

The consumption of stockholders and nonstockholders*

N. Gregory Mankiw

Harvard University, Cambridge, MA 02138, USA

Stephen P. Zeldes

University of Pennsylvania, Philadelphia, PA 19104-6367, USA

Only one-fourth of U.S. families own stock. This paper examines whether the consumption of stockholders differs from the consumption of nonstockholders and, if so, whether these differences help explain the empirical failures of the consumption-based CAPM. Household panel data are used to construct time series on the consumption of each group. The results indicate that the consumption of stockholders is more volatile and more highly correlated with the excess return on the stock market. These differences help explain the size of the equity premium, although they do not fully resolve the equity premium puzzle.

1. Introduction

Over the past two decades much research has attempted to build and examine models linking the consumption decision and the portfolio allocation decision. Contributions by Merton (1973), Lucas (1978), Breeden (1979, 1986), and Grossman and Shiller (1982) provided the theoretical underpinning of the consumption-based capital asset pricing model. Many empirical studies testing this model quickly followed. Unfortunately, the

*Mankiw is grateful for financial support from the National Science Foundation. Zeldes is grateful for financial support from the John M. Olin Fellowship at the National Bureau of Economic Research, the Wharton Junior Faculty Research Fund, and the Alfred P. Sloan Foundation. The authors thank Karen Dynan, Ahmed Taha, and Andrea Wicks for research assistance, and Robert Barro, Mariorie Flavin, Amy Salsbury, an anonymous referee, and Jerold Warner (the editor) for helpful comments. The PSID data used in this paper were made available by the Inter-University Consortium for Political and Social Research. The original source, the collectors of the data, and the Consortium bear no responsibility for the analyses or interpretations presented here.

Published in: Journal of Financial Economics

weight of the available evidence is that the standard model appears not to describe adequately the data on consumption and stock returns.

One of the most prominent empirical failures of the model is the equity premium puzzle. Over the past hundred years, the return on equity has averaged about 6% more than the return on short-term Treasury bills. Mehra and Prescott (1985) show that this equity premium is too large to be explained by a standard general equilibrium asset pricing model. More generally, the puzzle can be seen by examining the first-order conditions that arise in almost any asset pricing model that relates consumption and asset returns. A number of explanations of the equity premium puzzle have been proposed, but none has fully resolved the puzzle.¹

The source of the puzzle is that aggregate consumption growth covaries too little with the return on equities to justify the large observed risk premium on stocks. As a result, implausibly high levels of risk aversion are required to rationalize the size of the equity premium. Intuitively, if the random movements in stock returns are not associated with large changes in consumption, the randomness does not represent true riskiness to the consumer and therefore should not require a large risk premium.

An objection to the empirical work on consumption-based asset pricing models, including work on the equity premium, is that it relies on consumption data aggregated across stockholding and nonstockholding families. Estimates we present below indicate that about three-fourths of U.S. families hold no stock. These nonstockholding consumers are unlikely to satisfy the first-order conditions for the optimal holding of assets that underlie the consumption CAPM. Unless the consumption of stockholders fortuitously moves together with the consumption of nonstockholders, the standard practice of testing the model with data on aggregate consumption is inappropriate.

This paper is the first attempt to examine empirically the hypothesis that the consumption of stockholders differs from the consumption of nonstockholders. We use data from the Panel Study of Income Dynamics to construct a time series of the consumption of stockholders and a time series of the consumption of nonstockholders. The data have some serious limitations: substantial measurement error, a relatively short time series, and the availability of only food consumption. Despite these shortcomings, the evidence indicates that stockholders and nonstockholders differ substantially. In particular, stockholders' consumption is more volatile and more highly corre-

¹For recent work attempting to resolve the equity premium puzzle based on a representative consumer model, see, for example, Abel (1990), Cecchetti, Lam, and Mark (1989), Constantinides (1990), Kandel and Stambaugh (1990), Kocherlakota (1987), Rietz (1988) [and the response by Mehra and Prescott (1988)], and Weil (1989). For work on the equity premium puzzle based on individual heterogeneity, see Abel (1989), Ben-Zvi and Sussman (1988), Kahn (1988), and Mankiw (1986).

lated with the stock market. Although our data cannot provide a complete resolution of the equity premium puzzle, our findings suggest that the distinction between stockholders and nonstockholders may be crucial to an ultimate resolution of this puzzle and other asset pricing anomalies. Addressing this issue more fully will require better data on the consumption of stockholders.

