UCLA
Recent Work

Title

Investment Talent and the Pareto Wealth Distribution: An Experimental Analysis

Permalink

Authors

Levy, Shiki
Levy, Haim

Publication Date

1998-10-01

INVESTMENT TALENT AND THE PARETO WEALTH DISTRIBUTION:

AN EXPERIMENTAL ANALYSIS

Shiki Levy[*] and Haim Levy[**]

October 1998

[*] Anderson Graduate School of Management, University of California, Los Angeles, CA 90095-1481.
[**] Hebrew University, Jerusalem, Israel. The authors wish to thank Allon Cohen for his help and comments.

Abstract

It is a well-documented fact that wealth is distributed according to a power-law (Pareto) distribution at high wealth levels. Various models of wealth accumulation have been suggested in order to explain this empirical wealth distribution. Although these models are quite different one from the other, they are all based on a stochastic multiplicative process, and they all assume homogeneous talent: in these models the only source of inequality is the randomness of the process - luck! These models encounter two serious objections: a) it is claimed that the time it would take a stochastic homogeneous-talent process to generate the Pareto distribution is incredibly long, and b) many consider it unreasonable to assume homogeneous investment talent. Obviously, the provocative idea that inequality is primarily due to chance rather than talent has profound political, social, and philosophical implications. In this paper, we hope to shed some light on this controversy with evidence from a unique experiment in which the initial wealth of all participants is equal and real out-of-pocket money is involved. We find a convergence of the experimental wealth distribution to the Pareto distribution which is astonishing both in its speed (less than 10 trading rounds), and in its goodness-of-fit ($R^2 > 0.97$). Moreover, the Pareto parameter we find is similar in magnitude to the Pareto parameter of the actual wealth distribution in western countries. Analysis of the performance of the 63 participants in the experiment reveals that the differences in terminal wealth are primarily due to chance, rather than talent.

Introduction

At the end of the 19th century, the Italian-born Swiss economist, Vilfredo Pareto [1897] claimed that the population's wealth (and income) are distributed according to a particular functional form - a power function. The parameters of this distribution may change across societies, but regardless of the social or political conditions, taxation, etc., Pareto claimed that the wealth distribution obeys this general distribution law, which is named after him, Pareto's Law or Pareto's Distribution.

The discovery of a universal mathematical law for the distribution of wealth has lead to different theories about the origins of wealth inequality. Obviously, these theories have tremendous social and philosophical implications. The first to suggest an explanation for the Pareto distribution of wealth was Pareto himself [Pareto 1906]. Pareto suggested that the distribution of wealth corresponds to an underlying distribution of human abilities. However, Pareto has not offered a mathematical model that would explain the distribution of abilities and its relation to the Pareto Law. Pareto's explanation was advanced by Davis [1941] who introduced the "law of the distribution of special abilities" which asserts that the probability of an additional unit of ability was independent of the level of ability. This model, however, leads to a normal distribution of ability and therefore presumably to a normal, rather than Pareto, distribution of wealth. A different model for the distribution of ability was formulated by Boissevain [1939] who considered the distribution of abilities that could be represented as a product of several factors, each of which follows a binomial

1

distribution. Boissevain's model explains the positive skewness in the distributions of wealth and income, but leads to a log-normal distribution, not the empirically observed Pareto distribution.

The main models that offer an explanation for the precise form of the Pareto wealth distribution are the Markov chain model of Champernowne [1953], the stream model of Simon [1995] and the birth-and-death model of Wold and Whittle[1] [1957]. Although these models are quite different from each other in their details, they are all based on a stochastic multiplicative process of wealth accumulation, they impose some lower bound on wealth, and they all assume homogeneous wealth accumulation talent. Thus, in all the present models which can explain the empirical Pareto wealth distribution the **only** reason for inequality is the stochastic process − **chance**. In fact, S.Levy [1997] claims that homogeneous talent is a necessary condition if a stochastic wealth accumulation process is to lead to the Pareto distribution.

The above models for the process of wealth accumulation, and their provocative implication that chance, rather than talent, is the main reason for inequality at high wealth levels, encounter two major objections:

a)   It is claimed that the time it would take a stochastic homogeneous-talent process to generate the Pareto distribution is incredibly long (Shorrocks, [1973]). It is a natural question to ask: how long (or how many generations)

---

[1] For a historical review of the Pareto distribution see Persky [1992]. Review of models generating the Pareto distribution can be found in Steindl [1965], Arnold [1983], and Slottje [1989].

