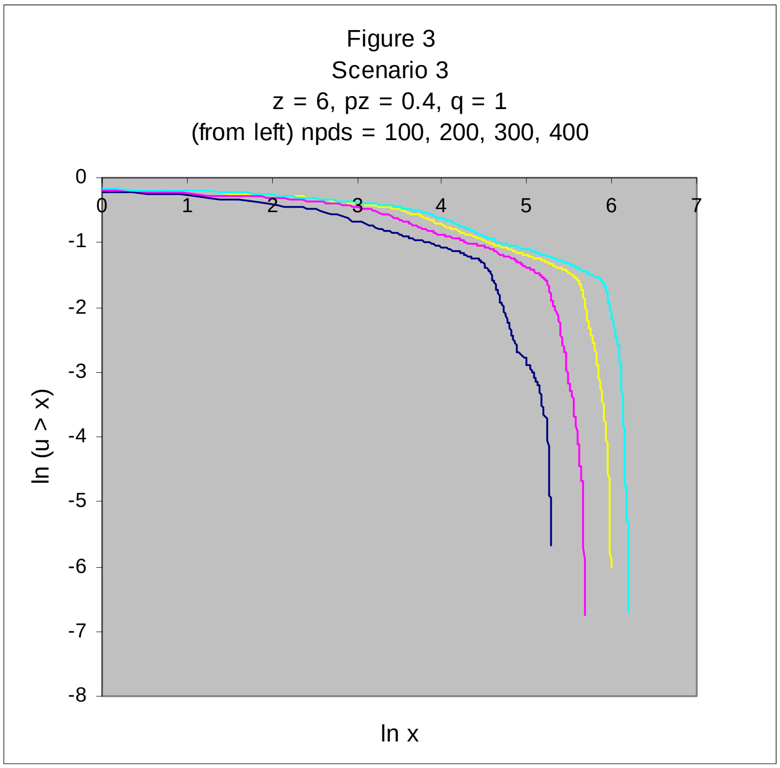

[4] If a period is interpreted as a year, then periods of 2-400 should represent a plausible length of "history" for the distribution of wealth to reach an asymptotic form. The small size of the economy (1000 agents) is not thought to be a problem because of the results presented in Table 4 above, which show the distribution to be reasonably independent of the number of agents.

[5] Casual inspection of the slope coefficients for periods of 200, 300, and 400 leads to what would seem like an obvious conclusion that the slopes are simply random deviations from an underlying common value. However, a pooled regression shows that the slope coefficient for 200 periods differs from the slope coefficient for 300 periods by a t-ratio of -3.19, and that the slope for 200 periods differs from a common slope for 300 and 400 periods by a t-ratio of -4.49. (The t-ratio for the difference between 300 and 400 periods is 1.80.) However, in the present context, this difference, while highly significant statistically, seems of little practical importance.

Table 7

Slopes of "Middle" Segments
of Graphs in Figure 3

NPDS	Slope Coefficient	t-ratio	R^2	Interval for x	1−F(x)
100	-0.4091	-108	0.9933	12 - 92	0.256
200	-0.5120	- 82	0.9760	35 - 191	0.189
300	-0.4945	-189	0.9933	32 - 277	0.196
400	-0.4861	-174	0.9896	49 - 370	0.190

The results in Figure 3 and Table 7 leave little question that a distribution of wealth with a *Pareto portion* (i.e., a linear log-log segment, with slope greater than -0.2, in the upper tail) can emerge in the absence of a Pareto "prime mover". However, whether this segment is adequately located in the upper tail to be able to conclude that the upper tail does in fact follow the Law-of-Pareto is another matter. The column headed by "Interval for x" in Table 7 gives the "wealth" intervals over which the linear regressions are fitted, 12 through 92 units of wealth for the period of 100, 35 through 191 for 200, etc. The column headed by "1 - F(x) gives the upper tails of the distribution of wealth that remains at the ends of the intervals -- 25.6 percent of agents have wealth greater than 92 units for the period of 100, and so on and so forth. Is 25 percent of the tail remaining (or about 20 percent in the case of periods of 200 - 400), which decidedly has an α much larger than 2, sufficient to foreclose the presence of a Pareto law?