The paper proceeds as follows. We begin in section 2 by summarizing some evidence on the frequency of stockholding and the demographic characteristics of stockholders and nonstockholders. We then review in section 3 the consumption-based capital asset pricing model. Focusing on the Euler equation relating consumption growth and asset returns, we discuss why the equity premium of 6% is puzzling. We argue that failures of the consumption CAPM might be rationalized by a model with two groups of consumers: stockholders and nonstockholders. In section 4 we examine the volatility of consumption and the correlation of consumption growth with stock returns for stockholders and nonstockholders. In section 5 we present conclusions and suggestions for further research.

2. Who holds stock?

We use data on a representative sample of families from the Panel Study of Income Dynamics (PSID). The 1984 survey included, for the first time, questions about the size and allocation of each family's financial wealth. One question asked for the current market value of 'shares of stock in publicly held corporations, mutual funds, or investment trusts, including stocks in IRA's'. Another question asked for the amount of money in 'checking or savings accounts, money market funds, certificates of deposit, government savings bonds, or Treasury bills, including IRA's'.

Table 1 presents the distribution of holdings of stocks and other liquid assets based on these questions from the 1984 survey. Of the total sample of 2,998 families, 27.6% hold a positive amount of their wealth in stocks, and 72.4% own no stock. The families that do not hold stock earn 62% of disposable income, account for 68% of food expenditure, and own 34% of total liquid assets (including stocks). Some stockholders own small amounts of stock. Only 23.2% of the sample holds equity in excess of \$1,000, and only 11.9% holds equity in excess of \$10,000.

The results in table 1 shed light on possible reasons for not holding stock. A large proportion of nonstockholders has few other liquid assets. In particular, 43.2% of nonstockholders have liquid assets of less than \$1,000. These consumers, who comprise 31.3% of all consumers, are very likely liquidity-constrained. Thus, in many cases, the failure to hold some wealth in the form of equity is simply due to the absence of any liquid wealth. Yet liquidity constraints are not the only reason for not holding stock: many individuals

Table 1

Distribution of stockholdings and liquid assets for U.S. households in 1984.

Based on 2,998 families in the 1984 Panel Study of Income Dynamics (PSID). Values are in 1984 dollars. The figure in each cell is the percentage of the population with the characteristics of that cell.

Stock value	Liquid assets (excluding stocks)					Total: Frequency #observations
	\$0	\$1-\$999	\$1,000-\$9,999	\$10,000-\$99,999	\$100,000 and up	
\$0	14.0%	17.3	25.0	14.6	1.5	72.4 2169
\$1-\$999	0.2	0.9	2.5	0.7	0.1	4.5 134
\$1,000-\$9,999	0.4	1.0	5.1	4.3	0.4	11.3 338
\$10,000-\$99,999	0.1	0.5	2.8	6.1	0.5	10.1 302
\$100,000 and up	0.0	0.1	0.5	1.0	0.3	1.8 55
Total: Frequency	14.7	19.8	35.9	26.8	2.9	100.0
# observations	440	592	1077	803	86	2998
% owning stock	4.8	12.3	30.6	45.5	47.7	27.6

who have substantial liquid assets also hold no stock. Of those consumers holding other liquid assets in excess of \$100,000, only 47.7% hold equity. For these wealthy consumers, the failure to hold equity is more puzzling. Substantial information costs or noneconomic reasons may be the explanation.

We can learn more about the reasons for not holding stock by examining the relationship between stockholding status and family characteristics. Table 2 shows the probability of being a stockholder for different categories of education and labor earnings averaged over three years.² We find that the fraction of households owning stock increases with average labor income, even holding education constant. For example, among households whose head has a high school degree but no college degree, 14.1% of those in the lowest income quartile own stock, whereas 47.8% of those in the highest income quartile own stock. In addition, more highly educated household heads are more likely to be stockholders, even holding income constant. For example, among households in the third income quartile, 17.5% of those with no high school degree own stock, whereas 51.2% of those with a college degree own stock. (The only exception to the latter finding is that having an advanced degree lowers the likelihood of stock ownership.) These findings are not surprising and are consistent with the presence of fixed information

²For earlier work on the demographics of stockholders, see Crockett and Friend (1963), Blume, Crockett, and Friend (1974), and Blume and Friend (1978).

Table 2

Fraction of U.S. families that hold stock in 1984 by education and average labor income.

Entries are the fraction of families in the corresponding cell that have stockholdings greater than zero. Numbers in parentheses are the total number of families (stockholders plus nonstockholders) in the corresponding cell. Labor income is equal to an average of the years 1981 to 1983.

This table is based on families included in all of the 1982 to 1984 PSID surveys.