2

a "new" economy with a uniform distribution of wealth has to operate in order for Pareto's law to emerge. For example, suppose that there was perfect equality in the Soviet Union before it adopted a capitalist economy (which is not really the case), how long would we need to wait before the initial equal wealth distribution transforms to the Pareto distribution with the extreme inequality which it implies?

b)   Many find it unreasonable to assume homogeneous investment talent. Indeed, most money manager compensation schemes implicitly assume that performance is directly linked with investment talent - there is no point in performance-based compensation if the performace is primarily due to luck.

The purpose of this paper is four-fold:

1) To examine whether a Pareto wealth distribution is obtained when wealth is accumulated solely in the capital market, where at time t_0, the initial wealth is distributed uniformly.

2) To examine the speed of convergence to Pareto's Law if it, indeed, exists.

3) Examining the magnitude of the Pareto distribution parameters. We also examine the change in these parameters across time, emphasizing the direction and speed of the changes.

4) And finally, we test directly whether the inequality of terminal wealth is due to differences in investment talent across investors or to chance. Did the wealthier subjects in

3

the experiment invest wisely or were they simply lucky?

Of course, analyzing such issues and, in particular, having an initial uniform distribution of wealth is possible only in laboratory experiments, which is the framework of this paper. However, what is unique to this experiment, as explained in detail in Section I, is that the subjects participating could gain or lose out-of-pocket money, which makes the incentive atmosphere in our experiment very realistic.

The structure of this paper is as follows. In section I, we describe the experiment. In section II we provide the main findings. We find that some investors accumulated a great deal of wealth while some remained "poor", so to speak. However, the striking result is that we find that this wealth distribution is mainly due to pure luck rather than to investment talent. We find both direct and indirect evidence for this result. After only a few rounds of trade, we find that the wealth distribution is in excellent agreement with Pareto's Law, even though all subjects start the investment process with an equall amount of wealth. As the parameters of the returns in the experiment fit closely annual mean returns and annual variances of returns, this implies that after only 11 years almost a precise Pareto distribution is obtained, in spite of the fact that the initial distribution of wealth is uniform. Thus, there is no need for many generations of wealth accumulation and differentiation in order to create the Pareto wealth distribution and the rather extreme inequality which it implies. The fact that wealth inequality is due primarily to luck, rather than to investment talent, is quite a surprising result which has strong implications for compensensation of mutual fund managers and for politicians contemplating various

alternative taxation systems.

Section III concludes the paper with a discussion of the main results and their implications.

I. The Experiment

In this experiment subjects could invest in a portfolio selected out of 20 available stocks, and could borrow or lend as much as they wished. There were 10 trading rounds, and in each round the market price of each stock was determined by supply and demand, exactly as in a market place with limit orders. After the trade, the book value of the firm changes randomly, reflecting the firm profit or loss from operation. Thus, as will be elaborated on below, we distinguish in this experiment between the market value of the firm and its book value. In each round these two values may differ exactly as occurs with closed-end mutual funds. However, after the 10th trading round, all the firms liquidate their assets and the market values are equal to the corresponding book values. Subjects trade at the beginning of the period (year), and at the end of the year (which is a week in the experiment), the firm reveals its profit or loss from operation (book value). Thus, we have 10 trading rounds and 10 book values reported. As only at the end of the last period (the liquidation date), the book value is equal to the market value, we record 11 market values of the subjects' wealth: 10 at the beginning of each period, and one at the end of the last period.

The subjects in the experiment were first year MBA students from the Hebrew University of Jerusalem who were studying a finance course. The experiment was not

mandatory, as loss of real out-of-pocket money was possible. Out of a potential of 67 students 63 participated in the experiment. The subjects could profit or lose money, depending on the realized returns on the investments which they selected. As they could lose out-of-pocket money, the subjects were warned at the beginning of the experiment to check their savings and other income resources to make sure they had enough money to cover potential losses. Making the experimental setup realistic is very important, as it has been shown that subjects behave completely differently in experiments in which only gains are possible, as apposed to experiments in which subjects may also incur losses[2]. The subjects reported their wealth and income to the experiment's manager at the beginning of the experiment. Most of the subjects were about 25 years old and most were employed either part or full time.

Each subject was given an initial investment allotment of $30,000 "paper" money, and was offered to buy stocks of 20 pure equity firms reported in Table 1. During the ten rounds of the experiment the subjects traded in the stock of these firms. As explained before, at the end of the tenth round all the firms liquidated their assets and divided them among the stockholders. The liquidation value was determined by the book value of each firm at the end of the tenth period, which we call the 11[th] round. Because the book value determines the terminal wealth (and hence the financial reward) of the subjects, we first turn to explain how the book values were determined.