The literature on thick-tailed distributions unfortunately does not provide much empirical guidance to this question. Mandelbrot, who has pioneered statistical study of fat-tailed distributions of the Pareto-Levy type, has always seemed to be most concerned about proper estimation of α. In a comment on the significance of the evidence provided by doubly logarithmic graphs, Mandelbrot had this to say:

> Limitations on the value of α lead to another quite different aspect of the general problem of observation. It concerns the practical significance of statements having only asymptotic validity. Indeed, to verify empirically the scaling distribution, the usual first step is to draw a doubly logarithmic graph: a plot of $\log_{10}[1 - F(u)]$ as a function of $\log_{10}[u]$. One should find that the graph is a straight line with the slope -α, or at least that this graph becomes straight as u increases. But, look closer at the sampling point of the largest u. Except for the distribution of incomes, one seldom has samples over 1000 or 2000 items; therefore, one seldom knows the value of u that is exceeded with the frequency $1 - F(u) = 1000^{-1}$ or 2000^{-1}. That is, the "height" of the sampling doubly logarithmic graph will seldom exceed *three* units of the decimal logarithm of 1 - F. The "width" of this graph will be at best equal to 3/α

units of the decimal logarithm of u. However, if one wants to estimate reliably the value of the slope α, it is necessary that the width of the graph be close to one unit. In conclusion, one cannot trust any data that suggest that α is larger than 3. [Mandelbrot (1997, p. 87.]

Note that Mandelbrot does not say anything in this comment, nor anywhere else in his writings that I have been able to find, about how far into the tail of the distribution that one needs to be concerned with doubly logarithmic linearity. I say this, *as an empirical matter*, because of the fact that, since in any real-world circumstance we are necessarily dealing with distributions with a finite number of elements, there must eventually reach a point in the distribution where the density function must unceremoniously drop to zero (and stay there). Is this beyond $1 - F(x) = 1000^{-1}$ or 2000^{-1}, as (possibly) suggested by Mandelbrot, or might one take such a drop as beginning as early as $1 - F(x) = 0.20$ (as in the graphs for periods of 200 - 400 in Figure 3)?

Before taking up a discussion of this question, it will be useful to consider a fourth scenario, in which the distribution of talents is assumed to be trinomial, rather than binomial. Specifically, this scenario allows for three categories of "productivity", z, v, and 1, for $z > v > 1$, with probabilities pz, pv, and 1-pz-pv, respectively. Double-log graphs of ln P(u>x) and ln x corresponding to simulations for n = 1000 and periods of 100-400, for z = 5, v = 3, pz = 0.2, pv = 0.6, and q = 1, are given in Figure 4. Comparison of the graphs in this figure with the graphs in Figure 3 show important similarities, but also an important difference. The similarities are: (1) the distributions for npds = 200-400 are essentially rightward translations of one another, and (2) there are again three reasonably well-defined linear segments. The important difference, however, is in the "middle" segment (for npds = 200-400), which occurs earlier in the distribution [between 1 - F(x) equal to 0.40 and 0.68,
as opposed to values of 0.55 and 0.80 for the distributions in Figure 3], and now has a slope of about -0.33, as compared with slopes of about -0.5 for the comparable segments in Figure 3. Thus, while introduction of a "hump" in the distribution of talents retains a Pareto-like segment in the distribution of wealth, the point at which segment stops leaves a tail containing more than 30 percent of the mass of the distribution. This seems a bit too much to support any conclusion that the tails of the distributions in Figure 4 are scaling.

However, that they are thick-tailed is another matter, and in investigation of this, it is instructive to examine the behavior of the "spread" of the distributions represented in Figure 4. This is done in Figure 5, which presents graphs of the standard deviations for the distributions in Figures 3 and 4. Of all the results presented to this point, this figure has to be the most dramatic[6], for, while a sequence of four elements hardly qualifies as determining asymptotic behavior, the two graphs are essentially perfectly linear, with slopes considerably greater than zero. In short, there is not even a hint of the standard deviation ever reaching an asymptote. Figure 5 refers to standard deviations.