Quartiles of average labor income	Education				Total (# observations)
	No high school degree	No college degree	College degree	Advanced degree	
Quartile 1	2.7%	14.1	36.7	28.6	11.1 (433)
Quartile 2	11.2	20.2	42.5	42.3	21.8 (432)
Quartile 3	17.5	28.7	51.2	33.3	31.6 (434)
Quartile 4	48.7	47.8	60.8	50.8	52.2 (433)
Retired	12.9	30.5	58.8	48.1	25.7 (501)
Total (# observations)	12.4 (591)	27.7 (1184)	53.5 (316)	45.1 (142)	28.4 (2233)

costs. Higher-income families are more likely to choose to pay the fixed cost because they have larger portfolios, and the fixed cost is lower for the more educated because information acquisition and processing are less costly.

We conclude this section by noting two potential problems with these data on stock ownership. The first is that some consumers may hold stock through pension funds and yet be called nonstockholders by the PSID. The existence of pension plans, however, does not substantially increase the prevalence of stock ownership. In 1978, 51% of the labor force had no pension plan at all. Moreover, 69% of those in pension plans had defined-benefit rather than defined-contribution plans. Thus, only 16% of the labor force had defined-contribution pension plans.³ Since the residual claimants in defined-benefit plans are the shareholders of the firms rather than the pension recipients, these pension recipients should not be viewed as stockholders. Therefore, the data in tables 1 and 2 do not substantially understate the incidence of stock ownership.

³These figures are reported in Beller (1983). The percentage of the labor force with defined-contribution plans has been trending upward and reached 31% in 1987 (Beller, private conversation). Thus, indirect stock ownership may be more important today than it has been historically. Note that these figures refer to individuals rather than families. It is likely that some of the defined-contribution plan members also hold stock outside of their pension plans and would thus already be counted by us as stockholders. In addition, some individuals in the PSID may have included stocks held in defined-contribution pension plans in the reported market value of stocks.

The second and more serious problem is that the figures from the 1984 PSID probably overstate the incidence of stock ownership throughout the 1970–1984 sample that we analyze below. The New York Stock Exchange (1986) reports that the fraction of the population owning stock almost doubled between 1965 and 1985. Thus, it seems likely that many of our reported stockholders were nonstockholders earlier in the sample. This imperfect separation of stockholders and nonstockholders makes it more difficult to detect differences in the consumption behavior of the two groups.

3. The consumption-based capital asset pricing model: Review

The consumption-based capital asset pricing model – the consumption CAPM – begins with the optimization problem of an individual choosing his portfolio to maximize

$$E_t \int_0^{\infty} e^{-\delta s} U(C_{t+s}) ds, \quad (1)$$

where C is consumption, $U(C)$ is the instantaneous utility function, and δ is the subjective rate of time preference. Between any two points in time, the first-order condition for this problem is

$$E_t \left[\frac{U'(C_{t+s}) e^{-\delta s}}{U'(C_t)} (1 + R_{t,t+s}^i) \right] = 1, \quad (2)$$

where $R_{t,t+s}^i$ is the rate of return on asset i between t and $t+s$. For simplicity, it is often assumed that the utility function takes the isoelastic form

$$U(C) = \frac{C^{1-A}}{1-A},$$

where A is the Arrow–Pratt coefficient of relative risk aversion.

One can use this first-order condition to derive the following relation between consumption and asset returns:

$$\begin{aligned} E(R_t^i) &= A E(GC_t) + \delta - ((1/2)A \cdot (A + 1))\text{var}(GC_t) \\ &\quad + A \text{cov}(R_t^i, GC_t), \end{aligned} \quad (3)$$

where GC_t is the instantaneous rate of growth of consumption and R_t^i is the instantaneous return on asset i . Eq. (3) leads to an equation for the

difference in the expected return on any two assets i and j :

$$E(R_i^i - R_i^j) = A \cdot \text{cov}(R_i^i - R_i^j, GC_i). \quad (4)$$

Grossman and Shiller (1982) show that this equation can be aggregated across individuals under quite general conditions. That is, eq. (4) is valid not only for a single individual, but also for the aggregate consumption of any set of consumers who are at interior solutions with respect to the holding of the relevant assets. If individuals have different coefficients of relative risk aversion, A is a weighted harmonic mean of those coefficients.

Much of the empirical literature on consumption-based asset pricing has examined whether conditions such as (2), (3), and (4) describe the data on aggregate consumption and asset returns. Hansen and Singleton (1983) and Grossman, Melino, and Shiller (1987) report rejections of the overidentifying restrictions implied by these equations. Mankiw and Shapiro (1986) report that the traditional CAPM outperforms the consumption CAPM in explaining mean returns in a cross-section of stocks. Breeden, Gibbons, and Litzenberger (1989) find that the performance of these two models is similar, but reject some important implications of the consumption CAPM. Campbell and Shiller (1988) test and reject the present value relation implied by eq. (2).