---

[2] Creating a situation where losses are possible greatly affects the subjects' behavior. For example, when losses are not possible, subjects tend to take leverage of hundreds of percents (see Kroll, Levy and Rapoport [1988]), but when losses are possible subjects become on average net lenders (see Kroll and Levy [1992]).

The book value of each firm's assets at the beginning of the experiment is shown in the right hand column of Table 1. The book value of the assets at time t_0 was determined by the experiment manager such that if all investors choose their portfolios by the Mean-Variance rule and diversify optimally, the market is in equilibrium at time t_0 (i.e., the total demand for each stock by all participants, with zero net borrowing, is exactly equal to the book value reported in Table 1). However, this fact was not known to the subjects. Moreover, in this stage of their study, the subjects did not learn yet the notion of optimal diversification and the optimization methods to find the mean-variance efficient portfolios.

While the first round book-value determination was not crucial for the experiment, this procedure guarantees that at least the first trade does not start with a severe disequilibrium; if all investors choose their portfolio by the Mean-Variance rule, the equilibrium price would be equal to the book value at time t_0, and the market would be cleared[3].

The method of determining the initial value of the firm was unknown to the subjects. However, the information given below, as well as the trade procedure, was fully known to the subjects and to the experiment manager[4]. Any information which was not known in advance to the subjects (e.g., the future buy-sell orders for a given

_____

[3] On the mean-variance efficient choices and equilibrium prices with a riskless asset see Markowitz [1952], Tobin [1958], Sharpe [1964] and Lintner [1965]. The above procedure for the book-value determination was also necessary for testing the CAPM with *ex-ante* parameters, a topic to which a separate article is devoted, see H. Levy [1997].
[4] On the role of public information in explaining laboratory asset market data, see Copeland and Friedman [1991] and Smith [1962].

stock), was also unknown to the experiment manager.

In each subsequent period, the book value of the firm's asset either grew or declined at random. The random variable which determines the book value was taken from a normal distribution with the corresponding mean and variance reported in Table 1. For example, denoting the book value of the i^{th} firm at the end of period t by $V_{i,B}^{t}$, the book value of firm 1 at the end of period 1 is given by $V_{1,B}^{1}$, where

$$V_{1,B}^{1} = V_{1,B}^{0}(1 + \tilde{R}_{1,B}^{1}) = \$78,930(1 + \tilde{R}_{1,B}^{1})$$

In this example i=1 indicates that we are dealing with the firm 1, Elite (see Table 1). The $78,930 is the book value of Elite at t_0 as reported in Table 1. $\tilde{R}_{1,B}^{1}$ is a random variable drawn in the first period form a normal distribution with a mean of 3% and standard deviation of 4%, which reflects the operating earnings that firm 1 makes on its assets. The subscript B indicates book value, namely, these are the rates of return which the firm earns on its assets. Similarly, at the end of the second period, the book-value of the assets of the first firm grows to $\$78,930(1 + \tilde{R}_{1,B}^{1})(1 + \tilde{R}_{1,B}^{2})$, where $\tilde{R}_{1,B}^{2}$ is a random variable drawn in the second period from the same normal distribution. Since there are ten periods (or trading rounds) in the experiment, the i^{th} firm's asset value at the end of the tenth period is given by

$$V_{i,B}^{10} = V_{i,B}^{0} \prod_{t=1}^{10} (1 + \tilde{R}_{i,B}^{t}),$$

where $\tilde{R}_{i,B}^{t}$ is the random variable corresponding to firm i in period t. Since all $\tilde{R}_{i,B}^{t}$

have positive means, it is likely that after ten rounds $V_{i,B}^{10} > V_{i,B}^0$. Indeed, this occurred in 19 out of the 20 firms.

To facilitate the subject's portfolio choice, all the random variables ($\tilde{R}_{i,B}^t$, $\tilde{R}_{j,B}^t$) are pairwise independent (zero correlation), and each random variable ($\tilde{R}_{i,B}^t$) is independent over time (this information was known to the subjects). To avoid differences between accounting profit and economic profit, the subjects were told that $\tilde{R}_{i,B}^t$ is the *cash* rate of return of the firm on the firm's assets and <u>not</u> an accounting return. This makes the book value at the liquidation date equal to market value (cash distributed to the subjects). It was also fully known that none of the firms would liquidate before the end of the experiment.

The subjects traded the stocks of the firms, and in each trading round determined the *market* value of the firms, which could be different from their book values. The subjects were told that at the end of the tenth round of trading all firms would liquidate and that the liquidation value would be $V_{i,B}^{10}$, namely, equal to the book value at that time. The liquidation value of each firm would then be distributed to the subjects in accordance to the proportion of shares they hold of each particular firm. Thus, this scenario is very similar to a trade in *term trusts*[5] (shares of a closed-end mutual fund with a given predetermined terminal liquidation date) which can exhibit an observable premium or discount between their market value and the total value of

---

[5] For more on term trusts, see for example: "The Terminator" by Mary Beth Grown, Forbes, Aug 11, 1997, page 130.

their assets.