---

[6] The point of this exercise is that, as discussed by Mandelbrot in his 1963 *JPE* paper, an implication of an distribution with an infinite variance is for sample variances to behave erratically as a function of increasing sample size, with no tendency to reach an asymptote.

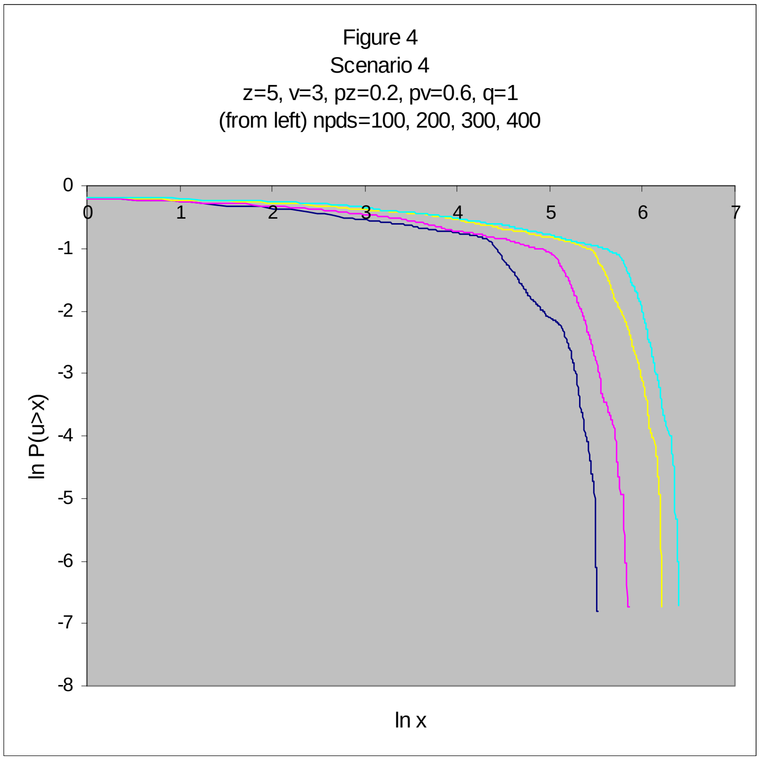
Figure 4
Scenario 4
z=5, v=3, pz=0.2, pv=0.6, q=1
(from left) npds=100, 200, 300, 400
ln P(u>x)
ln x
0
-1
-2
-3
-4
-5
-6
-7
-8
0
1
2
3
4
5
6
7

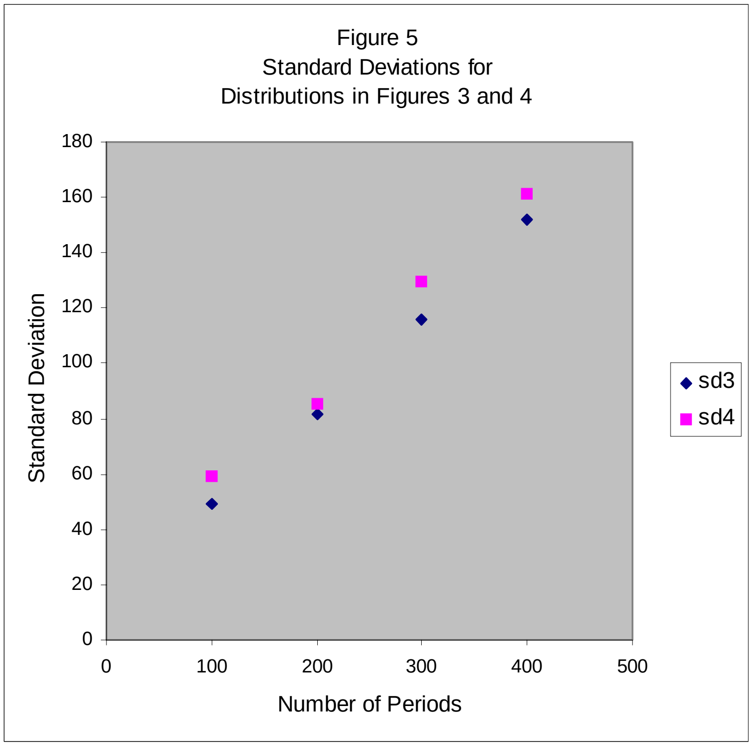
Figure 5
Standard Deviations for
Distributions in Figures 3 and 4
Standard Deviation
Number of Periods
sd3
sd4
180
160
140
120
100
80
60
40
20
0
0
100
200
300
400
500