Perhaps the most prominent anomaly for the consumption CAPM is the equity premium puzzle. To see what the model implies for the equity premium, consider the case in which asset i is the market portfolio of stocks and asset j is the short-term government bond rate. Eq. (4) then becomes the equation for the equity risk premium. Letting R_i^m and R_i^f denote the return on the market portfolio and the risk-free rate, eq. (4) can be rewritten as

$$E(R_i^m - R_i^f) = A \text{corr}(R_i^m - R_i^f, GC_i) \cdot \sigma(GC_i) \cdot \sigma(R_i^m - R_i^f). \quad (5)$$

One can use aggregate data to estimate the sample moments in eq. (5) and infer the coefficient of relative risk aversion A .⁴

Table 3 presents estimates of A from different estimates of the sample moments. The first row uses the Mehra and Prescott (1985) data, which are annual from 1890 to 1979. In these data, the correlation of the excess return on the Standard and Poors 500 and the growth of consumption is 0.40, the standard deviation of the growth of nondurables and services consumption is 0.036, the standard deviation of the excess return on the market is 0.167, and the average excess return is 0.062. These figures, together with eq. (5), imply that the coefficient of relative risk aversion is 26.3. These numbers are based

⁴This approach of calibrating the first-order condition is followed, for example, by Grossman, Melino, and Shiller (1987) and Mankiw (1986). It differs from the Mehra–Prescott approach of calibrating a general equilibrium model.

Table 3

Calibrating the equity premium: Aggregate data.

GC is the growth of consumption and A is the coefficient of relative risk aversion implied by the corresponding estimates.

	$\rho(GC, r^m - r^f)$	$\sigma(GC)$	$\sigma(r^m - r^f)$	$\text{cov}(GC, r^m - r^f)$	$E(r^m - r^f)$	Implied value of A
1890–1979 ^a (nondurables + services)	0.40	0.036	0.167	0.002349	0.062	26.3 17.5 ^c
1948–1988 ^b (nondurables + services)	0.45	0.014	0.140	0.000898	0.080	89.0
1948–1988 ^b (food only)	0.39	0.022	0.140	0.001201	0.080	66.5
PSID all families	0.26	0.021	0.148	0.000796	0.080 ^d	100.4

^aBased on annual data. All numbers are calculated from Mehra and Prescott (1985) data.

^bBased on 1st quarter to 1st quarter growth rates in National Income and Product Accounts and Ibbotson and Sinquefeld data. r^m is the return on the S&P 500, and r^f is the return on three-month Treasury bills. In each case, returns are calculated as the quarterly average of the monthly, twelve-month log return. Arithmetic (not log) returns are used to calculate the mean excess return. Further details are in the appendix in Mankiw and Zeldes (1990).

^cAdjusted for time aggregation.

^dUses value from 1948 to 1988.

on annual averages, however, and therefore do not necessarily correspond to the instantaneous moments in eq. (3). Grossman, Melino, and Shiller (1987) show that time aggregation biases the estimate of A upward. If consumption and the stock price index are each random walks, then the estimate of A should be multiplied by $2/3$, resulting in an estimate of A equal to 17.5.

The subsequent rows in table 3 present the same calculation using alternative estimates of the relevant moments. The second row uses only postwar data; although the correlation of consumption growth and the excess return is higher, consumption growth is less volatile, raising the implied value of A to 89. [Romer (1989) has provided evidence that prewar National Income Accounts output data are excessively volatile, suggesting that greater weight should be placed on the postwar calculations than on those based on the longer time series.] The third row uses consumption of food only, to provide National Income Accounts results that are most comparable to those based on the PSID data. Overall, the results are similar to those based on nondurables and services: the implied value of A is 66.5. The last row presents calculations based on aggregate food consumption for all families in the PSID sample. The correlation between consumption growth and the excess return is slightly lower, implying a coefficient of relative risk aversion of 100.4.

Most economists view the equity premium as puzzling because such large coefficients of relative risk aversion seem implausible.⁵ To judge the reasonableness of this parameter estimate, it is instructive to consider simple choices under uncertainty. For example, consider what value of X would make an individual indifferent between the following two gambles over consumption:

Gamble 1	\$50,000	with probability 0.5
	\$100,000	with probability 0.5
Gamble 2	X	with probability 1.0

Assuming constant relative risk aversion utility, here is the translation between the choice of X and the risk aversion parameter A .