The risk-free interest rate was r=2%, and there were no constraints on borrowing or lending. The subject could put all his/her money in the riskless assets and earn a sure profit of 2% per period. Moreover, he/she could switch during any trade round from risky assets to the riskless asset to ensure that he/she would not lose any profit already made on the stock market. On the other hand, subjects were free at any time to borrow as much as they wanted at the riskless interest rate and to invest the borrowed money in the risky assets.

The book value was reported to the subjects after they traded in the stock. Thus, we assume that trading takes place at the beginning of the year, and at the end of the year the firm reveals its profit or loss from operation. Therefore, for each firm we have 10 market values and 10 book values corresponding to the 10 trading rounds, and a liquidation value which is the last book value reported (at the beginning of the 10^{th} year we have a book value and a market value for each firm. At the end of this year, market value = liquidation value = book value, because the firm's profits are in cash and not accounting profits). Thus, we have 11 points of time where the subject's wealth can be measured.

The profit or loss of each subject at the termination date was determined as follows. Each subject received at time t_0 $30,000 in "paper" money with which he/she can buy stocks or invest in the riskless asset. If the subject does not go bankrupt during the experiment, at the end of the tenth round the k^{th} investor's wealth

is given by W_k^{10},

$$W_k^{10} = \sum_{i=1}^{20} (N_{i,k} / N_i) V_{i,B}^{10} - B_k^9 (1+r)$$

where

W_k^{10} is the wealth of the k^{th} investor at the end of the tenth round,

$V_{i,B}^{10}$ is the liquidation value of the i^{th} firm at the end of the tenth period,

N_i is the number of shares issued by the i^{th} firm,

$N_{i,k}$ is the number of shares of the i^{th} firm held by the k^{th} investor.

B_k^9 is the amount of money the k^{th} investor borrowed at the end of the ninth

trading round. Thus, $B_k^9(1+r)$, where r denotes the interest rate, is the amount of

money the subject should return to the bank at the end of the 10^{th} trading round. Note

that if the subject lends money $B_k^9 < 0$; hence $-B_k^9(1+r) > 0$ (and money is received

from the bank).

To calculate the actual dollar reward of each subject, each $1,000 in "paper"

money represented $1 in actual money; hence the actual financial reward of the k^{th}

subject at the end of tenth round is R_k,

$R_k = W_k^{10} / 1,000$

At the end of *each* trading round, the net market value of the assets of each

subject is examined. If it is negative, the subject goes bankrupt and pays out of his/her

pocket money to the experiment manger. The actual dollar penalty paid by the

investor should he/she go bankrupt at period t is equal to the amount of his/her

(negative) net market value of assets divided by 1,000:

$$W_k^t / 1,000 = [V_{k,M}^t - B_k^{t-1}(1+r)] / 1,000$$

where W_k^t is the market value of the k^{th} investor's wealth at the end of round t; $V_{k,M}^t$ is the market value of the stocks held by the k^{th} investor (as determined by the demand and supply of the subjects); and B_k^{t-1} is the amount borrowed at period t-1. For example, if the market value of the stocks held at the end of the fifth round is $100,000 and his/her outstanding borrowing is $B_k^{t-1}(1+r) = \$500,000$, he/she pays out of his/her pocket ($100,000-$500,000)/1,000=$400. When a subject goes bankrupt, he or she is eliminated from the experiment for the rest of the trading rounds. Thus, a penalty is paid by the subject whenever the net wealth is negative. However, a reward is paid only at the end of the tenth trading round. There are no transaction costs[6] on trading in the securities, and short sales were not allowed.

Market prices are determined at each trading round by demand and supply for the stocks. The subjects were allowed to submit buy and sell limit orders for stocks, and the equilibrium price was determined by the intersection of the supply and demand aggregate curves. (For more details on the type of orders and the mechanism of market price determination see Appendix A.) At the end of each trading round, each investor receives information on his/her portfolio composition and his/her net wealth. Each firm's new book value, new stock price, and the number of its shares traded, was

---

[6] On considerations for including or excluding transaction costs from an experimental market see Plott and Smith [1978] and Forsythe, Palfrey and Plott [1982].

also provided after each round as public information. Subjects however did not have direct access to the composition of other subjects portfolios (for a summary of the information and instructions provided to the subjects see Appendix B).