An idea of how variances vary with the number of periods is given in Figure 6, in which the variances divided by the number of periods for the distributions represented in Figures 3 and 4 are plotted against the number of periods. From the upward slopes of these graphs -- actually, there is even a hint of upward convexity! -- we once again see no evidence of an eventual asymptote.

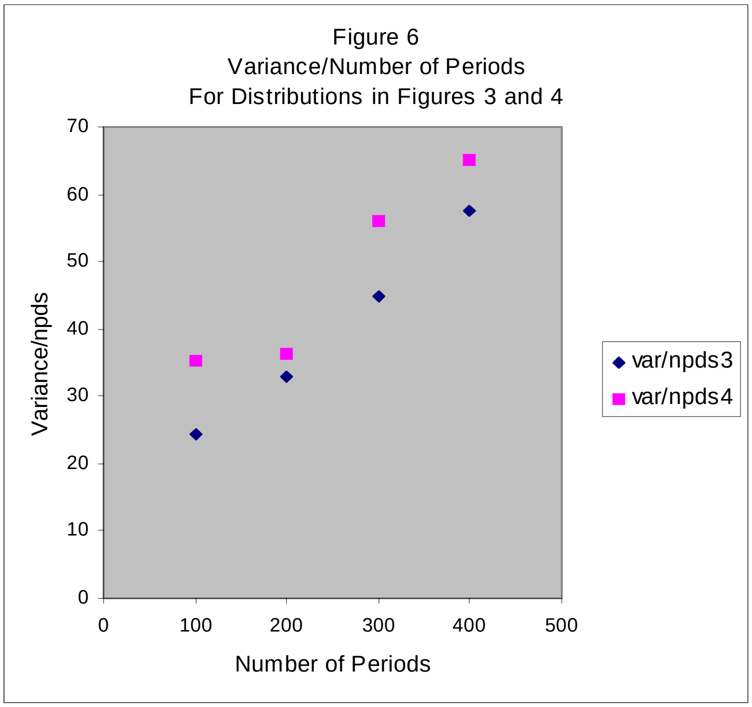

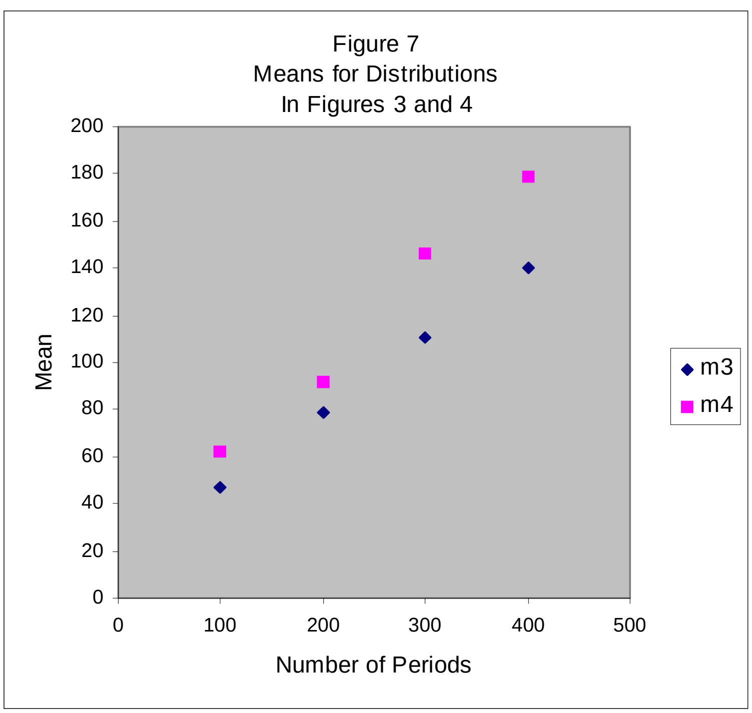

The rightward translation of the distributions represented in Figures 3 and 4 is obviously a reflection of increased mean holdings of wealth that occur with economic growth and heritability. This "aging" effect is quantified in Figure 7, which shows the means of the wealth distributions from Figures 3 and 4 plotted against the number of periods.

VI. Conclusions

The purpose of this note has been to investigate whether a Law-of-Pareto for the distributions of income and wealth can be obtained in the absence of assumptions concerning an initial Pareto "prime mover". Three scenarios involving highly stylized, artificial economies (all of which can be given a Darwinian interpretation) have been simulated under varying assumptions regarding heritability, distribution of talents, and stability of tastes. Simulations with the first two scenarios make it pretty clear that talent differentials and pure randomness of tastes cannot suffice to produce wealth distributions with sufficient thickness to be interesting. However, things change in the third scenario, in which there is an allowance for preference stability, in the sense that once an agent experiences a good, that good is consumed with a non-zero probability in subsequent periods (so long, of course, as the agent remains "alive"). What the results with the third scenario show is that, with strong preference stability and substantial productivity differences amongst agents, thick-tailed distributions of wealth can emerge that have certain Pareto features and are log-log translatable.

How Pareto-like these distributions actually are, on the other hand, is another matter. Simulations with a fourth scenario, which allows for a more realistic distribution of talents across agents (but which retains extreme "habit formation"), yields distributions that remain log-log translatable and "segment" scalable. However, the "Pareto" segment of these distributions now leave a tail that contains more than 30 percent of the mass of the distribution. Nevertheless, Mandelbrot-type exercises involving stability of variances suggest -- indeed, convincingly in my view -- that the distributions in question, whatever they might be, *do not have a finite variance*.