X	A
70,711	1
63,246	3
58,566	5
53,991	10
51,858	20
51,209	30

Values of X as low as 51,858 seem implausible, suggesting that the level of risk aversion necessary to generate the observed equity premium is too large to be believable.

This application of the consumption CAPM, like most of the empirical literature on this topic, assumes that aggregate consumption is the relevant measure with which to test the model. Yet many consumers hold no stock at all. To see the implications of nonstockholding for the consumption CAPM, consider an economy with two groups of individuals. One group is involved in the stock market and is at an interior solution with respect to the holding of stocks, and the other group holds no stocks at all. The relationship between aggregate consumption and the stock market considered above is no longer valid, because aggregate consumption includes the consumption of both the individuals who satisfy the first-order conditions and those who do not. Eq. (5) does hold for the total consumption of the stockholders, however, because the Grossman–Shiller aggregation theorem applies to this subset of consumers.

⁵Mehra and Prescott (1985) and Weil (1989) point out that an additional part of the puzzle relates to the low historical mean level of the riskless rate of return. This riskless rate puzzle can be viewed as the inability to fit eq. (3) to the aggregate data using the rate on short-term Treasury bills.

To implement empirically this model with two groups of consumers, one needs separate measures of the consumption of stockholders and nonstockholders. Unfortunately, aggregate data of this sort are not directly available. Below we use panel data on households from the PSID to construct the consumption of stockholders and contrast it with the consumption of nonstockholders.

4. Consumption comparisons

To examine the differences between stockholder and nonstockholder consumption, we begin with 17 years of data from the PSID. The consumption questions in the survey ask about the amount spent on food consumed at home and food consumed in restaurants (but not about total consumption expenditures). We deflate each component by its corresponding consumer price index and sum the two components to compute total real food consumption. These data show that stockholding families spend approximately 25% more per capita on food than nonstockholding families (approximately 12% more on food consumed at home and almost 80% more on food away from home), and that approximately 25% of stockholders' food expenditures and 17% of nonstockholders' food expenditures occur away from home.

The survey is administered sometime between late February and April, and the questions ask about consumption around the time of the survey. We interpret the responses as equal to consumption during the first quarter of the year, and time our stock returns and deflators accordingly. Certain consumption questions were not asked in the first and sixth surveys. As a result, growth rates could not be computed for the second, fifth, and sixth years. We are left with thirteen annual observations of growth rates between 1970 and 1984.

Because the question about the value of stocks was asked only in 1984, we categorize consumers as stockholders and nonstockholders throughout the sample on the basis of their 1984 stockholdings. We split the sample into stockholders and nonstockholders in three ways. In split 1, a household is a stockholder if it holds any stock at all. In split 2, a household is a stockholder if it holds at least \$1,000 of stock. In split 3, a household is a stockholder if it holds at least \$10,000 of stock. In each case, the families that do not satisfy the criterion are considered nonstockholders. We sum family consumption across stockholders, nonstockholders, and all families, and then divide by the total number of family members in each group, to obtain aggregate per-capita measures of stockholder, nonstockholder, and total consumption. The growth rates of these consumption measures are presented in the appendix. For a detailed description of the data construction, see Mankiw and Zeldes (1990) and Zeldes (1989).

Table 4

A comparison of the consumption of stockholders and nonstockholders.

GC is the growth of consumption (based on the PSID) and $r^m - r^f$ is the difference between the return on the S & P 500 and the return on three-month Treasury bills. Data are for 1970 to 1984.

	$\text{corr}(GC, r^m - r^f)$	$\sigma(GC)$	$\text{cov}(GC, r^m - r^f)$
<i>Total sample</i>			
	0.260	0.021	0.000796
<i>Split 1: Stockholders have stockholdings > \$0</i>			
Nonstockholders	0.093	0.020	0.000270
Stockholders	0.319	0.031	0.001440
<i>Split 2: Stockholders have stockholdings \geq \$1000</i>			
Nonstockholders	0.047	0.020	0.000137
Stockholders	0.410	0.031	0.001855
<i>Split 3: Stockholders have stockholdings \geq \$10,000</i>			
Nonstockholders	0.102	0.020	0.000305
Stockholders	0.488	0.032	0.002270

The values aggregated across all consumers in the PSID correspond reasonably well to the numbers reported in the National Income and Product Accounts (NIPA). For the 13 observations on growth rates, the correlation between the NIPA and PSID measures is 0.61 for total food consumption, 0.75 for food at home, and 0.51 for food away from home. The average real (1972 dollars) food consumption per capita in the PSID was \$1,323, whereas the corresponding value over the same period in the NIPA data was \$1,692. Note that in the 1984 NIPA data, food consumption is 51% of nondurable consumption and 19% of total consumer spending.