The experiment was designed to create an atmosphere close to that of real security markets. The main contributors to this atmosphere were the ability to gain and lose real money, and the investment over an extended period of time with many investment sub-periods.

II. <u>The Results</u>

In part (a) of this section we test the goodness-of-fit of the wealth distribution obtained in the experiment to the Pareto distribution. We estimate the value of the Pareto parameter, α, and we measure the time it takes the distribution to converge to the Pareto distribution. In part (b) of this section we test whether the main cause for inequality is differences in investment talent across investors, or simply luck.

a. <u>Goodness-of-fit to the Pareto Distribution</u>

The Pareto distribution can be written in various forms. Originally, it was stated as follows (see Johnson & Kotz [1970], Chapter 19):

(1) $P(x) = Pr(X \geq x) = \left(\dfrac{k}{x}\right)^{\alpha}$ where $k > 0$, $\alpha > 0$, and x k,

and $P(x)$ is the probability of the wealth being equal or greater than x. It can be easily shown that the cumulative distribution of wealth which follows from (1) is,

(2) $\quad F(x) = 1 - (\dfrac{k}{x})^{\alpha}$, where $k > 0$, $\alpha > 0$, $x \geq k$.

13

with a density function f(x) given by:

(3) $\quad f(x) = \dfrac{\alpha k^{\alpha}}{x^{\alpha+1}}, \ (\alpha > 0, \ x > k > 0)$

Pareto's Law can be formulated also as,

(4) $\quad n = A \, x^{-\alpha}$

where A is a constant, and n is the number of persons having wealth of x *or more*. Thus, n is the investor's rank by his/her wealth; the larger the wealth the smaller the rank.

Equation (4) can be written also as

(4′) $\quad x(n) = C n^{-1/\alpha}$

where x(n) is the wealth corresponding to the n ranked individual and $C = A^{1/\alpha}$

For these definitions and further analysis, see Johnson & Kotz [1970]. Showing the equivalence between (3) and (4′) is given in Appendix C.

Taking logarithms of both sides of equation (4′) yields:

(5) $\quad \ln(x(n)) = \ln(C) - (1/\alpha) \ln(n)$

We employ in this paper equation (5) to test whether Pareto's Law prevails with the subjects wealth as obtained in our experiment. After each round of trade we measure the wealth of each subject, which is given by the market value of his/her portfolio. Then, we rank all the subjects by their wealth. Having the pairs (x(n), n), i.e., the wealth of each subject and the corresponding rank n of this subject, we run the following regression, corresponding to eq.(5):

(6) $\quad \ln(x(n))_i = \hat{a} + \hat{b} \ln(n)_i + e_i$

14

If the subject's wealth is distributed by Pareto's Law, we expect to find such a linear fit (see eq.(5)). In order to examine Pareto's Law we analyze the data on the wealth of each subject as measured at the end of each trading round, as well as the wealth at the 11[th] round, i.e., at the liquidation date.

At the end of the experiment all the firms were liquidated and each subject received a reward according to his/her accumulated wealth. There were no bankruptcies in the expeirment[7]. Therefore, each subject received a positive financial reward at the end of the experiment, and no one paid money to the experiment's manager. There was a big dispersion of the reward. To illustrate, the highest reward of a subject was $547 (corresponding to $547,000 wealth) and the lowest was $33 (corresponding to $33,000 wealth).

Figure 1 illustrates the wealth distribution historgram at the liquidation date. Though the historgram intervals are determined arbitrarily ($10,000 width), the observed long right tail distribution is quite obvious; the wealth distribution is not symmetrical and seems to fit nicely the Pareto distribution and may be also the lognormal distribution. However, these assertions have to be statistically tested. Therefore, we next test the relationship between the Pareto distribution and the wealth distribution as obtained in the experiment. We first employ the regression suggested

---

[7] The reason for this fact is that the subjects were very careful in their borrowing policy. In fact, on average they were net lenders. In other experiments where no loss was possible, the subjects borrowed hundred percent of their wealth (see Kroll, Levy and Rapoport [1988]). This emphasizes the importance of creating realistic conditions in the experiment. On the importance of the subjects having a "stake" in the experiment see also Plott [1986] and Smith [1991].

by equation (6) to test whether wealth distribution obeys the Pareto Law. We run 11 separate regressions, based on the wealth data as measured at the end of the ten trading rounds as well as the liquidation date (the 11[th] round). Table 2 reports the coefficient $\hat{b}$, the T-values of the coefficient as well as the coefficient of correlation R^2, of these regression lines.