Obviously, the results of these exercises cannot settle the question of whether income and wealth distributions with a Law-of-Pareto upper tail can arise through normal birth, death, production, consumption, talent and taste differences, and heritability processes in the absence of an initial Pareto "prime mover". Nevertheless, the results are clear in suggesting that presence of the latter may not be necessary. While the exercises reported here embody what are thought to be plausible mechanisms concerning the distribution of talents and the formation of preference stability, heritability in the scenarios has been limited to agents keeping their stocks of wealth from one period to the next (so long, that is, as they survive). Heritability involving talents and tastes within a "family" is a scenario waiting to be explored. Adding such assumptions would seem almost certainly to increase the probability of "super large" fortunes, and could lead to a softening of the sharp "kinks" in the log-log graphs in Figures 3 and 4 that are evident at values of 1 - F(x) of about 0.20 and 0.30, respectively.[7] However, I feel that, when all is said and done, the strongest conclusion

[7] Also, as has been suggested to me by my colleague, Barbara Sands, another interesting exercise would be to eliminate the restriction that agents with negative wealth are removed the

arising from the exercises of this paper is the uncovering of what would appear to be a powerful link between large fortunes and strong preference stability. Preference stability is a feature of the real world,[8] as are large fortunes. A strong link between them would seem to open up an area of promising new research.

References

Mandelbrot, B. B. (1963), "New Methods in Statistical Economics," *Journal of Political Economy*, Vol. 71, pp. 421-440.

Mandelbrot, B. B. (1997), *Fractals and Scaling in Finance*, Springer-Verlag.

Taylor, L. D. (2003), "Stability of U.S. Consumption Expenditure Patterns: 1996 - 1999," Department of Economics, University of Arizona.

instant that their wealth turns negative. Allowing them to continue living, at least for a few periods, might give rise to the "lower hump" that is characteristic of real-world income and wealth distributions.

[8] Cf., Taylor (2003).