Table 4 presents some sample statistics on total food consumption aggregated for the entire sample and for subsamples of stockholders and nonstockholders. For each group, we present three statistics: the correlation between consumption growth and the excess return on equity, the standard deviation of consumption growth, and the covariance of consumption growth with the excess equity return. The excess equity return is the differential between the return on the S & P 500 and the return on 90-day Treasury bills.

Two findings in table 4 are noteworthy. First, the aggregate consumption of stockholders is more highly correlated with the stock market than is the aggregate consumption of nonstockholders. Second, the consumption of stockholders is more volatile than the consumption of nonstockholders. Both findings imply that the covariance of consumption growth with the excess return – the crucial moment for evaluating the equity premium – is much greater for stockholders. For split 1, this covariance is five times as great for

Table 5

Test for equality of covariances.

$$GC^{\text{stockholders}} - GC^{\text{nonstockholders}} = \alpha + \beta(r^m - r^f).$$

GC is the growth of consumption (based on the PSID) and $r^m - r^f$ is the difference between the return on the S&P 500 and the return on three-month Treasury bills. In split 1, a household is a stockholder if it holds any stock at all. In split 2, a household is a stockholder if it holds at least \$1,000 of stock. In split 3, a household is a stockholder if it holds at least \$10,000 of stock. In each case, the families that do not satisfy the criterion are considered nonstockholders. Data are for 1970 to 1984. Standard errors are in parentheses.

	Split 1	Split 2	Split 3
Constant (α)	0.0028 (0.0065)	0.0019 (0.0057)	0.0026 (0.0054)
$r^m - r^f$ (β)	0.054 (0.046)	0.079 (0.040)	0.090 (0.038)
\bar{R}^2	0.03	0.19	0.28
p -value (one-tailed test)	0.132	0.037	0.019

stockholders as for nonstockholders, and for splits 2 and 3 the stockholders' covariance is over seven times that of nonstockholders.⁶

To test whether these differences between stockholder and nonstockholder covariances are statistically significant, we run the regression:

$$GC^{\text{stockholders}} - GC^{\text{nonstockholders}} = \alpha + \beta(r^m - r^f).$$

The estimate of β equals

$$\text{cov}(GC^{\text{stockholders}} - GC^{\text{nonstockholders}}, r^m - r^f) / \text{var}(r^m - r^f),$$

which in turn equals

$$[\text{cov}(GC^{\text{stockholders}}, r^m - r^f) - \text{cov}(GC^{\text{nonstockholders}}, r^m - r^f)]$$

$$/ \text{var}(r^m - r^f).$$

Hence, β equals zero if and only if the covariances of stockholders and nonstockholders are the same. Table 5 reports the results. For split 1, the estimate of β is positive but significant at only the 13% level. For splits 2 and

⁶These calculations do not make any adjustments for sampling error. Sampling error biases the correlation downward and the standard deviation upward, but does not bias the covariance.

Table 6

Calibrating the equity premium: Stockholders vs nonstockholders.

GC is the growth of consumption (based on the PSID) and $r^m - r^f$ is the difference between the return on the S&P 500 and the return on three-month Treasury bills. A is the coefficient of relative risk aversion implied by the corresponding estimates. Data are for 1970 to 1984.

	$\rho(GC, r^m - r^f)$	$\sigma(GC)$	$\sigma(r^m - r^f)$	$\text{cov}(GC, r^m - r^f)$	$E(r^m - r^f)$	Implied value of A
PSID all families	0.26	0.021	0.148	0.000796	0.080 ^b	100.4
PSID ^a nonstockholders	0.10	0.020	0.148	0.000305	0.080 ^b	261.9
PSID ^a stockholders	0.49	0.032	0.148	0.002270	0.080 ^b	35.2

^aBased on split 3 (a household is a stockholder if it holds at least \$10,000 of stock and a nonstockholder otherwise).

^bUses value from 1948 to 1988.

3, the estimate is again positive but is now significant at conventional significance levels based on the appropriate one-sided test.⁷

The coefficient estimates in table 5 have a simple interpretation. For split 2, the estimate of β is 0.08. This implies that when the excess return on equity is 20% (which is about one standard deviation from the mean), the consumption of stockholders rises 1.6% relative to the consumption of nonstockholders. The difference between stockholders and nonstockholders is thus large economically as well as statistically.