The fit of wealth distribution to the Pareto distribution is striking and in the last few trading rounds it is almost perfect. The slope is, of course, negative (as expected, see eqs. (5,6)), in all 11 rounds. The T-values range from -3.73 (in the first round) to about -47 in the last two rounds. After only three trading rounds, R^2 is greater than 90%. In rounds 4-11 the R^2 is very close to unity, indicating almost a perfect fit to the Pareto distribution after only 4 rounds! Figure 2 illustrates the regression line (corresponding to equation 6) for the wealth distribution at the liquidation date, round 11. The regression lines for the first 10 rounds are reported in Appendix D. Note that the horizontal axis measures the subject rank (1-63) by his/her wealth, and as there are 63 observations and it is measured in $\log_{10}(\cdot)$ terms, the number ranges from zero (log 1 for the wealthiest subject) to 1.799 (which is log 63, the rank of the poorest subject). The horizontal axis range is, therefore, the same for all trading rounds. The vertical axis changes from one trading round to another because wealth changes across trading rounds. For example, in the last round, the wealthiest subject had $547,000 (or $547 in real money) with log (547,000) $\cong$ 5.73 and the wealth of the poorest subject was $33,000, with log (33,000) $\cong$ 4.518.

Contrasting equations (5) and (6) we see that $\hat{b}$ is an estimate of $-1/\alpha$. Therefore, to obtain the estimate of $\hat{\alpha}$ we need to calculate $-1/\hat{b}$ where $\hat{b}$ are given in Table 2. Thus, for round 1, $\hat{\alpha} = - (1/ - 0.019) = 52.63$, for round 2, it is $\hat{\alpha} = -(1/ -0.218) = 4.59$, etc. Note that $\hat{\alpha}$ tends to decrease as we advance in the trading ranks and it is 1.85 and 1.78 corresponding to the 10th and 11th rounds, (i.e. $-(1/-0.540)$ and $-(1/-0.563)$, respectively). It is interesting to note that $\hat{\alpha}$ corresponding to the U.S. wealth distribution is estimated to be 1.35, 1.06 for the U.K. and 1.83 in France (Levy 1998). Our $\hat{\alpha}$ estimates are a little higher than those of the U.S., U.K and similar to the France coefficient. As $\hat{\alpha}$ shows a tendency to decline and to converge in our experiment, one may suspect that the experiment's $\hat{\alpha}$ values tend to converge to a value lower than 1.78, closer to those $\hat{\alpha}$'s observed in the U.S. and U.K .

We next compare the theoretical Pareto distribution with the experimental cumulative distribution. To draw the theoretical Pareto distribution, we need first to estimate $\hat{\alpha}$ and $\hat{k}$. However, by equation (2) we have,

(7) $\quad 1 - F(x) = (\dfrac{k}{x})^{\alpha}$

Hence,

(8) $\quad \text{Log} [1 - F(x)] = \alpha \, \log k - \alpha \, \log x$

As 1-F(x) and log x can be estimated from the experiment data (or empirically) one can run a simple regression (treating $\alpha \, \log k$ as constant) to obtain an estimate for α :

17

$$(9) \quad \hat{\alpha} = \frac{-N\sum_{i=1}^{n}\log x_i \log[1-F(x_i)] + (\sum_{i=1}^{n}\log x_i)(\sum_{i=1}^{n}\log[1-F(x_i)])}{N\sum_{i=1}^{n}(\log x_i)^2 - (\sum_{i=1}^{n}\log x_i)^2}$$

(see Johnson and Kotz, p. 235)

Having the estimate of $\hat{\alpha}$ and the arithmetic mean values of the dependent and the independent variables, we estimate k by estimating $\alpha \log k$ from eq.(8) and α from eq.(9). Using the wealth at each round we can employ the above regression to estimate $\hat{\alpha}$ and $\hat{k}$ for each trading round. Figure 3 illustrates the theoretical Pareto distribution and the experimental cumulative distributions corresponding to the 11[th] round with $\hat{\alpha} \cong 1.728$ and $\hat{k} \cong 35,225$ as estimated by equation (9) above. Note that once again, with the different estimation method (see eq. (9)) we obtain a similar value of $\hat{\alpha}$ for the wealth distribution in the last round (compare with the value of $\hat{\alpha} \cong 1.776$ obtained by the rank-wealth estimation method from the coefficient $\hat{b}$ of -0.563, as reported in Table 2). Again, the value obtained in the experiment is larger than the corresponding $\hat{\alpha}$ for the U.S. and U.K. and smaller than the $\hat{\alpha}$ corresponding to France. The theoretical Pareto distribution given by the solid curve in Figure 3 and the experimental distribution are extremely close to each other (see Figure 3). To test the hypothesis that the wealth distribution is not significantly different from Pareto distribution we use the Kolmogrov-Smirnov statistic given by:

$$D = \text{Max} [F_{Ex}(x) - F_{Th}(x)] = 0.10567$$

where $F_{Ex}(x)$ and $F_{Th}(x)$ are the experimental and the theoretical distributions,

respectively. With n=63 observations, the 20% critical value is $1.07 / \sqrt{63} \cong 0.1348$.