Table 6 uses the moments estimated from the PSID to examine whether the distinction between stockholders and nonstockholders can resolve the equity premium puzzle. Recall from table 3 that using the entire sample for the PSID implies a coefficient of relative risk aversion of 100. Performing the same calculation using the consumption only of stockholders gives an estimate of A of 35. Although 35 is implausibly high, looking at the consumption of stockholders moves us substantially toward resolving the equity premium puzzle.

An important question is whether the distinction between stockholders and nonstockholders can explain the equity premium for a longer period. Al-

⁷The results in tables 4 and 5 are based on the growth rate of the average consumption for each group. An alternative approach would be to examine the average of the growth rates of consumption; this is equivalent to using a geometric rather than an arithmetic mean of consumption. This alternative approach yields results that are qualitatively similar but statistically much less significant. The difference in the approaches is that the growth rate of the arithmetic mean gives greater weight to the growth rate of high-consumption families. The results presented in the text are most analogous to the standard approach in the literature based on aggregate data.

though a century of panel data is not yet available, we can try to answer this question using the evidence in tables 3 and 6. In our short PSID sample, the covariance determining the equity premium is three times as large for stockholder consumption as for aggregate consumption. If this is also true for the 1890 to 1979 sample, a coefficient of relative risk aversion of only 6 would explain the size of the equity premium over that period.

Although these calculations suggest that the distinction between stockholders and nonstockholders can potentially resolve the equity premium puzzle, a final judgment requires better data. First, it would be preferable to have data on a consumption measure broader than food consumption. Looking back at the NIPA data in table 3, however, we find that the covariance of the excess equity return with consumption growth is roughly the same using food consumption as it is using nondurables and services consumption. Thus, we suspect that the use of food consumption does not substantially alter the estimated covariance.

Second, and more important, our ability to separate stockholders and nonstockholders is imperfect. Because the PSID first asked about stockholding only in 1984, we undoubtedly include many nonstockholders in our stockholder category in the early years of our sample. We suspect that a more accurate separation of consumers would yield an even greater covariance of stockholder consumption with equity returns.

5. Conclusions

Only about a quarter of U.S. families own stock. In this paper, we examine the differences between the consumption patterns of stockholders and nonstockholders. Our analysis suffers from the fact that our data measure only food consumption, cover only a relatively short period, and contain substantial measurement error. Nevertheless, our examination suggests that the distinction between stockholders and nonstockholders is important for explaining the empirical failure of the consumption-based capital asset pricing model. We find that the aggregate consumptions of these two groups differ substantially. Furthermore, since we find that stockholder consumption covaries more strongly with excess equity returns than does total consumption, the differences between these two groups help explain the equity premium. The implied coefficient of relative risk aversion based on stockholder consumption is only about one-third of that based on the consumption of all families. Although the resulting coefficient is still too large to be plausible, our work goes in the direction of resolving the equity premium puzzle.

A number of questions remain unanswered. First, why do many wealthy households hold no stock at all? Second, is there a way to approximate the consumption of stockholders using data that are available as a long time series? We leave these questions open for future research.

Appendix

Table A.1

Data on the consumption growth of stockholders, nonstockholders, and all families, 1970–1984.

The consumption growth and the excess return for year t are from the first quarter of year $t - 1$ to the first quarter of year t . The excess return is the differential between the return on the S & P 500 and the three-month Treasury bill rate. In split 1, a household is a stockholder if it holds any stock at all. In split 2, a household is a stockholder if it holds at least \$1,000 of stock. In split 3, a household is a stockholder if it holds at least \$10,000 of stock. In each case, the families that do not satisfy the criterion are considered nonstockholders.

Year	PSID wave	All	Split 1		Split 2		Split 3		Excess return
			Stock	Nonstock	Stock	Nonstock	Stock	Nonstock	
1970	3	0.017	0.004	0.042	0.003	0.041	0.000	0.036	-0.169
1971	4	0.020	0.050	0.004	0.045	0.008	0.044	0.013	0.082
1972	5	-0.010	-0.012	0.004	-0.011	0.002	-0.001	-0.001	0.071
1975	8	-0.033	-0.035	-0.026	-0.046	-0.021	-0.049	-0.025	-0.199
1976	9	0.023	0.033	0.026	0.038	0.025	0.034	0.027	0.213
1977	10	0.026	0.036	0.016	0.037	0.017	0.041	0.020	-0.019
1978	11	-0.008	0.007	-0.007	0.006	-0.006	0.007	-0.005	-0.126
1979	12	-0.006	-0.002	-0.008	-0.002	-0.008	-0.002	-0.007	0.093
1980	13	-0.008	-0.022	0.002	-0.022	0.001	-0.022	-0.002	0.054
1981	14	-0.036	-0.046	-0.032	-0.036	-0.036	-0.030	-0.037	0.125
1982	15	-0.005	-0.008	-0.006	-0.008	-0.006	-0.007	-0.006	0.218
1983	16	0.018	0.034	0.009	0.037	0.010	0.059	0.009	0.215
1984	17	0.016	0.039	0.009	0.035	0.013	0.018	0.019	0.031