Thus, we obtain a smaller D value than this critical value, hence the null hypothesis cannot be rejected even at a relatively high significant level. The Pareto and the experimental distributions are so close to each other such that even at a significant level which exceeds 20%, we cannot reject the null hypothesis asserting that the two distributions are not equal. (For comparison, the null hypothesis asserting that the experimental distribution is a lognormal distribution leads to a D value of 0.153, and is rejected at a 10% significance level). This confirms our previous results revealing an excellent fit between the experimental wealth distribution and the Pareto distribution.

b. Investment Talent

An important implication of this study relates to the distribution of investment talent across investors. How much of the wealth inequality is due to differential investment talent, and how much is due to luck? Levy [1997] claims that differential investment talent leads to a distribution of wealth which is different than the Pareto distribution. He suggests that a Pareto wealth distribution can evolve only when the main factor driving the inequality is chance (which is indeed the case in all the models generating the Pareto distribution, see Champermowne [1953], Simon [1995], Wold and Whittle [1957], and Levy [1997]). Thus, the fact that we find such a striking fit between the experimental wealth distribution and the Pareto distribution constitutes indirect

evidence that luck is the main force driving the investment success; i.e., the wealth distribution inequality.

The question whether there is a significant differentiation in investment talent across investors is related to the issue of market efficiency. If the market is efficient, no matter how talented the investor is at analyzing the available information, he/she will not be able to achieve abnormal returns. Though most empirical studies support the Efficient Market Hypothesis there is no conclusive evidence regarding this issue. For evidence showing that the market is efficient, see Fama [1970], [1991]. However, several market "anomalies" have been observed, which suggest that company specifics such as size or book-to-market ratio (Fama and French [1992]) or the stock's past performance (Jegadeesh and Titman [1993]) can be used in order to obtain abnormal returns. This implies that the market is inefficient and that talented investors may be able to outperform by exploiting these anomalies. Moreover, even if the standard market efficiency tests reveal that the market is efficient, we still can not safely conclude that there are no talented investors which can "beat the market" systematically, because it is possible that such investors use complex investment methods which are not tested for in the standard EMH tests. In other words, the factors usually employed in the EMH tests (e.g. book-to-market, P/E ratios, autocorrelations, etc.) may reveal market efficiency, while talented investors employ more sophisticated methods to obtain abnormal returns.

In order to reach a more definitive conclusion regarding the existence and degree of investment talent differentiation one can employ a more direct approach. Sharpe (1966)

measures directly whether the performance of mutual funds in one period are indicative of the next period's performance (as we would expect if performance is driven by talent rather than luck). He measures the reward to volatility of mutual funds in two consecutive five-year periods. Sharpe finds a slight positive relationship between the performance in the two periods (R=0.36). This indicates that if differential investment talent exists, it does not seem to be a very dominant factor in explaining mutual fund performance. Recently, Beckers [1997] finds that the major factor explaining mutual fund performance is luck. Samuelson [1989] summarizes the empirical evidence on the performance of money managers:

> *"Those lucky money managers who happen in any period to beat the comprehensive averages in total return seem primarily to have been merely lucky. Being on the honor roll in 1974 does not make you appreciably more likely to be on the 1975 honor roll."* (Samuelson [1989] p. 4).

The studies cited above analyze the role of investment talent in the investment performance of mutual funds. A parallel direct analysis of the role of talent versus the role of luck in personal investment is generally very difficult, because the data regarding individuals' portfolios is usually not available. In the framework of our experiment, we have a unique opportunity to perform such a direct analysis. The data we recorded includes not only the wealth of each individual at the end of each trading round, but also the composition of his/her portfolio, and the rate of return achieved by each investor at each round. In what follows we describe the method employ in order to test directly for differential investment talent. The idea is simple: if differential talent exists, success in one

round should (on average) predict success in future rounds. In this case, we would expect a positive autocorrelation of investment performance. However, if it is all luck, success in one period has no predictive power whatsoever regarding future performance.