References

- Abel, Andrew B., 1989, Asset prices under heterogeneous beliefs: Implications for the equity premium, Rodney L. White working paper (The Wharton School, University of Pennsylvania, Philadelphia, PA).
- Abel, Andrew B., 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review* 80, 38–42.
- Beller, Daniel, 1983, Estimates of participant and financial characteristics of private pension plans (U.S. Department of Labor, Washington, DC).
- Ben-Zvi, Shmuel and Oren Sussman, 1988, The equity premium and the volatility of the rate of return on stocks, Unpublished working paper (MIT, Cambridge, MA).
- Blume, Marshall E., Jean Crockett, and Irwin Friend, 1974, Stockownership in the United States: Characteristics and trends, *Survey of Current Business*, 16–40.
- Blume, Marshall E. and Irwin Friend, 1978, The changing role of the individual investor, A Twentieth Century fund report (Wiley and Sons, New York, NY).
- Breeden, Douglas T., 1979, An intertemporal capital asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.
- Breeden, Douglas T., 1986, Consumption, production, inflation, and interest rates: A synthesis, *Journal of Financial Economics* 16, 3–39.
- Breeden, Douglas T., Michael Gibbons, and Robert Litzenberger, 1989, Empirical tests of the consumption-oriented CAPM, *Journal of Finance* 44, 231–262.
- Campbell, John Y. and Robert J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.

- Cecchetti, Stephen, Pok-sang Lam, and Nelson Mark, 1989, The equity premium and the risk free rate: Matching the moments, Unpublished working paper (Ohio State University, Columbus, OH).
- Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.
- Crockett, Jean and Irwin Friend, 1963, *Characteristics of stock ownership*, *Proceedings of the Business and Economic Statistics Section, American Statistical Association*, 146–179.
- Grossman, Sanford, Angelo Melino, and Robert Shiller, 1987, Estimating the continuous-time consumption-based asset-pricing model, *Journal of Business and Economic Statistics* 5, 315–327.
- Grossman, Sanford and Robert Shiller, 1982, Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information, *Journal of Financial Economics* 10, 195–210.
- Hansen, Lars Peter and Kenneth J. Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of Political Economy* 91, 249–265.
- Institute for Social Research, 1972–86, A panel study of income dynamics, procedures and tape codes, Waves I–XVII (Survey Research Center, University of Michigan, Ann Arbor, MI).
- Kahn, James, 1988, Moral hazard, imperfect risk sharing, and the behavior of asset returns, Unpublished working paper (University of Rochester, Rochester, NY).
- Kandel, Shmuel and Robert F. Stambaugh, 1990, Asset returns, investment horizons, and intertemporal preferences, Rodney L. White working paper (The Wharton School, University of Pennsylvania, Philadelphia, PA).
- Kocherlakota, Narayana R., 1987, The equity premium: A puzzle?, Unpublished working paper (Kellogg Graduate School of Management, Northwestern University, Evanston, IL).
- Lucas, Robert E. Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1429–1445.
- Mankiw, N. Gregory, 1986, The equity premium and the concentration of aggregate shocks, *Journal of Financial Economics* 17, 211–219.
- Mankiw, N. Gregory and Matthew D. Shapiro, 1986, Risk and return: Consumption beta versus market beta, *Review of Economics and Statistics* 68, 452–459.
- Mankiw, N. Gregory and Stephen P. Zeldes, 1990, The consumption of stockholders and non-stockholders, NBER working paper no. 3402 (National Bureau of Economic Research, Cambridge, MA).
- Mehra, Rajnish and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Mehra, Rajnish and Edward C. Prescott, 1988, The equity risk premium: A solution? *Journal of Monetary Economics* 22, 133–136.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- New York Stock Exchange, 1986, *Shareownership 1985* (NYSE, New York, NY).
- Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.
- Romer, Christina D., 1989, The prewar business cycle reconsidered: New estimates of gross national product, 1869–1908, *Journal of Political Economy* 97, 1–37.
- Weil, Philippe, 1989, The equity premium puzzle and the riskfree rate puzzle, *Journal of Monetary Economics* 24, 401–421.
- Zeldes, Stephen P., 1989, Consumption and liquidity constraints: An empirical investigation, *Journal of Political Economy* 97, 305–346.