In each trading round we measure the rate of return obtained on the portfolio of the k^{th} investor, denoted by R_k. Then, we run the regression:

(10) $R_{k,L} = a + b\, R_{k,M} + e_k$

where $k = 1, 2, \ldots 63$, and $L \neq M$ are two trading rounds. If investors have different talents, we expect b to be significantly positive. That is, given that an investor has exceptionally good investment talent with a relatively high rate of return in round M implies that a relatively high rate of return is expected also in round L. By this argument, if differential investment talent exists, we expect b to be positive. If b is negative, it implies that a success in round M predicts a failure in round L. If b is not significantly different than zero we tend to conclude that success is due to pure luck.

Table 3 reports the regression coefficients for all pairs $L \neq M$ of regression (10). Table 4 provides the corresponding correlations, R. The interpretation of these two tables is a little tricky: there are more positive significant slopes b than negative significant slopes. Out of the 110 T-values reported in Table 3, there are 30 significant positive values (at a 5% level), 17 significant negative values, and 63 non-significant values. Thus, though for some rounds talent seems to exist, over all the experiment, it seems that the wealth was determined more by luck than by investment talent, because there are 80 out of 110 coefficients which are either not significantly different from zero or significantly negative.

Also, in some rounds of our experiment, there are significant positive slopes followed by significantly negative slopes (or vice versa).

Table 4 complements table 3 by showing the relevant correlations, R, for all rounds $L \neq M$. As we can see from the table, the correlations are relatively low, with the lowest one (-86%) corresponding to round 1 and 2, and the highest on (R=69%) corresponding to rounds 8 and 9. The bottom line of Table 3 and 4 gives the averages. Though we do not attribute any economic meanings to these averages, they indicate that there is not even one trading round with dominating large T-values.

The explanation for the positive slope is mainly due to different leverage policy rather than to stock selection ability. To illustrate, suppose that investors are ranked by their proportion of investment in the riskless asset (which can be positive or negative). Then, if the stock market goes up, say in round 3 and 6, then we expect to have a positive slope corresponding to these two rounds due to the leverage effect, even though there is no stock selectivity ability at all. As the stock market (as well as in our experiment) generally goes up more periods than it goes down, we should expect to find more positive (and significant) slopes b, than negative slopes. This effect is also expected to be present in the actual financial markets. If there is no significant differentiation in investment talent, and the stock market goes up more years than it goes down, we should expect to find more positive slopes than negative slopes because of the leverage effect.

In order to control for the differential leverage effect, we follow Sharpe's [1966] test. We calculate the reward to volatility ratio for the first 5 periods and for the next 6

periods. We run the regression:

$$(11) \quad (R/V)_{k,2} = a + b \ (R/V)_{k,1} + e_k$$

where $k = 1,2,...63$, and 1 and 2 stand for the two periods. Like Sharpe we find a slight positive correlation with R=0.22, which is non-significant (t=1.72).

To sum up, it seems that the general characteristic of this experiment is that luck rather than investment talent is the dominating factor in the wealth accumulation process. This conclusion may seem surprising at first. On the other hand, it also seems to be a natural consequence of the Efficient Market Hypothesis. If the market is efficient, no rates of return prediction and no abnormal investment capabilities should be possible.

III. Discussion

Some researchers claim that it would take an enormous amount of time for a Pareto wealth distribution to evolve from a stochastic investment process (Shorrocks [1973]). Our results are striking: we start with perfect equality (a uniform distribution of wealth where each investor has \$30,000) and after only a few trading rounds (years) we obtain an extremely good fit to the Pareto distribution. From the 9[th] trading round (9[th] year) onwards, the fit is almost perfect with more than 97% coefficient of correlation, R^2. Moreover, the parameter of the Pareto distribution, α, obtained in the experiment is similar in magnitude to the estimated parameter of the real wealth distribution in the U.S., U.K. and France. Thus, even if one distributes wealth uniformly among investors, it takes only about one

24

decade for the Pareto wealth distribution, and the strong inequality which it implies, to evolve. This leads one to suspect that the rather extreme inequality in modern western society is a very fundamental and robust outcome of the nature of the investment process. This inequality can emerge within several years - it is not necessarily the accumulative result of many generations of wealth accumulation.

As heterogeneous investment talent is claimed to be inconsistent with the Pareto Law, the finding of a Pareto wealth distribution constitutes indirect evidence that luck, rather than investment talent, is the main reason for the observed wealth distribution inequality. Our experimental framework gives us a unique opportunity to go further and test this hypothesis directly. If talent plays a significant role in the investment process, one would expect good performance in one period to be indicative (on average) of high investment talent, and therefore to be followed by more good performance in another period (again, on average). In our experiment we find that past performance generally does not help predict future performance. This leads to the conclusion that investment performance, and the resulting observed wealth distribution and inequality, are primarily due to luck, not to talent. This conclusion conforms with the Efficient Market Hypothesis.